Facilitating Better Negotiation Solutions using AniMed*

Raz Lin[1], Yehoshua Gev[1] and Sarit Kraus[1,2]
[1] Department of Computer Science
Bar-Ilan University
Ramat-Gan, Israel 52900
{linraz,sarit}@cs.biu.ac.il

[2] Institute for Advanced Computer Studies
University of Maryland
College Park, MD 20742 USA

ABSTRACT

Mediation is an important paradigm for dispute resolution. If done properly, it can lead to "win-win" situations and benefit all parties. Thus, the advantage of designing a proficient automated mediator capable of interacting with people during their negotiations is of great importance. Yet, succeeding in this task is difficult due to the diversity of people and their bounded rationality. To be successful, the mediator must take this into account, and propose solutions deemed relevant, or otherwise, lose the focus and trust of the negotiators. In this paper we present *AniMed*. *AniMed* is 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 more than 100 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 state-of-the-art automated mediator or no mediator was involved.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Multiagent systems*

General Terms

Bilateral Negotiations

Keywords

bilateral negotiations, automated mediation, incomplete information

*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.

The Fourth International Workshop on Agent-based Complex Automated Negotiations (ACAN). Taipei, Taiwan, May 2011

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 [11] and has evolved to assist the negotiating parties in reaching beneficial solutions and increasing their social welfare. In many occasions, 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), and the mediator is usually neutral (unbiased) and objective. This emphasizes the importance of a successful mediation, and thus nowadays it is widespread for dispute resolution ([12], Chapter 2).

Automated mediators, intelligent agents that take the role of an active mediator in the negotiation process, can play an important role in the mediation process between people. 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 the computerized (perhaps "cold") nature of them.

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 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. Another original feature implemented in *AniMed* is its capability to propose 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 with many possible types of negotiators without any restriction to any specific domain. Lastly, as *AniMed* was built on top of GENIUS, a generic negotiation framework [7], it will be available for the public and can be modified and used in numerous domains and settings.

AniMed was evaluated in experiments with more than 100 people who negotiated face-to-face on a neighbor dispute domain 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 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 another state-of-the-art mediator or no mediator were involved. The results also indicate that while people are content with the agreements they achieve without any mediator involved, better agreements can be achieved when *AniMed* is present, which only emphasizes the benefits of using it in human-to-human negotiations. The animated design of *AniMed* and the structure of the experiments was also motivated by findings of Nass and Moon [9], with respect to human-human versus human-computer interaction.

The rest of the paper is organized as follows. Section 2 provides an overview of automated mediators. Section 3 describes the negotiation context, followed by Section 4 that presents the design of *AniMed*, including the user-interface design. Section 5 describes the experimental setting and methodology and reviews the results. Finally, Section 6 provides a summary and discusses future work.

2. RELATED WORK

Few automated mediators are mentioned in the literature. Some are discussed in the context of online dispute resolutions, which are mostly alternative services to litigation. For example, eBay's resolution center[1] tries to facilitate the resolution of conflicts between buyers and sellers.

A number of negotiation support systems are also described in the literature. Family Winner [1], for example, is software that assists divorcees to rationally negotiate their disputes. It does this by advising rational options for trade-offs of assets between opposing parties. However, it is focused on a single domain and cannot be generalized. The HERMES system [6] is a collaborative decision support system that acts as an assistant and advisor by facilitating communication and recommending solutions to members of a decision makers group aiming at reaching a decision. It uses an argumentation framework that provides an issue-based discussion forum [5] whereby users can propose and discuss alternative solutions. Like HERMES, *AniMed* uses the issue-based discussion approach, yet *AniMed* is implemented to support face-to-face negotiation, and not as a collaborative decision support system.

PERSUADER [13] is a computer program that acts as an automated labor mediator in hypothetical negotiations relying on Case-Based Reasoning methods (i.e., logic formulation of the problem). PERSUADER is topic-embedded, and requires data from previous negotiation sessions to reason. In addition, it employs manipulation methods as a mean of manipulating the parties. In contrast, *AniMed* enables the parties to reach satisfactory agreements without the need to resort to manipulations and without the need of a historic database. In addition, unlike PERSUADER, *AniMed* was evaluated with people.

e-Alliance [2] 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.

