2017 Seventh International Conference on Affective Computing and Intelligent Interaction (ACII)

Towards Modeling Agent Negotiators by Analyzing Human Negotiation Behavior

Yuyu Xu, Pedro Sequeira and Stacy Marsella

College of Computer and Information Science, Northeastern University
360 Huntington Ave., Boston, MA 02115

Email: {yuyuxu, psequeira, marsella}@ccs.neu.edu

Abstract—**Negotiation is a fundamental aspect of social interaction. Our research aims to contribute towards the creation of artificial agent negotiators that can be used for training purposes to improve human negotiation skills. To achieve that, we address the challenge of identifying differences in human negotiation styles and relating those differences to individuals' personality traits. In particular, we follow a data-driven approach by collecting data on how people negotiate against an agent using a fixed-response strategy during a task involving the partition of a set of items. We then use different machine learning techniques to: 1) analyze the relationship between negotiation styles and personality traits; 2) characterize changes in the human negotiation behavior during the game; 3) discover human behavior patterns in response to different offers by the agent player. Our analyses show how different personality traits lead to distinct behaviors during the negotiation. In turn, this data will allow us to build agent negotiators that have a rich behavioral repertoire and are able to adapt to human negotiation trainees, thus fostering more interesting learning experiences.**

1. Introduction

We live in a world of bounded resources, whether it is constraints on our time, actions, material things, or space. Due to these constraints, negotiation plays a central role in both our professional and personal lives. Not surprisingly, this has lead to a wide range of research on computationally modeling negotiation. Some of this research seeks to build automated agents that can negotiate for us [1, 2], while other concerns building systems that can train people in the social skills that will make them better negotiators [3, 4].

Our ultimate goal in this work is to use data-driven approaches to build agent-based facsimiles of human negotiators that can be used to train people by having them negotiate with the agents. In particular, we are interested in modeling sequential negotiation tasks, in which people negotiate over multiple items during a series of offers and counter-offers. For the purpose of training, the agents will have to exhibit a wide range of negotiation behaviors, *i.e.*, have different goals, exhibit different negotiation styles, etc. On the other hand, they should be able to predict and adapt to the human negotiators' behavior. One of the key challenges here is that there are considerable differences in *how* people negotiate [5], which will ultimately lead to distinct negotiation outcomes. These differences in negotiation styles are influenced by personality traits [6, 7] as well as cultural factors [8, 9]. As such, in order to effectively train people we first need to computationally characterize these differences within the agent models. To better capture such differences, we are exploring a data-driven approach, collecting data on how people negotiate and using that data to build models of human negotiators.

In the work reported here, we take an important step towards this modeling goal by classifying patterns of negotiation styles from human negotiation data. In particular, we address two key challenges:

- A *data collection* challenge: In order to capture diverse and interesting negotiation behavioral patterns, data needs to be collected from a population with a wide range of personality traits.
- A *data analysis* challenge: Human negotiation poses a hard modeling problem. It is a sequential process where each step depends on both individual internal factors—such as one's culture, social predispositions or preferences—as well as external ones—such as the negotiation partners' prior offers during the negotiation task.

Addressing these challenges requires the creation of data collection methodologies and systematic ways to extract meaningful information from human negotiation data. To address the data collection challenge we used the Amazon Mechanical Turk (AMT) crowdsourcing platform to acquire large amounts of data from people with diverse cultural and social backgrounds. Each participant was asked to play an online turn-based negotiation game based on the *Auction Wars* game [10]. At each turn, they had to choose a partition over a set of items being negotiated, including 3 records, 2 lamps and 1 painting, against a simple autonomous agent using a strategy based on [7]. Fig. 1 presents the interface of the negotiation game. We collected demographic information from the human players as well as personality trait instruments related to negotiation style, including *Social Value Orientation* (SVO) [11], corresponding to how much weight a person attaches to the welfare of others in relation to their own, and *Machiavellianism* (Mach) [12] which relates

978-1-5386-0563-9/17/$31.00 ©2017 IEEE

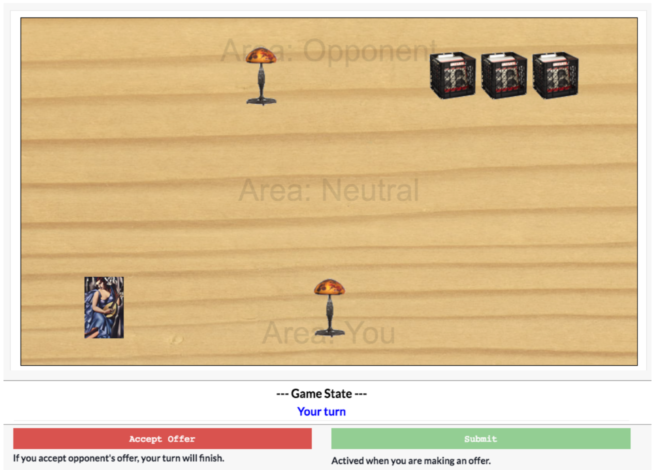

Figure 1: The negotiation game interface.

