Gender-Sensitive Automated Negotiators

Ron Katz and Sarit Kraus
Department of Computer Science and The Gonda Brain Research Center
Bar-Ilan University, Ramat-Gan 52900, Israel, sarit@cs.biu.ac.il

Abstract

This paper introduces an innovative approach for automated negotiating using the gender of human opponents. Our approach segments the information acquired from previous opponents, stores it in two databases, and models the typical behavior of males and of females. The two models are used in order to match an optimal strategy to each of the two subpopulations. In addition to the basic separation, we propose a learning algorithm which supplies an online indicator for the gender separability-level of the population, which tunes the level of separation the algorithm activates. The algorithm we present can be generally applied in different environments with no need for configuration of parameters. Experiments in 4 different one-shot domains, comparing the performance of the gender based separation approach with a basic approach which is not gender sensitive, revealed higher payoffs of the former in almost all the domains. Moreover, using the proposed learning algorithm further improved the results.

Introduction

As a result of the rapid expansion of automated agents in many negotiation environments, settings in which human and automated agents interact with each other are increasingly prevalent (e.g. (Davidson *et al.* 2000; Katz & Kraus 2006)). These heterogeneous environments differ from conventional environments of pure computerized agents, where most competitors have strong computational power, and act as rational negotiators. In particular, when interacting with humans, the theoretical equilibrium strategy is not necessarily the optimal strategy since human subjects, who are inherently rationally bounded as well as computationally restricted, commonly do not behave according to the perfect equilibrium (Erev & Roth 1998). In addition, simulation of agents interactions typically runs for thousands of trials until satisfactory performance is reached, while we cannot run trials with humans and our agent for such durations. Thus the convergence towards an optimal strategy must be fast.

Nevertheless, interacting with people enjoys the advantage of the ability to project from the behavior of one human opponent onto another. The existence of psychological theories shows that although people behave differently, there are many common patterns in the behavior of most people. Therefore, an agent that interacts with a series of different people, can learn the general pattern of its opponents' population, and develop a suitable strategy. Recently, this '*Generic Opponent Modeling*' attitude has been successfully used in environments, such as poker (Davidson *et al.* 2000), auctions and the Ultimatum Game (UG) (Katz & Kraus 2006).

In this paper, we propose to construct a negotiating agent which performs ***gender-based*** *opponent modeling*, rather than generic modeling. This approach enables negotiating with various human opponents, as does the generic approach, but furthermore it models the opponents more specifically by considering their gender. The main idea is to predict the behavior of opponents by constructing two separate databases of previous opponents' behavior according to their gender. In contrast to the generic methodology, in which one database includes all previous opponents' behavior information, in this methodology each gender group is matched with a specific strategy. As far as we know, this attitude is unique, and has not been implemented neither for academic nor for commercial use.

Despite the intuitiveness of the suggested approach, its efficiency is not guaranteed. A naive approach would suggest to design an agent which holds two **completely separate** databases that consist of the behavior of previous male and female opponents with whom the agent has already competed. However, dividing the data of previous opponents' behavior into sub databases may cause loss of information, since about half of the information is ignored. This is especially critical during the first interactions, when there are not many samples. Thus, the profitability of this *full separator* agent crucially depends on the existence of different behavior patterns for different gender groups. Nevertheless, we believe (and we empirically demonstrate here) that gender affiliation can be a very efficient separator for two main reasons: First, numerous previous studies support the existence of significant gender differences in various negotiation environments (for surveys see: (Rubin & Brown 1975; Walters, Stuhlmacher, & Meyer 1998)). Second, gender identity is prevalently common and accessible in many environments (and usually can be inferred from users' names), and therefore easy to utilize.