Olive *et al.* [10] formalize the functionalities an automatic mediator should be able to activate when operating in a multi agent environment. However, as *AniMed* operates in bilateral human environments, some of the defined functionalities are not implemented. For example, the storage of the dialog protocol and its specifications, as well as resolving disputes over the protocol's rules, are irrelevant since *AniMed* uses a pre-defined protocol.

AutoMed [4] is an automated mediator that most resembles our proposed mediator. *AutoMed* monitors the exchange of offers and actively suggests possible solutions, during the negotiation process. It uses a qualitative model for the negotiator's preferences, without past knowledge. The suggestions it makes are basically Pareto-optimal solutions that maximize the social welfare of the negotiating parties. *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 agreements that are deemed more relevant by the negotiating parties.

3. NEGOTIATION PROBLEM DESCRIPTION

We consider the problem of a proficient automated mediator as a key to improving the performance of two human negotiators, who strive to reach an agreement 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. The incomplete information is expressed as uncertainty regarding the utility preferences of the opponent, as explained below.

The negotiation can end when (a) the negotiators reach a full agreement, (b) one of the negotiators opts out, thus forcing the termination of the negotiation with an opt-out outcome, or (c) a predefined deadline, denoted dl, is reached, whereby, a status quo outcome, denoted SQ, is implemented.

[1] `http://resolutioncenter.ebay.com/`

Given a set of issues, $I = \{I_1, I_2, \ldots, I_m\}$ and a set of values $dom(I_j) = \{v_1^j, v_2^j, \ldots, v_{b_j}^j, \perp\}$ for each $I_j \in I$ (since we allow partial solutions to be proposed, $\perp \in dom(I_j)$) and let O be a finite set of discrete values for all issues $(I_1 \times I_2 \times \ldots \times I_m)$. A solution is denoted as a vector $\vec{o} \in O$, where its utility is calculated as a sum of its values. While the utility is known to each negotiator, it is unknown to the mediator. A full order, $\prec$, exists over the values of $dom(I_j)$ and this is the only information shared between the negotiators and the mediator, that basically captures the preference values in the sense of which is "more important than".

It is assumed that the negotiators can take actions during the negotiation process until it terminates. If a partial agreement was accepted it is then implemented. While we did not incorporate time costs in our settings, they can be easily generalized to include time costs which are assigned to the negotiators' utilities. In each period each agent can propose any number of possible agreements, and the other agent can accept the offer, reject it or opt out. Each agent can either propose an agreement which consists of all the issues in the negotiation, or a partial agreement.

To make the problem more realistic the negotiation we consider a setting in which the negotiation is done face-to-face using a video conference and a negotiation system. Thus, the parties negotiated freely and discussed the different issues until they arrived at potential solutions to agree upon.

The negotiation problem also involves incomplete information with regards to the preferences of the opponent. While each side knows its own utilities, the utilities of the other side are private information. Formally, we denote the utility of each side $l \in \{A, B\}$ as u_l, and $u_l : \{O \cup \{SQ\} \cup \{OPT\}\} \rightarrow \mathbb{R}$.

4. 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 [7]. 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 offers deemed relevant by the negotiating parties. To achieve this, *AniMed* utilizes recent offers proposed by the negotiators when generating its own offers, thus centering its offer 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 motivation behind the strategy design of *AniMed* was to generate offers that would maximize the social welfare of

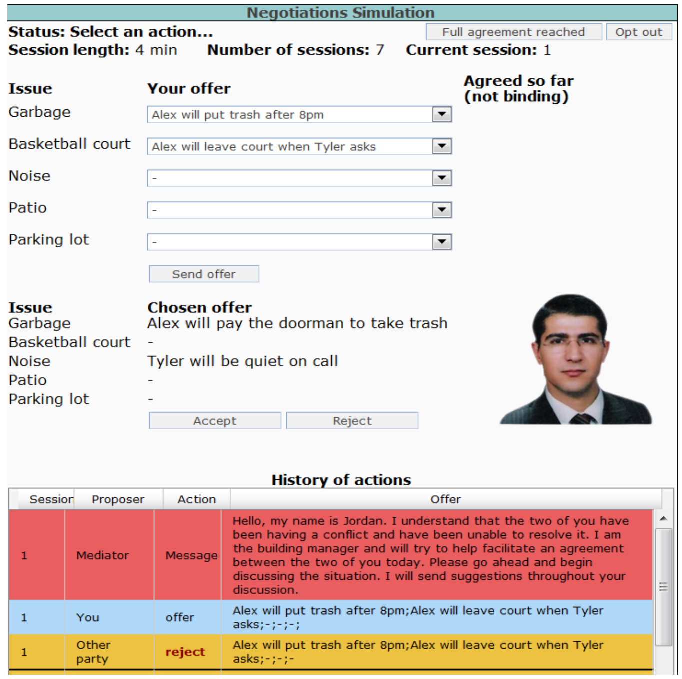

Figure 1: An example of a negotiation snapshot using Genius.