to a person's tendency to be deceptive and manipulative. Regarding the data analysis challenge, we decided to break it down into three levels, each analyzing a particular aspect of the negotiation data, namely:

1) *Personality traits analysis*, where we analyze the relationship between a set of features summarizing each person's negotiation style and their personality traits;
2) *Negotiation dynamics analysis*, where we investigate how the personality traits influence changes in item partitioning throughout the negotiation;
3) *Response tendency analysis*, where we explore human player's responses to the agent's counter-offer to the human's previous offer, *i.e.*, action-reaction triplets, in order to identify the probability of specific responses in the human players in the context of prior offers.

The results of our analyses showed correlations between individual traits and specific characteristics of their negotiation behavior. For example, the *personality traits analysis* revealed that a person's SVO is correlated with the value of their first offer, maximum offer made and persistence in repeating an offer. The *response tendency analysis* revealed that individualists tended to persist in making very unfair offers far more frequently than prosocials after a sequence of very unfair offers and counter-offers. Also, low-Mach players more frequently persisted in re-stating a fair offer after an agent's unfair counter-offer while high-Mach individuals responded by exploring different kinds of offers in the negotiation space. In the context of our research, we will use the resulting models to build agent negotiators capable of improving human negotiation skills, both by providing a range of opponents with different negotiation styles and goals based on personality traits characteristics, as well as by predicting human negotiation behavior given their traits in order to tailor the learning experience during training.[1]

1. A deeper discussion on how we can use the results in this paper to build agent negotiators for training will be provided in Sec. 5.

2. Related Work

In the context of human-computer negotiation, there have been many efforts in modeling negotiation agents using various methods such as equilibrium strategies [13] and opponent modeling [14–16].

Other works have adopted a data-driven method for learning various aspects of human negotiation behavior [3, 17–20]. They normally apply various machine learning (ML) algorithms on data collected for a certain negotiation task to predict human player's preference or action, using game-specific information as input features. In particular, Peled et al. [17] collected data through the colored-trails game framework under the incomplete information game setting. They then applied classical ML algorithms to predict what the human player reveals at certain steps. This predictive model was used to provide a probabilistic estimation of the player's action within their MERN agent. Gal and Pfeffer [18] collected data on human bidding task and presented several novel learning algorithms to predict the chosen bid given features of all the candidate proposals. Haim et al. [20] applied ML to predict how people reach and fulfill agreements in negotiation task given their cultural background, where cultural information was incorporated as extra features along with game specific features. As in these works, we could also learn predictive models of human negotiation behavior by incorporating the personality traits as additional features. However, in this paper we are mainly interested in characterizing the negotiation behavior explicitly for each trait.

There is also work in psychology and social science that investigates human behavior patterns and their relationship to the Mach and SVO traits, *e.g.*, [6, 7]. Nonetheless, not many computational models of human negotiation incorporate these two traits. Exceptions include the work of Nazari et al. [21], which presented a recognition model for Machiavellianism derived from multimodal data, *e.g.*, gesture and language. Our work both seeks to predict the Mach/SVO score from negotiation behavioral features, as well as learn different sequential models according to different trait types.

3. Collecting Human Negotiation Data

In this section we address the challenge of collecting human negotiation data from a population with a high range of individual traits. In that regard, we used AMT and developed a multi-stage, interactive web-application hosted on Amazon's AWS EC2 server. AMT provides a good solution to gather a high quantity of inexpensive data. We implemented a negotiation game where a player (crowd participant) negotiates against an autonomous agent using a fixed strategy similar to the one used in [7], which can be described as a strategy between cooperativeness and competitiveness. In this manner, we focus our analysis on the human players' negotiation behavior. Notably, recall that at this stage of our research we are interested in investigating which personality traits contribute more to human negotiation behavior variability, and what are the negotiation dynamics stemming from such individual characteristics.

3.1. Interactive Application Flow and Design

3.1.1. Pre-Game Data Collection. The participants started by taking a HIT (Human Intelligence Task) where they had to fill in their demographic information. In particular, we collected data on participants' gender, age, education and ethnicity. After that, we collected data on individual traits. As mentioned earlier, the two traits that are closely related to negotiation tasks are Mach and SVO. For the former, we used the widely adopted *MACH-IV score* personality test by Christie and Geis [12]. For the latter, we used a continuous slider SVO test proposed by Murphy et al. [11] involving several resource allocation tasks.