In addition, in order to reduce the loss of information involved in the full separator, we propose an algorithm, termed the *gender-sensitive negotiator* (GSN), which uses a combination of the separated gender databases with a generic database which consists of all the previous opponents. In

Copyright © 2007, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

order to determine when to use each database, the algorithm performs online learning during the interactions, which indicates when the gender-based databases can "stand alone" and supply a more accurate evaluation of the population than the generic database. Moreover, the algorithm can indicate online whether the gender partition is a good separator for the given opponents' population, and if not - it automatically performs a generic opponent modeling. In this manner, we are able to reduce the damage that may occur if gender is not a good separator in certain populations and domains. The proposed algorithm can extend any agent which is based on generic opponent modeling.

The applicability of the GSN was examined in 4 different one-shot interaction domains from the set of Cliff-Edge (CE) environments. In CE environments, which include common interactions, such as sealed-bid auctions and the UG, the probability of success decreases monotonically as the expected reward increases. Our agent competes repeatedly in one-shot interactions, each time against a different human opponent. It performs online learning of the population's behavior and does not apply any examples of previous interactions in the environment. In (Katz & Kraus 2006) it was shown that an algorithm named *Deviated Virtual Reinforcement Learning* (DVRL) yields the highest reward in all the examined CE environments. The DVRL algorithm uses generic opponent modeling. In the current paper we convert the DVRL to be gender sensitive, according to the GSN algorithm. This version is shown to achieve a reward significantly higher than the generic agent, in almost all the examined domains. It also has better results than a full separator of males and females, and than a GSN which was given a random partition, rather than the real opponents' gender.

In the next section, we survey related work. In the subsequent section we introduce our gender-sensitive algorithm. Then we describe the CE environments and the DVRL agent. Afterwards we detail the experimental setting in which we compared the performance of the agent with and without gender considerations, and present the results. In the last section we conclude and outline directions for future work.

Related work

The relation between gender and negotiation style or decision making in negotiations has been intensively explored throughout the years. Rubin & Brown and Walters et al. conducted large surveys and meta-analysis of gender differences in negotiations, which cover dozens of previous studies (Rubin & Brown 1975; Walters, Stuhlmacher, & Meyer 1998). Most studies have supported the existence of gender differences in various domains. In UG, for example, which is examined in the current paper, it was found that males have lower acceptance rates than females (Solnick & Schweitzer 1999). Gender differences were also found in computer mediated negotiations, where human and artificial agents interact (Savicki & Kelley 2000). These evidences reinforce the potential of automated agents which are sensitive to their opponents' gender, as discussed above.

The idea of designing negotiating agents which are sensitive to gender or to other demographic affiliation of the opponents is innovative, as mentioned before. However, *agent-based computational economics* studies also consider demographic distribution when designing agents (see (Tesfatsion 2001)). In these studies they develop groups of agents to model the market population, in order to computationally simulate market dynamics. Some of the agents representing the population are constructed according to various profiles consisting of demographic and psychographic characteristics (Billari & Prskawetz 2003). Other agents use demographic information to assist clients in information searching (Krulwich 1997) and in learning (Boff, Reategui, & Viccari 2005). These agents are more similar to ours, since they refer to each client according to her demographic affiliation. However, we could not find any previous work which directly negotiates with human opponents considering their gender or any other demographic parameter.

The proposed approach - GSN

In this section we present a general description of the proposed algorithm for competing with human opponents by segmenting the population to different gender groups. In this study we focus on one-shot interactions, where the modeling of each specific opponent cannot be done. We assume that the basic algorithms which use generic opponent modeling, learn the actions of previous opponents and store their data. In particular, as can be seen in **Algorithm 1**, in order to select their actions, the basic algorithms manage and use an evaluation of the expected utility, *U(i)*, of each offer, *i*, provided it is chosen. The evaluation of *U(i)* is determined according to the results of previous interactions. Actually, each basic algorithm consists of its own UPDATE and SELECT procedures. The UPDATE procedure determines how to update the expected utilities vector, *U*, after observing the results of the latest action (line 6). The SELECT procedure determines how to select the next action (apparently according to both the current expected utility evaluation, and considerations of exploration of the optimal action) (line 4). In the first round when the U vector is empty, picking an action randomly is the best alternative (line 3).