both negotiators. However, this is a difficult task due to conflicting interests between the negotiators. To overcome this, *AniMed* starts proposing only after both negotiators have proposed or accepted an offer 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 offers as irrelevant, *AniMed* does not propose any offer if it is identical to the last offers made by the negotiators.

The strategy used by *AniMed* consists of five steps, which we describe hereafter using an example to better illustrate it. Assume that in a given negotiation domain there are two agents, A and B, and 7 possible solutions. Also assume agent A and B just proposed solutions indexed 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 used by *AniMed* is taken before the negotiation starts. While the utilities of the solutions are private information of each negotiator and unknown to the mediator, *AniMed* uses a linear function, $order(\cdot)$, to obtain an ordinal scale of all solutions. Each issue I_j is given a cost, λ_{I_j}, which is its ranking compared to all other issues, based on the preference relation between the issues. Each issue's value is also ranked based on the preference relation between the possible values of the given issue. Then, the mediator multiplies the costs of issues and values to obtain the linear preference relation. Note that this order may be different from the actual order of the values of the ne-

Solution Idx	$order_A(\vec{o})$	$order_B(\vec{o})$	joint order	diff order
1	7	1	8	6
2	6	2	8	4
3	5	3	8	2
4	4	4	8	0
5	3	6	9	3
6	2	5	7	3
7	1	7	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 offers made by the negotiators are marked in bold.

gotiators. Then, during the game play, *AniMed* chooses its suggestions based on these orderings and on the last offers made by the negotiating parties, which are marked in bold 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 ranking than her last proposal. Thus, *AniMed* removed solutions #1, #2 (higher ordering for agent A) and #7 (higher ordering for agent B) while it kept the four other offers. Then, in its third step, *AniMed* searches for any non-Pareto optimal offers and removes them as well. In our example, one of the solutions was non-Pareto optimal (#6).

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, to allow *AniMed* to propose solutions which are in context with the solutions previous suggested by the negotiating parties, solutions that are more similar to previously suggested solutions, 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 describes the pseudo-code of the algorithm for generating a full proposed solution.

Algorithm 1 Generating A Full Proposed Solution

```
for all o⃗ ∈ O do
    Insert o⃗ to OrderedList_A, OrderedList_B
    Using order_A(o⃗), order_B(o⃗)
end for
if Both sides interacted then
    for all o⃗ ∈ O do
        if order_A(o⃗) < lastOffer(A) then
            remove o⃗ from OrderedList_A
        end if
        if order_B(o⃗) < lastOffer(B) then
            remove o⃗ from OrderedList_B
        end if
    end for
    OffersList =
    Intersect(OrderedList_A, OrderedList_B)
    OffersList = ParetoOffers(OffersList)
    Define jointOrder(o⃗) = order_A(o⃗) + order_B(o⃗)
    Define diffOrder(o⃗) = abs(order_A(o⃗) - order_B(o⃗))
    Sort OffersList
    Using jointOrder(o⃗)
    Then SimilarityToRecentOffers(o⃗)
    Then diffOrder(o⃗)
end if
```

AniMed now has a full solution that it can propose. However, from preliminary experiments, we observed that the dynamics of the negotiation mainly involves partial agreements. Thus, the fifth step in *AniMed*'s strategy is to generate partial solutions that could still benefit the negotiators. *AniMed* incorporates two mechanisms for generating partial solutions. The first is based on joint-value issues. That is, issues with values that are estimated as generating higher utilities for both parties, based on their orderings, can be suggested by the mediator. The second is based on a trade-off between the issues. This is done by calculating the distances between the ordering of given issues, denoted by "diff order" in Table 2. *AniMed* then continues to calculate the difference between the orderings of each pair of issues, that is, $\forall i, j \in I, (order_A(i) - order_B(i)) - (order_A(j) - order_B(j))$. For example, the result for the pair $\langle Issue_1, Issue_5 \rangle$ 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. This evaluation is motivated by our belief that people tend to perceive issues as "important" or "not important", thus they do not use the full possible ranking of solutions. The aforementioned evaluation tries to capture this observation, that is, a pair of issues that would be deemed "important" and "not important" to one of the negotiators, and the opposite to the other party.