After completing the personality questionnaires, participants were led to an introduction to the negotiation game itself. We provided a description of the game, explaining to the participant that there is another party looking at the items so that the player is unaware that he/she is playing against an automated negotiator. We indicated that a record is worth the highest value, as much as all the other items combined, that a lamp is worth a medium value, while the painting has a residual value when compared to the other items. Notably, we did not provide quantitive information about the items' value and also did not refer to how the other party valued the items. We instructed the participant to negotiate to the best of his/her capabilities in order to get the best outcome. We stated that the player would receive $40¢$ for participating in the study, and that upon good performance he/she would receive an extra $10¢$ in order to incentivize their engagement in the negotiation task—in reality, in the end of the game all AMT participants received $50¢$ independently of their performance. By vaguely stating the relative value of the items and what constitutes a good negotiation performance, we aimed at creating a complex multi-issue, multi-level negotiation that would not be reduced to a simple calculus of "maximize the total value". Moreover, we added a quiz after the introduction to ensure that the game setup was well understood—in particular, we prevented participants from proceeding to the game without correctly answering a set of questions about the relative values of the items.

3.1.2. Negotiation Game. We designed a minimal interface for the negotiation game, which is depicted in Fig. 1. As we can see, the game was played over a virtual 2D table, where items were placed either in the bottom part, corresponding to human player's side, or at the top, corresponding to the autonomous agent's (the opponent) side. We let the human player make the first move to avoid the agent of establishing the initial offer, which could frame the rest of the negotiation [22]. During each turn, a player could either accept the offer previously made by its opponent or make a counter-offer by dragging the items to the side they feel adequate—this task was automated for the agent's turns. An offer was considered to be valid only if it corresponded to a full partition over the items, in which case the "Submit" button would become active. The game ended when both parties accepted some particular partition offered by one of the players, or a total number of 9 negotiation rounds (offer and counter-offer) was reached. Once the game was over, the participants got a reward code and left the web-application server to return to the AMT page.

TABLE 1: Summary of collected data before/after filtering.

Summary	**Before filtering**	**After filtering**
Num. data points	404	292
Mean age	36.50 ± 11.33	37.67 ± 11.85
Gender (#male / #female / #missing)	216 / 182 / 6	160 / 130 / 2
Mean Mach score	58.28 ± 4.99	58.20 ± 4.93
SVO score (#individualist / #prosocial)	165 / 239	107 / 185
Mean total session duration in min.	11.70 ± 5.15	12.20 ± 4.86
Mean game sequence length	5.49 ± 2.22	6.31 ± 1.41
Mean participant's last offer payoff	6.62 ± 3.04	7.76 ± 1.86

3.2. Autonomous Agent Player

As mentioned earlier, we adopted a fixed agent negotiation behavior, which is similar to the one in [7]. In particular, the agent followed a fixed sequence of offers where for the first six turns, it went from trying to keep every possible item until ending up in a slightly-unfair value split in the 6$^{\text{th}}$ turn.[2] From there on, the agent would make that last offer until the game ended. In the context of this work, we consider a partition to be fair if one player gets either 2 records, or 1 record, 2 lamps and 1 painting, while the other gets the remainder of the items. In each turn, we sampled the agent's response time in seconds from a uniform distribution of interval $[4, 12]$ to make the agent's behavior more realistic.

4. Analyzing the Negotiation Data

In this section we analyze the data according to the different levels, namely the personality traits, the negotiation dynamics and the response tendency analyses. Before going through all collected data and analyses some notations are introduced here. Let $O_i = (o_i^1, ..., o_i^t, ..., o_i^{T_i})$ denote the offer sequence of human player i, where each o_i represents a particular set of items proposed by the player and $T_i \in [3, 9]$ represents the length of the negotiation game in number of rounds. Given the number of the different available items, there are $4 \times 3 \times 2 = 24$ possible item partitions in total. During a negotiation game, we refer to each o_i^t as the *negotiation state* for a player in round t. In addition, for convenience of analysis we assigned a concrete value to each item according to the game description given to the participants. Specifically, we attributed a value of 4.1 to a record, 2.0 to a lamp and 0.1 to the painting. In this manner, each offer o_i^t made by a player has an associated payoff, denoted by $r(o_i^t) \in \mathbb{R}$, obtained by summing the values of each item in that offer.

4.1. Summary of Collected Data

In total, we collected data from 404 participants, where each interaction session corresponding to a different datapoint. A summarization of the data can found in Table 1.[3]

2. The fixed sequence of offers we used for the agent player was: $[3, 2, 1]$, $[2, 2, 1]$, $[2, 2, 0]$, $[2, 1, 1]$, $[2, 1, 0]$, $[1, 2, 1]$, where each offer is in the form $[\#records, \#lamps, \#paintings]$.