Algorithm 1 A TYPICAL STRUCTURE OF 1-SHOT NEGOTIATION USING GENERIC OPPONENT-MODELING

1: *round*=1
2: **For** each interaction, **Do**
3:    **If** *round*=1, Select an action i randomly
4:    **Else** Select an action i according to a SELECT procedure, based on U
5:    Observe result of the action, calculate reward
6:    Update vector *U* according to an UPDATE procedure
7:    $round = round + 1$

According to our general approach, which is presented in **Algorithm 2**, there are different evaluations of the expected utilities for males and for females. Therefore, we should maintain different vectors of expected utilities U_m and U_f, which are updated according to the regular UPDATE procedure given the data of males and females opponents, respectively (lines 11-13). Nevertheless, as mentioned above, we also want to exploit our knowledge about the behavior of the other gender group, especially in the first stages of the learning, when we do not have many previous samples of the opponents' actions from each gender separately. Therefore, we also hold a generic database U_g (line 10), which is identical to the U-vector in Algorithm 1. The idea is to start by selecting the first actions according to the generic U_g (line 4), and to examine the separability level between males and females. A high separability level is determined by the existence of clear and consistent differentiation between males' and females' optimal strategies. If the groups of females and males are indeed well-separable, i.e. each group is more homogenous than the general population. There should be different optimal strategies against each sub-population. If, for example, in UG female responders are tougher and set a higher acceptance rate, as shown in (Solnick & Schweitzer 1999), the agent should offer them higher amounts than males. Thus, we should observe a consistent difference between the estimated optimal actions of males and females. In line 14 we present a difference function *D*, which can testify to the difference between the current optimal action (offer) against females ($\arg\max_j U_f(j)$) and against males, in UG and in other CE domains. If, for the last 3 interactions, the optimal action for males was always lower (or higher) than that of the females, we should try to use the separated U_m and U_f in order to select the next action (lines 6,7). On the other hand, if the two groups are not well-separable, the optimal actions in U_m and U_f will frequently swap, and thus we would rather make our move based upon U_g (line 8). In other domains, where there is no monotonic dependency between strategies as in CE environments, the difference between optimal strategies can simply be indicated if U_m consistently inspires a certain optimal strategy which is different from the strategy inspired by U_f. The decision to examine the **3** last interactions rather than the minimal 2 interactions, was based on empirical experiments in CE domains that have shown that determining consistency after 3 interactions yields better results. Considering the last 4 or 5 interactions rather than 3 (in line 6), yielded similar results as 3.

Algorithm 2 THE GENERAL PROPOSED APPROACH - GSN

```
1:  round=1
2:  For each interaction, Do
3:    If round=1, Select an action i randomly
4:    Else If round≤3, Select an action i according to the SELECT
          procedure, based on U_g
5:    If round>3,
6:      If (D(round-1)>0 and D(round-2)>0 and D(round-3)>0)
        or (D(round-1)<0 and D(round-2)<0 and D(round-3)<0),
7:        If current opponent is Female, Select an action i according
          to the SELECT procedure, based on U_f, Else based on U_m
8:      Else Select an action i according to the SELECT procedure,
        based on U_g
9:    Observe result of the action, calculate reward
10:   Update vector U_g according to the UPDATE procedure
11:   If current opponent is Female
12:     Update vector U_f according to the UPDATE procedure
13:   Else Update vector U_m according to the UPDATE procedure
14:   D(round)=arg max_j U_f(j) − arg max_j U_m(j)
15:   round=round+1
```

The Cliff-Edge environments

Cliff-Edge (CE) environments are characterized by the conflict between the desire to maximize profits while preventing the entire deal from falling through. Consider, for example, a proposer in the Ultimatum Game who needs to decide how to divide an amount of money with his opponent (Guth, Schmittberger, & Schwarz 1982): Decreasing the share offered to the opponent increases the profits the proposer accrues, so long as the offer exceeds the opponent's acceptance threshold. A slightly greedier proposal causes the proposer to lose the whole deal. Similarly, a bidder in a sealed-bid first-price auction (e.g. (Ockenfels & Selten 2005)) attempts to bid an amount that is only slightly higher than those put forward by opponent players. This situation is somewhat similar to that of a person standing on the edge of a cliff, trying to see the panoramic view. The closer he approaches the cliff's edge, the better the view. However, one step too many causes the viewer to fall off the cliff. Hence, interactions and games of this type are referred to as Cliff-Edge interactions.