Issue Idx	$order_A(issue)$	$order_B(issue)$	diff order
1	1	4	3
2	2	2	0
3	3	3	0
4	4	5	1
5	5	1	-4

Table 2: An example of deciding the trade-off between issues. For each negotiator, the order of importance of an issue is determined by the maximum value she can achieve from that issue.

From those possible partial solutions, *AniMed* tries to make proposals based on trade-offs between issues, or joint-value issues, that were agreed upon or discussed by the negotiators.

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 text message stating that if the negotiators make the suggested trade-off 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

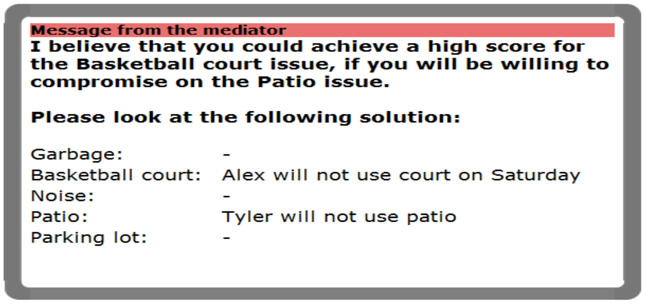

Figure 2: *AniMed* avatar example.

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 (translation to native language was given in the text) speaking avatar (see Figure 2), using a commercial text-to-speech engine, to convey a more "humanized" appearance and a less distant and computerized one.

5. EXPERIMENTS

We matched people on a given domain using the GENIUS environment, an integrated environment for supporting the design of generic automated negotiators [7]. The two negotiators were given a task to negotiate a beneficial agreement. Four different experiments were conducted using the same domain, in order to compare the proficiency of *AniMed*. One experiment involved matching people without any mediator. In another experiment we matched two people with a simple automated mediator, *AutoMed*, suggested by [4]. Chalamish and Kraus demonstrated that this mediator enables the negotiators to achieve more satisfactory agreements in environments where only messages are exchanged.

Finally, we matched people in a setting which included our proposed mediator, once while providing them a facilitation mechanism that provided them a "negotiation calculator" which enabled them to calculate the utilities of each solution at any given time, and once without that mechanism. They played in only one of the experiments in order not to bias the results. We begin by describing the domain which was used in all the experiments and then continue with the experimental methodology and the results.

5.1 The Negotiation Domain

For the negotiation domain we chose a neighbor dispute domain. 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:

1. **Trash**. This issue dictates the solution to the fact that Alex puts its trash every morning on the stairwell, making Tyler angry as it attracts flies and stinks. The possible values are (a) Alex will continue to put trash on the stairwell, (b) Alex will put trash at 5pm, (c) Alex will get a friend to take out the trash, (d) Alex will pay the doorman to take out the trash, or (e) Alex will put trash after 8pm.

2. **Basketball court**. This issue describes how the basketball court will be shared between Alex and Tyler on Saturdays. The possible values are (a) Alex will continue to use the court on Saturdays at any given time, (b) Alex will use court for two hour only, (c) Alex will use court for one hour only, (d) Alex will leave court when Tyler asks him to, or (e) Alex will not use the court on Saturdays.

3. **Noise**. The noise issue tries to resolve the problem of Tyler making noise at nights, disturbing Alex's sleep. The possible values are (a) Tyler will be quiet after 11pm, (b) Tyler will be quiet after 12am, (c) Tyler will be quiet after 1am, (d) Tyler will be quiet upon request by Alex, or (e) Tyler will continue to be loud.

4. **Patio**. This issue describes how the patio will be shared between Alex an Tyler. The possible values are (a) Tyler will not use patio, (b) Tyler will use patio for one hour every other night, (c) Tyler will use patio for one hour every night, (d) Tyler will use patio for two hours every night, or (e) Tyler will continue to use patio whenever he wants to.

5. **Parking lot**. This issue describes the resolution for using the parking lot by Tyler's guests. The possible values are (a) Call the police to give tickets or tow away unauthorized cars, (b) Alex and Tyler will try to recruit other residents to move unauthorized cars, (c) Alex and Tyler will donate money to install "non-parking" signs , (d) complain to the owner about the situation, or (e) Tyler and Alex do nothing.