3. The collected data can be found at the project's code webpage: `https://github.com/yuyuxu/auction-war-solo.git`.

To make sure our analysis was performed on *reliable* data and given our goal of creating *good* negotiator agents, we filtered out games where we judged the players as not being engaged in the task according to the following criteria:

- *The game was played for too few rounds*. Specifically, we filtered out games with only one or two negotiation rounds—this meant that we could not observe interesting behaviors since the agent tries to keep almost every item to itself in the first two rounds;
- *The time spent on the entire session by the participant was too short*. By examining the mean session duration, detailed in Table 1, we chose a threshold of 5 minutes;
- *The human player had an extremely bad final performance*. According to the values in Table 1, we defined a bad performance as corresponding to an offer of $[0, 1, 1]$ or worse, *i.e.*, to a payoff less than or equal to 3.1—this could either mean that the player did not understand the items' relative values or did not take the task seriously.

After this filtering step, a total of 292 data-points remained for further analysis. As can be seen from Table 1, the distribution of both Mach and SVO score remains relatively the same before and after filtering.[4]

4.2. Personality Traits Analysis

In this first level of analysis we are interested in evaluating the relationship between a set of features summarizing each participant's negotiation style and their personality traits. By doing such analysis, we can answer the question of whether one can make predictions regarding human players' traits based only on their demographics and generic information about the negotiation game. This is an important phase towards our goal of creating intelligent negotiator agents that can predict their negotiation partner's personality and social predispositions during a negotiation task.

In the context of our analysis, we considered 3 categories of features regarding each data point, as listed in Table 2. The first group of features reflects how the human player behaved in the game. Regarding the *negotiation style* features, recall that a game only ends if both parties accept the offer, meaning that the agent may accept the human player's offer, but in turn the human may propose a different counter-offer. As such, the purpose behind feature f_7 is to analyze whether the human player was "testing the limit of the opponent". Features f_8 and f_9 count the number of decreases or increases in the payoff between two consecutive offers, respectively. The intuition is to evaluate whether the player was searching the negotiation space for the best possible outcome. Feature f_{10} counts the number of offers in a sequence that had an equal payoff, thus measuring the persistence of the human's negotiation style. Feature f_{11}'s purpose is to assess whether the negotiation behavior of the human was consistent or erratic. All features in this group are normalized according to the sequence length. Finally, we used the demographics of each individual as features to

TABLE 2: Features used in the trait analysis. We dropped the individual's reference for ease of notation.

Category	ID	Description
General features on payoff	f_1	Negotiation length, T
	f_2	Payoff of first offer, $r(o^1)$
	f_3	Payoff of last offer, $r(o^T)$
	f_4	Average offer payoff, $\bar{r}(o^T) = \sum_t r(o^t)/T$
	f_5	Minimum offer made, $\min_t r(o^t)$
	f_6	Maximum offer made, $\max_t r(o^t)$
Negotiation style (strategy) features	f_7	Number of offers until agent's acceptance
	f_8	Number of decreases in payoff
	f_9	Number of increases in payoff
	f_{10}	Number of repeated sequential offers
	f_{11}	Avg. absolute difference between two offers
Metadata features	f_{12}	Total interaction session time (seconds)
	f_{13}	Gender: $0: female, 1: male$
	f_{14}	Age
	f_{15}	Education (bachelors): $0: below, 1: equal, 2: above$

discover whether metadata information could interact with behavioral features and together correlate with the traits.

4.2.1. Methods. Let $y_{m_i} \in \mathbb{R}$ denote the score for of the Mach trait of an individual and $y_{s_i} \in \{0, 1\}$ represent the category for their SVO trait, where 0 means individualistic and 1 refers to a prosocial player. We first used statistical analysis to see how these features help explain the personality traits by applying a linear regression model on y_{m_i} and a generalized linear regression model with binomial distribution for response variable y_{s_i}. Both models contained linear and pairwise interaction terms, and the features were standardized before applying the models.[5] We then ran a step-wise model selection on the two models to select the best feature subset. We also explored how to use these features to best predict trait characteristics by trying ML models with different expressive power and report the performance using 10-fold cross-validation.

4.2.2. SVO Trait Results. By applying generalized linear regression statistical analysis on the features and performing step-wise model selection, we got a 78-terms subset. The significant main effects ($p < 0.05$) include f_2 (payoff of first offer), f_6 (maximum offer made), f_{10} (persistence in repeating the same offer). For 60 interaction effects terms, 34 involve features from the first category, 27 terms involve features from the second category and 38 terms involve features from the third category. To perform a prediction task on features f_1 to f_{15}, we applied logistic regression with L2 regularizer,[6] which gives an accuracy of $Acc = 64.49\%$ and gradient boosting[7] (as a highly non-linear model), which generates $Acc = 63.10\%$. Here the L2 regularized logistic regression worked better, suggesting that gradient boosting could be overfitting severely on this dataset for this task.