The general pattern of one-shot CE interactions considers a competitor required to choose an offer *i*, being an integer $0 \leq i \leq N$, where *N* is the maximum optional choice. Then a positive reward, *r*, corresponding to the offer, *i*, is determined, depending on whether the offer passed a certain threshold, t, set by the opponent. Specifically:

- In the sealed-bid first-price auction, an amount, *N*, is auctioned.[1] The bidder is required to place a bid *i*, being an integer $0 \leq i \leq N$, which will enable the bidder to gain a reward, *r*, if it exceeds the highest bid, t, made by all other bidders in the current auction. If $t \leq i, r = N - i$ (the amount gained less the bid amount), otherwise r=0. In the all-pay version, where all the bidders must pay their bids, even if they have not won the auction (Krishna & Morgan 1997): if $t \leq i, r = N - i$ otherwise $r = -i$. The bidder is never informed of the size of the bid offered by an opponent.
- In the UG, a proposer needs to divide an amount (*N*) with an opponent by offering the latter an integer amount, $i, 0 \leq i \leq N$. The reward the proposer, *r* will gain is determined according to the opponent's acceptance threshold, t. If $t \leq i, r = N - i$, otherwise *r*=0.

  In this paper we also consider a variant of the classic UG, proposed by Guth and Huck (1997). The only difference in this setting is that the responder always receives his own allocation, and he decides only whether the proposer will receive his share. The reward calculation for the proposer remains exactly the same as in the classic UG.

In this paper, we consider environments with a large set of decision options, and set *N* = 100.

To demonstrate the challenge facing a competitor in CE environments, let *R(i)* be the reward corresponding to a successful offer *i* and let *P(i)* be the probability of the offer, *i*, succeeding (i.e. the probability that the offer will be higher than the other bids, in the case of an auction, or will be accepted by the responder, in the case of the UG). Obviously, there is a trade-off between *R* and *P*: choosing an offer *i* which increases the expected reward, *R(i)*, decreases the probability of success, *P(i)*, and vice versa.

In a recent work where CE environments were considered it was found that the Deviated Virtual Reinforcement Learning (DVRL) algorithm achieved the best performance (Katz & Kraus 2006). The algorithm is presented in **Algorithm**

[1] In order to avoid considerations of value estimations (Ockenfels & Selten 2005) the item auctioned is an amount of money.

3, and is briefly described here.[2] DVRL is based on the basic principle of Reinforcement Learning (RL), according to which an action is selected on the basis of its expected utility, U-value, that each offer would yield if chosen. The U-value of the chosen offer is reinforced after a successful interaction, while, after an unsuccessful interaction, it is decreased. One problem, however, in applying basic RL to CE environments is its disregarding of the fact that in CE the probability of an offer is gradually influenced by the size of the offer. Thus, a reasonable approach for the U-vector update procedure in CE environments is *Virtual Learning* (VL) (Vreind 1997). According to the VL principle, the proposer in the UG, for example, treats all offers higher than an accepted offer as successful (virtual) offers, not withstanding that they were not actually proposed. Similarly, it considers all offers lower than a rejected offer as having been (virtually) unsuccessfully proposed. The rationale behind this principle is that the higher the amount proposed to the opponent, the higher the probability of the proposal being accepted. However, while VL proceeds towards less risky offers after unsuccessful interactions, it performs no exploration of offers which are greedier than the current optimal offer. This is a deficiency it shares in common with the basic RL.