In this scenario, a total of 3,125 possible solutions existed ($5 \times 5 \times 5 \times 5 \times 5 = 3125$). 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 one of the issues both negotiators received the same utility. On two other issues 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, the negotiators could use tradeoffs between the values of both issues to gain the same utilities[2]. The utility values ranged from 0 to 1,000 for both negotiators, where the Pareto-optimal solution generated a utility of 720 for both.

Each turn in the scenario equated to two minutes of the negotiation, and the negotiation was limited to 28 minutes.

[2]Detailed score functions for the domain can be found at `http://u.cs.biu.ac.il/~linraz/Papers/neighbors-utilities.pdf`

If the negotiators did not reach an agreement 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 agents as the status quo outcome. Each negotiator could also opt-out of the negotiation if it felt that the prospects of reaching an agreement with the opponent were slim and that it was no longer possible to negotiate. The status quo value equaled the opting out value and which was 280 for both negotiators.

5.2 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 negotiators negotiated face-to-face on the different negotiable 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 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 the offer was also done using the user interface. The mediator in turn 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. 104 human subjects participated in the experiments (52 pairs). A total of four sub-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, e.g., whether it was a human or an automated one and what type of strategy it used. The outcome of each negotiation was either they reached a full agreement, they opted out or the deadline was reached.

Prior to the experiments, the subjects were given oral instructions and were shown an instruction video regarding the experiment and the domain. The subjects were instructed to play based on their score functions and to achieve the best possible agreement for them.

5.3 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 administrated questionnaires inquiring about the satisfaction of the negotiators from the outcome and their view on the helpfulness of the automated mediator.

Throughout this section, we also evaluate the significance of the results, compared to the results of *AniMed* without the facilitation mechanism. With respect to the utility values, the significant test was performed by applying the *t-test* on the results, which is a statistical hypothesis test in which the test statistic has a *t-distribution* if the null hypothesis is true. This test requires a normal distribution of the measurements ([3], Chapter 3). Thus, in our analysis it is used to compare the utility values of the negotiation parties in the different experiments conducted, which have continuous values. We applied the *Mann-Whitney U-test* on the results of all other parameters [8]. The *Mann-Whitney U-test* is a non-parametric alternative to the paired t-test for the case of two related samples or repeated measurements on a single sample, suitable for data without normal distribution (as in our case).

		Alex	*Tyler*	*SW*
AniMed w/o facilitation	Average	723	665	1388
	Stdev	92	69	75
AniMed with facilitation	Average	735	669	1404
	Stdev	52	45	59
	p-value	0.35	0.43	0.28
AutoMed	Average	651	590	1241
	Stdev	80	103	145
	p-value	0.022	0.02	0.002
No Mediator	Average	645	595	1240
	Stdev	130	121	150
	p-value	0.045	0.041	0.002

Table 3: Average utility scores, standard deviations, social welfare (SW) and p-values in the different experiments. p-values are compared to experiments involving *AniMed* without the facilitation mechanism.

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. When *AniMed* was involved, the average utility for both negotiators was significantly higher (735 and 669 or 723 and 665 for Alex and Tyler with and without facilitation, respectively) than in any of the cases in which it was not involved (that is, with *AutoMed* – 651 and 590 for Alex and Tyler, respectively – or without any mediator – 645 and 595).

Comparing the sum of utility values of both negotiators when *AniMed* was involved to cases in which it was not involved, also reveals that the sum was significantly higher in cases in which *AniMed* was involved (1241 and 1240 with *AutoMed* or without any mediator, respectively, as compared to 1,388 and 1,404 with *AniMed*). These results were found to be significant (using the 2-sample *t-test*: $p < 0.002$ for both cases). It is interesting to note that though the utility scores were symmetric for both negotiators, on average Alex received higher scores than Tyler. When analyzing the results and videos we can see that there were two issues (noise and garbage) for which non-symmetric agreements were reached. We believe this was due to possible values of the issues and their scores. It seems that the *content* of the value caused subjects to choose them since they seemed reasonable enough, though other values could have generated higher utilities. For example, for the noise issue there were two values – being quiet after 1am, which yielded equal utilities for both Alex and Tyler, or being quiet after 12am, which yielded a higher utility for Alex, yet was preferred by both negotiators. It seems that the country where the negotiations were conducted, being quiet after 12am seemed reasonable enough to be chosen, even though it generated lower utilities for Tyler.

