

Negotiation as a Challenge Problem for Virtual Humans

Jonathan Gratch,¹ David DeVault,¹ Gale Lucas¹ and Stacy Marsella²

¹University of Southern California and ²Northeastern University
{gratch,devault,lucas}@ict.usc.edu; marsella@ccs.neu.edu

Abstract. We argue for the importance of negotiation as a challenge problem for virtual human research, and introduce a virtual conversational agent that allows people to practice a wide range of negotiation skills. We describe the multi-issue bargaining task, which has become a *de facto* standard for teaching and research on negotiation in both the social and computer sciences. This task is popular as it allows scientists or instructors to create a variety of distinct situations that arise in real-life negotiations, simply by manipulating a small number of mathematical parameters. We describe the development of a virtual human that will allow students to practice the interpersonal skills they need to recognize and navigate these situations. An evaluation of an early wizard-controlled version of the system demonstrates the promise of this technology for teaching negotiation and supporting scientific research on social intelligence.

1 Introduction

Negotiation is an indispensable skill for any social creature. Civil society frowns on those that simply take what they need from others. Whether in the home, the market place or virtual market place, people achieve what they need through discussion and compromise. Unfortunately, most of us are poor negotiators. Research has documented a range of cognitive biases that undermine the quality of negotiated agreements [1], and companies invest billions of dollars in training their employees to negotiate and resolve conflict [2]. Automated decision tools might help people avoid these limitations, but to the extent computers are able to negotiate at all, it is only through very restrictive protocols that simplify away many of the complexities faced by human negotiators.

In this article, we argue for the importance of negotiation as a challenge problem for virtual human research. Negotiation engages a wide range of skills that are not only crucial for people, but essential for machines that socially engage with humans:

- **Intelligence:** Negotiations bring together a number of cognitive skills. They involve tradeoffs across multiple goals; they require one to infer and reason about the goals of one's negotiation partner, engaging theory of mind reasoning; they evoke emotion and these shape outcomes for good and ill.
- **Language:** Negotiations create unique challenges for natural language research as parties often violate the standard Gricean Maxims of cooperative communication. Negotiators hedge, obscure or outright lie about their preferences, or adopt sophisticated strategies such as reciprocal disclosure to build trust. A skilled negotiator is

attuned not only to explicit statements, but also to the implications of what is not said and the subtleties of how information is conveyed.

- **Embodiment:** Negotiation research emphasizes that expressions, postures and paralanguage can convey your partner's preferences and power, and improved skill in reading these signals leads to better outcomes. On the flip side, regulating one's own nonverbal signals can strongly influence how a negotiation will unfold.

Creating virtual humans with this range of capabilities can have important practical and scientific benefits. In practical terms, virtual human negotiators can help teach interpersonal skills [3-5]. More broadly, the capabilities needed to successfully negotiate can inform the design of a wide range of machines that socially interact with people. From a scientific perspective, the act of creating a virtual negotiator can serve to advance theories of human cognition. This can occur through the act of concretizing social theories into working artifacts [6], but also because virtual agents enable a level of experimental control unobtainable in most social science research [7].

This paper describes the development of a virtual human that allows people to practice a wide range of negotiation skills. We first describe the *multi-issue bargaining problem*, a formulation of negotiation that has been adopted by the social science, education and the multi-agent communities. We then describe the Conflict Resolution Agent, a conversational agent that performs this task and allows students to practice an array of negotiation concepts (see Fig. 2). The agent is being developed through an iterative design process, starting with face-to-face data collecting, moving to a wizard-assisted system, and then finally moving to a fully-automated virtual human. Currently, we are part-way through this design process but already the system has supported a number of scientific findings. We report on our progress and argue for the importance of negotiation as a challenge problem to advance virtual human research.

2 Definitions

Negotiations are dialogues aimed at reaching an agreement between parties when there is a perceived divergence of interests, beliefs, or in ways to achieve joint ends [8]. Although this definition is broad, researchers have sought to abstract essential elements of negotiations into more structured formalisms that are suitable for both teaching and scientific enquiry. In this paper, we focus on one useful and common abstraction known as the multi-issue bargaining task [9], which has become a *de facto* standard for both teaching and research on negotiation in both the social and computer sciences (e.g., see [2, 10, 11]). Multi-issue bargaining generalizes simpler games developed in game theory, such as the ultimatum game, and more closely approximates many of the challenges found in real-life negotiations. This task has received so much attention amongst educators and researchers because, with only a small number of mathematical parameters, one can evoke a wide range of psychologically-distinct decision-tasks. Thus, multi-issue bargaining has been used to teach and study a wide range of negotiation concepts.