In contrast, DVRL deviates from the strict rationale underlying the VL principle, and extends the range of offers updated after each interaction. Thus, after successfully offering amount *i*, DVRL would increase the U-values of all the offers higher than the actual offer *i*, as well as $\lfloor \frac{i}{round+1} \rfloor$ offers **below** the actual offer, as described in line 8 of **Algorithm 3**. Respectively, we would reduce the U-values of all the offers lower than this new threshold (line 9). Similarly, after an unsuccessful interaction, we would reduce the U-values of all the offers above the actual offer *i*, as well as a few offers **above** *i*, up to $\lfloor \frac{N-i}{round+1} \rfloor$ (line 12). Respectively, all the offers above this new threshold would be reinforced (line 13). This is the innovative UPDATE procedure of DVRL (lines 6-13). In the SELECT procedure we simply choose the offer with the maximal current U-value, in a greedy manner (line 4). The usage of a greedy algorithm in learning is quite unique, and is possible thanks to the special UPDATE procedure, as explained below.

The main challenge of an on-line learning algorithm is to efficiently balance between the need for exploration of new options, and the will to exploit current information in order to maximize payoffs. The DVRL distorts observed information in a manner which actually outlines a direction of the optimal solution searching, rather than the random trial-and-error approach that underlies conventional methods, such as RL. A DVRL agent that offered, for example, 80% of the amount *N* to its UG opponent in the first interaction, and its offer was accepted, would offer 40% in the next interaction. The agent continues to decrease its offer until it is rejected. During the learning process the evaluated model comes closer to the real distribution of the opponents population, and therefore the deviation extents are gradually decreased (by the positioning of *round + 1* in the denominators in lines 8 and 12).

In **Algorithm 4** we present the GSN version of the DVRL algorithm, which was used in the experiments of this paper.

[2]The version presented here is an improved version of DVRL which eliminates the usage of configurable parameters.

Algorithm 3 THE GENERIC DVRL ALGORITHM

Notation: s(j) denotes the corresponding reward for a successful offer j and f(j) is the corresponding reward for offer j when it fails.

1: *round*=1 For j=0 to N, **Do** U(j)=0
2: For each interaction, **Do**
3: **If** round=1, Select an action i randomly
4: **Else** offer i=$\arg\max_j U(j)$
5: Observing opponent's move, calculate reward
6: **If** offer i has succeeded **Then**
7: For j=0 to N, **Do**
8: **If** j $\geq$(i - $\lfloor \frac{i}{round+1} \rfloor$) $U(j) = \frac{U(j)(round-1)+s(j)}{round}$
9: **Else** $U(j) = \frac{U(j)(round-1)+f(j)}{round}$
10: **If** offer *i* has failed **Then**
11: For j=0 to N, **Do**
12: **If** j <(i + $\lfloor \frac{N-i}{round+1} \rfloor$) $U(j) = \frac{U(j)(round-1)+f(j)}{round}$
13: **Else** $U(j) = \frac{U(j)(round-1)+s(j)}{round}$
14: *round*=*round*+1

If the estimated optimal offers are steady (as calculated in line 13), the offer with the highest expected utility for the relevant gender group is selected as the next offer (lines 7-8), or else the generic optimal offer is chosen (line 9). Afterwards, the results of each interaction update the U_m or the U_f vector, according to gender of the opponent (line 12), and the U_g vector (line 14). In our experiments we also considered a ***full-separator*** algorithm which does not use U_g. This algorithm always chooses the next offer according to U_m and U_f, which means that it always executes lines 8-9, and ignores lines 4-7 and 10 in the SELECT procedure, and lines 13-14 in the UPDATE procedure.