4.2.3. Mach Trait Results. After applying the linear regression model and step-wise model selection, we got

4. We filled in missing the values for *age* using its mean value and for *gender* by sampling from a Bernoulli distribution.

5. Matlab Statistics and Machine Learning Toolbox: www.mathworks.com/products/statistics.html.

6. Sklearn: scikit-learn.org.

7. Pyramid: github.com/cheng-li/pyramid.

$R^2 = 0.466$, adjusted $R^2 = 0.326$ and a subset of 62 terms. The significant main effects ($p < 0.05$) correspond to features f_3, f_7, f_8, f_{11}, f_{12}, f_{13} and f_{15}. For 44 interaction effects terms, 26 of such terms involve features from the first category, 23 terms involve features from the second category and 26 terms involve features from the third category. This suggests that the features we extracted all have impact and correlate to the Mach trait through interaction. In order to perform a prediction task on features f_1 to f_{15}, we applied one linear model—linear regression—and one highly non-linear model—gradient boosting. Linear regression gives $RMSE = 5.11$ while gradient boosting gives $RMSE = 4.87$. The non-linear model provided better results in comparison to linear model given the complex interactions between features.

4.3. Negotiation Dynamics Analysis

At this level we are interested in analyzing how human players change their negotiation behaviors from moment to moment, *i.e.*, evaluate the dynamics across negotiation steps. In particular, we want to gain insights on how the trait characteristics of players influence such dynamics. In turn, this information may be used to model the behavior of our negotiator agents. In light of the fixed-response pattern of the agent, we start by considering here just the dynamics of the human offers, while in the following section we also consider the agent's moves.

4.3.1. Methods. To simplify the analysis of dynamics, we divided the possible offers into categories of fairness. In particular, we considered 5 types of offer states with different payoff ranges, denoted by the set $S = \{EG, G, F, U, EU\}$:

- *Extremely generous* (EG) offer, where $r(.) \in [0.1, 6.1)$;
- *Generous* (G) offer, where $r(.) \in [6.1, 8.1]$;
- *Fair* (F) offer, where $r(.) = 8.2$;
- *Unfair* (U) offer, where $r(.) \in [8.3, 10.3]$;
- *Extremely unfair* (EU) offer, where $r(.) \in (10.3, 16.4]$;

We arranged the types so that each range had a similar number of possible item partition configurations.

We examined the behavioral changes by looking at how the human players transitioned between different states and analyzed whether the transitions differ across individual traits. In essence, this is equivalent to building a Markov chain given the relative frequency of transitions provided from the negotiation data. The differences in state transitions across the traits were tested for statistical significance by applying Fisher's exact test on a 2×2 contingency table where we counted the number of times a player transitioned between two particular offer states, *i.e.*, $s_i \rightarrow s_j$, versus the number of times it transitioned to other states, *i.e.*, $s_i \rightarrow s_k, s_k \neq s_j$, where $s_i, s_k, s_j \in S$, according to two partitions over the trait scores. While the SVO already contained only two possible values, either individualist, denoted by S_{ind} or prosocial, denoted by S_{pro}, for the Mach score we partitioned the data according to the mean trait value. In particular, players with a Mach score below average were considered to have a *low-Mach*, denoted by M_{low} while above-average individuals were put in the *high-Mach* group, denoted by M_{high}.

4.3.2. SVO Trait Results. The resulting Markov chain transitions regarding the SVO partition of the data are shown in Table 3. Each entry in the table provides the frequency of the transition from the corresponding row offer state to the column state relative to all transitions originating from the row state. Significant differences between the two SVO groups ($p < 0.05$) are presented in bold. As we can see, individualistic players tend to transition to extremely unfair states after proposing either extremely generous or extremely unfair offers more often than prosocial players.

4.3.3. Mach Trait Results. The resulting Markov chain for the Mach score is shown in Table 4. We see that, on average, low-Mach players tend to remain in a fair offer state (transition from fair state to fair state) more often than high-Mach players, which tend to transition to either unfair or generous offers. We also see that low-Mach players transition to generous offers after being in a unfair state less frequently than high-Mach players. These results thus seem to suggest a difference in preference for negotiation states—low-Mach individuals tend to remain more in a "fair partition" negotiation state compared to high-Mach players.

4.3.4. Trait Prediction. For the sake of completeness, we also used the resulting Markov chain models to perform prediction of the trait types. The prediction accuracy using the Markov chain for the Mach trait using 10-fold cross-validation was 55.70%, while for the SVO trait it was 56.31%. The results are worse than those reported in the previous analysis, suggesting that the human negotiation dynamics by itself is not a good predictor of their personality traits, although specific negotiation preferences arise when considering different traits.