Side A Payoff						Side B Payoff					
Record Crates		Lamps		Painting		Record Crates		Lamps		Painting	
Level	Value	Level	Value	Level	Value	Level	Value	Level	Value	Level	Value
0	\$0	0	0	0	\$0	0	\$0	0	0	0	\$0
1	\$20	1	\$10	1	\$100	1	\$10	1	\$30	1	\$0
2	\$40	2	\$20			2	\$20	2	\$60		
3	\$60					3	\$30				

Table 1. Example 3-issue bargaining problem

In its basic form, multi-issue bargaining requires parties (typically 2) to find agreement over a set of issues. Each issue consists of a set of levels and players must jointly decide on a level for each issue (levels might correspond to the amount of a product one player wishes to buy, or it might represent attributes of a single object, such as the price or warranty of a car). Each party receives some payoff for each possible agreement and each player's payoff is usually not known to the other party. The payoff is often assumed to be additive (i.e., a player's total payoff is the sum of the value obtained for each issue) and presented to players through a payoff matrix. For example, Table 1 illustrates the two payoff matrices for a hypothetical negotiation over items in an antique store. In this case, players must divide up three crates of records, two lamps and one painting, but each party assigns different value to items.

Preference weights: The weight each party assigns to issues defines one class of parameters for creating qualitatively different classes of negotiation. The payoff structure in Table 1 defines an *integrative* (or win-win) negotiation. For example, as player A receives the most value from the painting and records, whereas player B receives the most value from the lamps, the joint payoff is maximized when player B gets all the lamps and player A gets the rest (also known as the *Pareto efficient* solution). A *distributive* (or zero-sum) negotiation arises when both parties have conflicting preferences. For example, if both parties had the same payoff as side A, any gain in value to one side would result in an equal loss to the other side. The painting represents a special type of issue known as a *compatible issue* as one party doesn't incur a cost if the other party receives their preferred level. Compatible issues create an opportunity for *misrepresentation*. Specifically, if player B, claims that the painting has value to them, they can offer this 'invented' value in exchange for other items they want [12].

BATNA: The second important class of parameters is the Best Alternative to a Negotiated Agreement (BATNA) for each player. This represents how much a party would receive if the negotiation fails. For example, if player A already has a tentative deal with another player that affords him \$150, there is no reason to accept a deal worth less than \$150 from player B (e.g., 2 records and a painting). The BATNA represents the player's bargaining power, and as with preference weights, these are typically unknown to the other player. If player B's BATNA is only \$20, then player A has more potential power in the negotiation, although whether this translates into better outcomes depends on how each party shapes the other party's perceptions and how carefully they attend to the structure of the negotiation.

Figure 1 summarizes several basic negotiation concepts. The graph shows all 24 possible agreements defined in Table 1 in terms of the value each player receives. The Pareto frontier defines the set of *efficient* agreements. Expert negotiators should not accept any deal below this frontier as inefficient solutions can always be improved for one party without harming the other (thus increasing joint value), although inexperienced negotiators often fail to discover efficient solutions. The BATNAs define a *zone of agreement*. Any deal outside this zone should be rejected by one player as it is below their BATNA, however inexperienced negotiators often fail to follow this principle. The fact that the Pareto frontier is convex means there is integrative potential: players can improve on a 50-50 split by understanding each other's preferences and allocating each player their most important issue. Inexperienced negotiators often assume negotiations are distributive (a 'fixed-pie' bias) and fail to realize integrative potential.

It is important that preference weights and BATNA are typically unknown to the other party and must be estimated through language and/or the pattern of offers. Much of the skill of negotiation comes from learning when to reveal truthful information or how to elicit truthful information from the other party. A player that reveals too much information without reciprocation could be exploited, creating a tension between cooperation and competition. However, even when all information is public, players often fail to find efficient agreements.