Algorithm 4 THE DVRL ALGORITHM - GSN

1: $round$=1,$round_m$=1, $round_f$=1
2: For j=0 to N, **Do** $U_g(j)$=0, $U_m(j)$=0, $U_f(j)$=0
3: For each interaction, **Do**
4: **If** round=1, Select an action i randomly
5: **Else If** round$\leq$3, offer i=$\arg\max_j U_g(j)$
6: **If** round>3,
7: **If** (D(round-1)>0 and D(round-2)>0 and D(round-3)>0) or (D(round-1)<0 and D(round-2)<0 and D(round-3)<0),
8: **If** current opponent is Female, offer i=$\arg\max_j U_f(j)$
9: **Else** offer i=$\arg\max_j U_m(j)$
10: **Else** offer i=$\arg\max_j U_g(j)$
11: Observing opponent's move, calculate reward
12: **If** current opponent is Female execute lines 6-14 in Algorithm 3, when U is changed to U_f and round to $round_f$ **Else** execute lines 6-14 in Algorithm 3, when U is changed to U_m and round to $round_m$
13: D(round)=$\arg\max_j U_f(j) - \arg\max_j U_m(j)$
14: Execute lines 6-14 in Algorithm 3, when U is changed to U_g

Experimental design and analysis

In order to evaluate the performance of the proposed GSN, and to compare it with a generic, non gender sensitive approach, we experimentally examined the agents interaction with human opponents in 4 different CE environments, as follows:

1. Sealed-bid 2-bidders auction for 100 New Israeli Shekels (NIS, where 1 U.S. \$ $\approx$ 4.5 NIS), i.e. N=100.
2. All-pay sealed-bid auction for 100 NIS, where all the bidders must pay their bids (Krishna & Morgan 1997). In this version risk taking considerations are added to the bidder's decision.
3. UG, where the players had to divide 100 NIS.
4. A variant of UG in which the responder determines only whether the proposer receives his share. The responder himself always receives his allocation (Guth & Huck 1997). In this version the responder's decision is influenced by fairness and vindictiveness considerations, rather than profitability as in the original UG version. Here also N=100.

The examination of different environments which activate different personality characteristics of human opponents, strengthen the generality of the results.

In our experiments we surveyed 49 students (25 males and 24 females) participating in the UG and in the auction, and 69 students (30 females and 39 males) participating in the UG variant and in the all-pay auction. The participants were students at Bar Ilan University, from various faculties, aged 20-28, who were not experts in negotiation strategies nor in economic theories directly relevant to the experiment (e.g. game theory, decision theory).

In the auctions the participants were required to propose a bid, which could be any integer from 0 to 100 NIS. The winner gained a virtual 100 NIS. In the ultimatum games, the full amount to be shared was N=100 NIS, as well.

After extracting the bids from people in the auctions, as well as their minimal acceptance thresholds in the ultimatum games, we constructed sets of opponents' reactions for each environment. At this stage we examined the performance of four algorithms; a GSN with gender-based segmentation, a GSN with random segmentation, a generic non-GSN, and a full separator, which were all run serially against the sets of opponents' bids or thresholds. Thus, each agent had one interaction with each of the human opponents in each of the four environments, without knowing in advance the number of interactions. Since there is importance to the order of the opponents, we constructed 100 random permutations of the human decisions series, for each environment, and compared the average payoffs of the different algorithms for each permutation.

Results

Table 1 presents the average payoffs of the **GSN** algorithm and of the gender-based **full-separator** in the 4 CE domains, and compares them to the baseline performance of the **Generic** algorithm (which is not sensitive to gender) and to the GSN which was given **random** segmentation of the population, rather than their real gender. In all the environments except for the UG the payoff of the GSN algorithm and the full-separator were higher than the Generic's. Non-parametric Wilcoxon tests indicated significance for all the pairwise differences ($p < 0.001$). In the UG the Generic was significantly better then the Separator, while no significant difference was revealed between Generic and GSN. In comparison to the random segmentation, the gender-based GSN always yielded significantly higher payoffs. Thus, it can be concluded that the gender parameter was always a meaningful separator, and therefore segmenting the population by gender was worthwhile. Even more interesting, in the auctions and in the UG variant the random GSN was significantly better than the *Generic*. It appears that in all the domains except for UG, the opponents' population is very distributed, and therefore constructing two (or even more) different types of opponents is more worthwhile than only one generic model, as in the *Generic*. This claim is true even without gender considerations, although constructing 2 models based on gender is always better than a random partition, as shown in table 1. In UG, where the diverseness among the population is not high (since all acceptance rates are between 0-50), one model is adequate and thus the generic method performs well.