4.4. Response Tendency Analysis

In the previous section we considered the transitions in the negotiation behavior of individuals independently of the offers of the agent player. In this section, we explore how trait characteristics of the human players influence the offers they make in response to their opponents' counter-offers. This analysis is extremely important in the context of our research, as it helps predicting how human players with certain personality traits respond to specific proposals by the other party during a negotiation.

4.4.1. Methods. We counted the occurrence of all possible sub-sequences in the form (*human's offer – agent's offer – human's counter-offer*), what we refer to as *action-reaction triplets*. Just like with the *negotiation dynamics analysis*, we considered the 5 different categories to qualify the offers made by the human player. Given that the agent's offers fall only into 3 out of the 5 possible categories due to its fixed strategy, this originates a total of $5 \times 3 \times 5 = 75$ possible triplets. Also similar to the previous analysis, here we want to see whether different action-reaction triplets patterns arise

TABLE 3: Markov chain for individualistic vs. prosocial SVO trait. Entries represent relative frequency of transitions from the row offer state into the column offer state. Bold results indicate significant differences between trait values.

	Extr. generous		**Generous**		**Fair**		**Unfair**		**Extr. unfair**		*Total count*	
	S_{ind}	S_{pro}	S_{ind}	S_{pro}	S_{ind}	S_{pro}	S_{ind}	S_{pro}	S_{ind}	S_{pro}	S_{ind}	S_{pro}
Extr. generous	0.154	0.200	0.154	0.250	0.230	0.350	0.077	0.150	**0.385**	**0.050**	10	21
Generous	0.073	0.090	0.293	0.451	0.293	0.262	0.268	0.180	0.073	0.016	67	158
Fair	0.008	0.015	0.205	0.186	0.425	0.485	0.244	0.227	0.118	0.087	174	338
Unfair	0.000	0.003	0.085	0.128	0.316	0.348	0.415	0.363	0.184	0.157	211	306
Extr. unfair	0.022	0.005	0.038	0.020	0.168	0.245	0.384	0.470	**0.389**	**0.260**	138	133
Total count	13	20	41	122	127	264	234	350	185	200	600	956

TABLE 4: Markov chain for low-Mach vs. high-Mach trait. Entries represent relative frequency of transitions from the row offer state into the column offer state. Bold results indicate significant differences between trait values.

	Extr. generous		**Generous**		**Fair**		**Unfair**		**Extr. unfair**		*Total count*	
	M_{low}	M_{high}	M_{low}	M_{high}	M_{low}	M_{high}	M_{low}	M_{high}	M_{low}	M_{high}	M_{low}	M_{high}
Extr. generous	0.083	0.238	0.250	0.190	0.417	0.238	0.083	0.143	0.167	0.190	11	20
Generous	0.083	0.088	0.347	0.462	0.264	0.275	0.236	0.176	0.069	0.000	103	122
Fair	0.004	0.026	0.171	0.225	**0.513**	**0.390**	0.213	0.265	0.100	0.093	311	201
Unfair	0.000	0.004	**0.086**	**0.148**	0.340	0.329	0.403	0.354	0.170	0.165	311	206
Extr. unfair	0.014	0.012	0.018	0.042	0.209	0.206	0.464	0.381	0.295	0.358	155	116
Total count	12	21	72	91	240	151	347	237	220	165	891	665

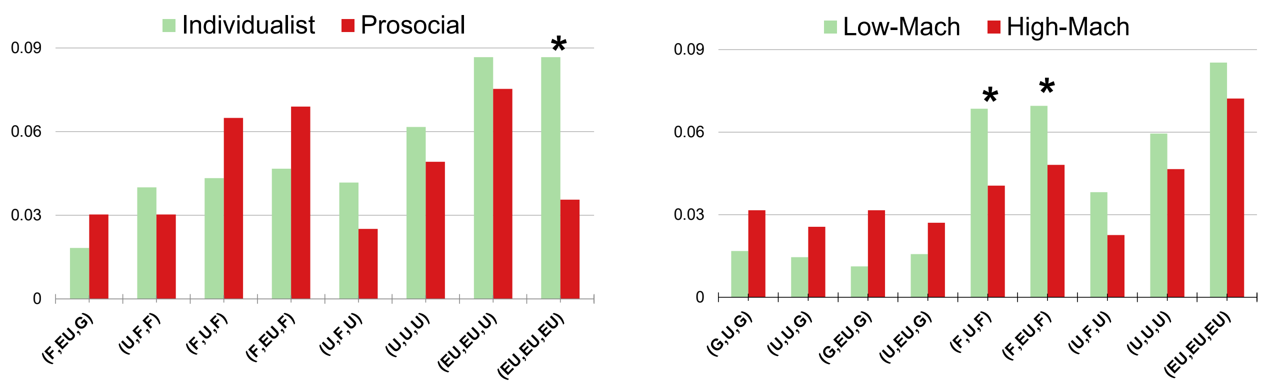

Figure 2: Action-reaction triplet probabilities for individualistic vs. prosocial trait (left) and high-Mach vs. low-Mach trait (right). Only triplets with prob. $> 1\%$ and pronounced trait differences are shown. Highlighted results are marked with $*$.

across the individual traits. However, the difference is that here we utilize the distribution over possible triplets to aid in identifying differences in behavior since data is more sparse given there are more possible combinations.