It is also noteworthy that in the two cases in which *AniMed* was involved, once with the facilitation mechanism, and once

	AniMed w/o facilitation	*AniMed* with facilitation	*AutoMed*	*No Mediator*
Average Proposals	7.38	6.46	5.77	6.38

Table 4: Average number of proposals exchanged.

	Outcome Satisfaction	p-value	Mediator's Helpfulness	p-value
AniMed w/o facilitation	3.29		1.08	
AniMed with facilitation	3.31	0.05	2	0.01
AutoMed	3.11	0.18	0.35	0.03
No Mediator	3.19	0.31	N/A	

Table 5: Average satisfaction levels (with 0 being the lowest and 4 the highest) and p-values in the different experiments. p-values are compared to experiments involving *AniMed* without the facilitation mechanism.

without, similar results were revealed without any statistical differences between them, both for the individual utilities of the parties and for the social welfare.

We then analyzed the number of proposals exchanged between the negotiating parties (see Table 4). More proposals were exchanged when *AniMed* was involved than in the other cases, though the differences were not statistically significant. We believe this could be due to the fact that when *AniMed* intervenes in the negotiation process it makes the parties aware of more resolution possibilities, which they later propose themselves.

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 5). The satisfaction levels ranged from 0 (lowest) to 4 (highest). The results significantly demonstrate that the negotiators perceived *AniMed* as more helpful than *AutoMed* ($p < 0.03$). 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 *significantly 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.

6. 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 more than 100 people demonstrated the benefits of *AniMed* compared to another automated mediator and to settings without any mediator. 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 on additional scenarios, including ones with nonlinear utility functions and ones with a larger number of issues. We will also extend *AniMed* to present the negotiators with threats and the ability to enforce solutions and penalties. These features will extend the functionality and the richness of the mediator. Experiments are needed to validate whether these capabilities will still allow the mediator to be successful and whether better agreements can be achieved compared to the current design. Moreover, this kind of manipulative mediator can be used in interesting studies on the impact of different mediation styles on negotiations. In addition, mechanism for obtaining information from the video conference will facilitate the negotiation and will allow the negotiation turn more realistic.

7. REFERENCES

[1] E. Bellucci and J. Zeleznikow. Family winner: A computerised negotiation support system which advises upon australian family law. In *ISDSS2001*, pages 74–85, London, 2001.

[2] M. Bratu, J. M. Andreoli, O. Boissier, and S. Castellani. A software infrastructure for negotiation within inter-organisational alliances. In *AMEC-IV*, pages 161–179, Berlin, Germany, 2002.

[3] S. R. Brown and L. E. Melamed. *Experimental Design and Analysis*. Sage Publications, Inc., CA, USA, 1990.

[4] M. Chalamish and S. Kraus. AutoMed - an automated mediator for multi-issue bilateral negotiations. *JAAMAS*, 2011.

[5] T. F. Gordon, N. Karacapilidis, H. Voss, and A. Zauke. Computer-mediated cooperative spatial planning. In *Decision Support Systems in Urban Planning*, pages 299–309, 1997.

[6] N. Karacapilidis and D. Papadias. Computer supported argumentation and collaborative decision making: The HERMES system. *Information Systems*, 26(4):259–277, 2001.

[7] R. Lin, S. Kraus, D. Tykhonov, K. Hindriks, and C. M. Jonker. Supporting the design of general automated negotiators. In *ACAN*, 2009.

[8] H. B. Mann and D. R. Whitney. On a test of whether one of two random variables is stochastically larger than the other. *Annals of Mathematical Statistics*, 1947.

[9] C. Nass and Y. Moon. Machines and mindlessness: Social responses to computers. *Journal of Social Issues*, 56(1):81–103, 2000.

[10] E. Oliva, P. McBurney, A. Omicini, and M. Viroli. Argumentation and artifacts for negotiation support. *International Journal of Artificial Intelligence*, 4(S10):90–117, 2010.

[11] D. Roebuck. *Ancient Greek Arbitration*. Holo Books The Arbitration Press, 2001.

[12] B. Starkey, M. A. Boyer, and J. Wilkenfeld. *Negotiating a Complex World: An Introduction to International Negotiation*. Rowman & Littlefield Publishers, 2 edition, 2005.

[13] K. P. Sycara. Problem Restructuring in Negotiation. *Management Science*, 37(10):1248–1268, October 1991.