Domain	Generic	Random	Full Separator	GSN
Auction	20.27	21.35	21.45	**22.17**
All-pay auction	3.46	3.85	4.16	**4.31**
UG	**47.08**	46.84	46.88	47.01
UG variant	43.39	45.47	46.31	**46.8**

Table 1: Average payoff of various algorithms against human opponents in 4 domains

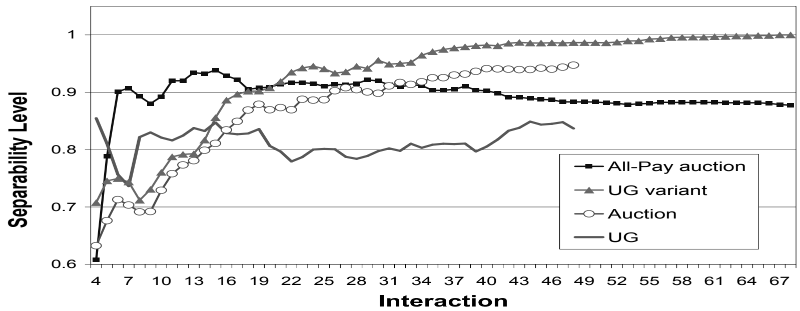

Figure 1: The separability-level values during the interactions in the 4 CE domains

In the rest of this section we discuss the separability-level between males and females in the 4 domains examined here. As mentioned above, the GSN algorithm indicates how well males and females in a given population are separable, according to the consistency of the relationship between the optimal offer in U_f and in U_m (line 6 in Algorithm 4). Thus, we can indicate the separability-level by examining whether the condition in line 6 is true or false. A true value indicates a high separability level. In order to calculate the separability-level estimation at each stage of the learning, we counted the occurrences when the condition became true at each interaction in 100 runs of different permutations. In **Figure 1** we present the probability that the condition was true during the interaction in the 4 CE domains. It begins from the 4th interaction, since the calculation is based on the results of the 3 preceding interactions. It can be seen in the y-axis that the probability is higher than 0.5 during all the interactions. In all the domains except for the UG, there is a general increase in the separability-level during the interactions. This indicates that higher number of samples enabling a clearer separation of the two gender clusters.

The left column of **Table 2** explicitly shows the final separability-level observed in each of the 4 domains (which can be seen in figure 1). There is a clear correlation between

Domain	final separability-level	offer for females	offer for males
Auction	0.95	66.66	75.38
All-pay auction	0.88	75.25	79.32
UG	0.84	50.09	46.77
UG variant	1.0	54.38	68.36

Table 2: The separability-level value and the offers during the final interaction in each domain

the separability-level and the relative success of the GSN algorithm in Table 1. In the auction and in the UG variant, where the payoffs of GSN are noticeably higher than the *Generic*'s, the values are very close to 1. In the UG and in the all-pay auction, on the other hand, the advantage of the GSN is much lower, and so are the final separability-level values. As explained above, in the UG domain the GSN mechanism prevented the agent from totally separating the population as was done by the *Separator*.

The two right columns of **Table 2** present the final offers made by GSN for males and for females. In our experiments, males offered higher bids in both auctions, and therefore the counter-offers of the agent were higher for males. In UG, on the other hand, as demonstrated in previous studies as well (Solnick & Schweitzer 1999), males were less tough and had lower acceptance rates than females, and therefore the agent offered males smaller amounts. In the UG variant the opposite occurred. Males' acceptance rate was much higher than females', and therefore our agent offered them 69.38 vs. only 54.71 offered to females. It is noteworthy that in the UG variant domain we found a significant difference between males' and females' acceptance rates (t-test, $p < 0.05$), which caused the large gap between our agent's offers. However, in other domains we also succeeded in improving the payoffs although the differences between males' and females' acceptance rates or bids did not meet statistical significance. Again, it can be noted that in the auction and in the UG variant the differences between males and females' acceptance rates / bids are much higher than in the UG and in the all-pay auction domains, in correlation with the corresponding separability levels.