4.4.2. SVO Trait Results. The distribution of triplets for individualistic and prosocial traits is shown in Fig. 2 (left). When partitioning our data according to the SVO score, we can see that in cases where the human player starts with an extremely unfair offer and the agent responds with an extremely unfair offer to the player, individualists tend to stick with the extremely unfair offer more often than prosocial players. In particular, the proportion of the individualist group is twice that of the prosocial group. Such difference was indeed found significant according to a Fisher's exact test, resulting in $p = 28 \times 10^{-4}$.

4.4.3. Mach Trait Results. The distribution of triplets for low-Mach and high-Mach is shown in Fig. 2 (right). Some interesting results stem from this analysis. In cases where the human player starts with a fair offer and the agent responds with an unfair or extremely unfair offer, low-Mach players have a significantly-higher (Fisher's exact test resulted in $p = 5 \times 10^{-3}$) chance of going back to a fair offer when compared to high-Mach individuals.

5. Discussion

In this paper we presented a detailed analysis of data collected from humans playing a turn-based negotiation game against an automated opponent using an online web-application. We used standard personality instruments to assess their trait characteristics. Our analysis of the players' negotiation dynamics shows that people classified as low-Mach tend to remain more in a fair partition negotiation state when compared to high-Mach players. In addition, individualists tend to persist with unfair offers more often than prosocial players. Finally, human negotiation response patterns show that low-Mach players tend to target fair states while high-Mach players seem to be exploring more of the negotiation state-space. Regarding the SVO trait, prosocial players tend to respond less negatively to extremely unfair offers from the agent than individualistic players. All these results suggest that both Mach and SVO are important factors in characterizing the goal of human negotiators, *e.g.*, aiming for fairness vs. aiming for a resolution, and how they respond to specific offers by the other party, *e.g.*, retaliate vs. comply.

The current analyses reveal a range of behaviors that suggest several ways forward to building agent models based on trait characteristics. For example, one might simply develop a probabilistic policy based on the above patterns, *i.e.*, by using the action-reaction triplets as states, we can build Markov chain models according to different traits where the agent's decisions are based not only on its previous state but also on the opponent's state. Potentially, one could try to infer goals from the behavior and traits, that in turn could be validated by asking the crowd participants what their goals were regarding the negotiation.

In the context of negotiation training, we could use as opponents for the learner a range of agent types bootstrapped from the human behaviors identified in the current work. Additionally, the gathered knowledge of human behavior provides valuable information for the creation of interesting learning opportunities tailored to the learner. Specifically, it can be used by agents with a Theory of Mind capacity to reason about their human opponent, *e.g.*, to craft rewards for a goal-based agent. Furthermore, it may help in identifying negotiation situations in which we predict humans may behave poorly, *e.g.*, when people with certain traits respond negatively to an opponent's unfair offer and because of that end up having a low final payoff. In turn, such situations can be used by the trainer agent to improve the trainee's negotiation skills.

Our next steps are to relax the constraints used here to simplify data collection and analysis. Going beyond using a fixed opponent strategy, we have begun collecting human-human interactions. Critically, the methodologies developed in this work, *negotiation dynamics analysis* and *response tendency analysis*, will still hold for human-human negotiation data. In addition, we are incorporating communications between the players using a fixed repertoire of messages. Another aspect we would like to explore is to go beyond the constraint of full partition offers to investigate the exchange of partial offers.

Acknowledgments

This work is supported by the Air Force Office of Scientific Research, under grant FA9550-14-1-0364. The content does not necessarily reflect the position or the policy of any Government, and no official endorsement should be inferred.

References

[1] S. Fatima, S. Kraus, and M. Wooldridge, "The negotiation game," *IEEE Intelligent Systems*, vol. 29, no. 5, pp. 57–61, Sept 2014.

[2] K. Fujita, R. Aydoğan, T. Baarslag, T. Ito, and C. Jonker, "The fifth automated negotiating agents competition (anac 2014)," in *Recent Advances in Agent-based Complex Automated Negotiation*. Springer, 2016, pp. 211–224.