Conclusion and future work

In the paper we examined the possibility of automated negotiation with human opponents considering their gender. A simple separation between males' and females' samples was found to be worthwhile in almost all 4 domains examined here. With additional use of a learning algorithm proposed here, we succeeded to further improve the average payoff. The algorithm also supplies an indicator for the gender separability-level of the population, which calibrates on-line the level of separation between the males' and the females' databases.

In the future, we would like to extend the GSN for repeated interactions, in which several negotiation rounds can be conducted against each opponent. When competing repeatedly against the same opponent, a specific modeling of the current opponent must be done, in addition to the generic modeling of her gender group. Moreover, in contrast to the one-shot version, a current move may influence the future behavior of the opponent, a fact that must be taken into account. Therefore, we intend to design an agent that develops several optional models of typical opponents from each gender group, and matches the appropriate model to each opponent with which it interacts.

The approach presented here may also be applied to other demographic characteristics. The relation between human behavior in negotiations and various demographic groups affiliation such as profession (Carter & Irons 1991), culture, age, social background and even religion (Rubin & Brown 1975) has been intensively explored. Some of these works reveal considerable differences in negotiating habits of different demographic groups, which should be taken into account when designing agents for the growing human-machine negotiation world.

Acknowledgements: This work is supported in part by NSF #1211-04. Kraus is also affiliated with UMIACS.

References

Billari, F., and Prskawetz, A. 2003. *Agent-Based Computational Demography: Using Simulation to Improve Our Understanding of Demographic Behaviour*. Heidelberg: Physica-Verlag.

Boff, E.; Reategui, E.; and Viccari, R. 2005. An affective based virtual character for learning environments. In *AI-ED Workshop on motivation and affect in educational software*.

Carter, J., and Irons, M. 1991. Are economists different, and if so, why? *Journal of Economic Perspectives* 5:171–177.

Davidson, A.; Billings, D.; Schaeffer, J.; and Szafron, D. 2000. Improved opponent modeling in poker. In *IC-AI'2000*, 1467–1473.

Erev, I., and Roth, A. 1998. Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibrium. *Amer. Econ. Rev.* 88(4):848–881.

Guth, W., and Huck, S. 1997. From ultimatum bargaining to dictatorship - an experimental study of four games varying in veto power. *Metroeconomica* 48(3):262–279.

Guth, W.; Schmittberger, R.; and Schwarz, B. 1982. An experimental analysis of ultimatum bargaining. *J. Econ. Behav. Organ.* 3:367–388.

Katz, R., and Kraus, S. 2006. Efficient agents for cliff-edge environments with a large set of decision options. In *AAMAS'06*.

Krishna, V., and Morgan, J. 1997. An analysis of the war of attrition and the all-pay auction. *J. Econ. Theory* 72:343–362.

Krulwich, B. 1997. Lifestyle finder: Intelligent user profiling using large-scale demographic data. *AI Magazine* 18(2):37–45.

Ockenfels, A., and Selten, R. 2005. Impulse balance equilibrium and feedback in first price auctions. *Games Econ. Behav.* 51(1):155–170.

Rubin, J. Z., and Brown, B. R. 1975. Bargainers as individuals. In *The social psychology of bargaining and negotiation*. New York: Academic Press. 157–196.

Savicki, V., and Kelley, M. 2000. Computer mediated communication: Gender and group communication. *CyberPsychology and Behavior* 3(5):817–826.

Solnick, S., and Schweitzer, M. 1999. The influence of physical attractiveness and gender on ultimatum game decisions. *Organ. Behav. Hum. Dec.* 79(3):199–215.

Tesfatsion, L. 2001. Special issue of agent-based computational economics. *J. econ. dyn. control* 25:281–654.

Vreind, N. 1997. Will reasoning improve learning? *Econ. Lett.* 55(1):9–18.

Walters, A.; Stuhlmacher, A.; and Meyer, L. 1998. Gender and negotiator competitiveness: A meta-analysis. *Organizational Behavior and Human Decision Processes* 76(1):1–29.