[3] J. Kim, R. W. Hill, P. Durlach, H. C. Lane, E. Forbell, M. Core, S. C. Marsella, D. V. Pynadath, and J. Hart, "BiLAT: A Game-Based Environment for Practicing Negotiation in a Cultural Context," *International Journal of Artificial Intelligence in Education*, vol. 19, no. Issue on Ill-Defined Domains, pp. 289–308, 2009.

[4] J. Gratch, D. DeVault, and G. Lucas, "The benefits of virtual humans for teaching negotiation," in *International Conf. on Intelligent Virtual Agents*. Springer International Publishing, 2016, pp. 283–294.

[5] L. Thompson, "Negotiation behavior and outcomes: Empirical evidence and theoretical issues." *Psychological bulletin*, vol. 108, no. 3, p. 515, 1990.

[6] W. R. Fry, "The effect of dyad machiavellianism and visual access on integrative bargaining outcomes," *Personality and Social Psychology Bulletin*, vol. 11, no. 1, pp. 51–62, 1985.

[7] C. K. De Dreu and P. A. Van Lange, "The impact of social value orientations on negotiator cognition and behavior," *Personality and Social Psychology Bulletin*, vol. 21, no. 11, pp. 1178–1188, 1995.

[8] M. Gelfand and N. Dyer, "A cultural perspective on negotiation: Progress, pitfalls, and prospects," *Applied Psychology*, vol. 49, no. 1, pp. 62–99, 2000.

[9] M. J. Gelfand, M. Higgins, L. H. Nishii, J. L. Raver, A. Dominguez, F. Murakami, S. Yamaguchi, and M. Toyama, "Culture and egocentric perceptions of fairness in conflict and negotiation." *Journal of Applied Psychology*, vol. 87, no. 5, p. 833, 2002.

[10] D. DeVault, J. Mell, and J. Gratch, "Toward natural turn-taking in a virtual human negotiation agent," in *AAAI Spring Symposium on Turn-taking and Coordination in Human-Machine Interaction. AAAI Press, Stanford, CA*, 2015.

[11] R. O. Murphy, K. A. Ackermann, and M. Handgraaf, "Measuring social value orientation," *Judgment and Decision Making*, vol. 6, no. 8, pp. 771–781, 2011.

[12] R. Christie and F. L. Geis, *Studies in machiavellianism*. Academic Press, 2013.

[13] R. Lin, S. Kraus, J. Wilkenfeld, and J. Barry, "Negotiating with bounded rational agents in environments with incomplete information using an automated agent," *Artificial Intelligence*, vol. 172, no. 6-7, pp. 823–851, 2008.

[14] K. Hindriks and D. Tykhonov, "Opponent modelling in automated multi-issue negotiation using bayesian learning," in *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 1*. International Foundation for Autonomous Agents and Multiagent Systems, 2008, pp. 331–338.

[15] Y. Oshrat, R. Lin, and S. Kraus, "Facing the challenge of human-agent negotiations via effective general opponent modeling," in *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*. International Foundation for Autonomous Agents and Multiagent Systems, 2009, pp. 377–384.

[16] T. Baarslag, M. Hendrikx, K. Hindriks, and C. Jonker, "A survey of opponent modeling techniques in automated negotiation," in *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, ser. AAMAS '16. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems, 2016, pp. 575–576. [Online]. Available: http://dl.acm.org/citation.cfm?id=2936924.2937008

[17] N. Peled, K. Gal, and S. Kraus, "Learning to reveal information in repeated human-computer negotiation," in *Human-Agent Interaction Design and Models (HAIDM) workshop*, 2012.

[18] Y. Gal and A. Pfeffer, "Predicting people's bidding behavior in negotiation," in *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems*, ser. AAMAS '06. New York, NY, USA: ACM, 2006, pp. 370–376. [Online]. Available: http://doi.acm.org/10.1145/1160633.1160699

[19] G. Haim, Y. Gal, S. Kraus, and Y. Blumberg, "Learning human negotiation behavior across cultures," in *HuCom10-Second International Working Conference on Human Factors and Computational Models in Negotiation*, 2010.

[20] G. Haim, D. Nisim, and M. Tsatkin, *Human-Computer Agent Negotiation Using Cross Culture Reliability Models*. Cham: Springer International Publishing, 2017, pp. 118–133. [Online]. Available: http://dx.doi.org/10.1007/978-3-319-57285-7_8

[21] Z. Nazari, G. Lucas, and J. Gratch, "Multimodal approach for automatic recognition of machiavellianism," in *Affective Computing and Intelligent Interaction (ACII), 2015 International Conference on*. IEEE, 2015, pp. 215–221.

[22] A. D. Galinsky and T. Mussweiler, "First offers as anchors: The role of perspective-taking and negotiator focus." *Journal of Personality and Social Psychology*, vol. 81, no. 4, pp. 657–669, 2001.