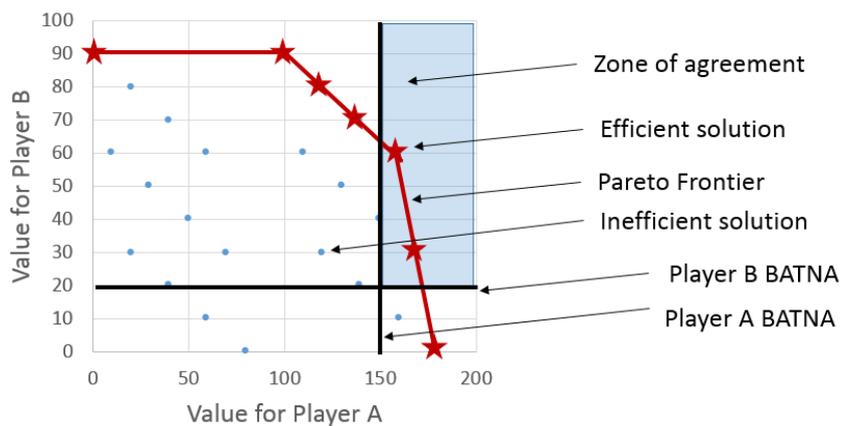


Fig. 1. Summary of key bargaining concepts

Preference weights and BATNA define the basic mathematics of the multi-issue bargaining task¹ but several other factors can be varied which are irrelevant from the perspective of classical rationality, but that can have profound differences on human decision-making, especially for inexperienced negotiators. For example, preferences can be presented as losses or gains. Issues can carry moral significance [13]. Parties can negotiate for themselves or as representatives of their organization or as part of a team [14].

¹ Time pressure can be introduced by adding a deadline or a temporal discounting function. Automated negotiation agents usually require parties to alternate complete offers. Generalizations are also possible, e.g., by relaxing the assumption of an additive utility function.

The amount of information available (e.g., the other player's preferences and/or BATNA) can also be varied. All of these – and other – factors have been shown to influence the negotiation processes, especially for inexpert negotiators.

Uses: Researchers can create a vast number of psychologically-distinct negotiation problems with only a small number of mathematical parameters and a bit of textual framing. Thus, multi-issue bargaining has proven an especially rich tool for the study and teaching of human social skills, as well as a tool for advancing artificially intelligent agents. For example, in emotion research, bargaining tasks are used to examine how signaled or induced emotion shapes joint outcomes [11]. In conflict-resolution research it is used to study various social processes involved in resolving disputes [8]. In social neuroscience, it is used to examine specific brain regions associated with social cognition [15]. In game theory, it is used to advance rational models of multi-party decision-making [16]. In artificial intelligence, it serves as a standard challenge problem for advancing automated models of social decision-making [10]. Finally, in educational settings, this extensive body of research provides a firm theoretical basis for informing pedagogy as bargaining games are used to teach a wide range of interpersonal skills including negotiation, conflict-resolution, teamwork, emotional intelligence and inter-cultural fluency (e.g., see the leadership exercises at the Northwestern Dispute Resolution Research Center at negotiationexercises.com). Therefore, virtual humans that can perform bargaining tasks in a general way will have broad impact on science and education.

3 Related Work and Current Limitations

Research on negotiation within computer science has already yielded tangible benefits. Machines can predict the outcome of a negotiation by analyzing the verbal and nonverbal cues of negotiators [17, 18]. Even stylized virtual humans evoke physiological threat [19] and influence negotiation outcomes with their emotional expressions [20]. Virtual humans can use language to establish beneficial long-term relationships with other negotiators [21]. An analysis of the language of negotiation has advanced dialog research on turn-taking and incremental speech production [22]. We build on this research, but also extend it in ways that enhances its connections with the larger body of research on teaching and understanding human negotiation skills.

Several algorithms have been developed to automate the decision-making of an artificial negotiator [10]. Unfortunately, most of this work has focused on agent-agent interaction and adopts assumptions that may not apply in human negotiations. Such agents only communicate through formal representations of offers, whereas people rely heavily on language. Agents usually only allow the exchange of complete offers, whereas people often focus on a subset of issues at a time (indeed, learning how to “package” different issues is a key skill taught to negotiators). Agents assume money is the only source of value whereas people often assign value to intangible considerations like fairness or maintaining relationships [23]. These restrictions avoid many of the challenges that face human negotiators (although see [24] for one attempt to relax these restrictions). Nevertheless, this research can serve as an important basis for the reasoning techniques that inform a virtual human negotiator.

Education researchers have looked at the potential of bargaining agents to teach negotiation, though none of these systems have tackled spoken interaction. For example, the pocket negotiator uses preference-elicitation techniques and visualizations of the Pareto frontier to help students better prepare for a face-to-face negotiation [25]. ELECT BiLAT allows students to practice a series of negotiations with virtual characters that use sophisticated decision-theoretic and theory-of-mind techniques to guide their behavior. However, the main pedagogical focus of BiLAT, like the pocket negotiator, is on the preparations leading up to a negotiation [26]. Kraus and colleagues have shown that negotiating with a disembodied rational agent can help students learn [3]. This research serves to inform how to use virtual humans to teach negotiation skills.

Finally, a line of research within the virtual human community has explored natural language negotiations with embodied agents. For example, the SASO system allowed student-soldiers to negotiate with a local leader over how best to conduct a peacekeeping operation [27]. However, this class of approaches adopts a very different formalism of negotiation, building more on planning and shared-plans frameworks (e.g., [28]), and thus has only limited relevance to the larger body of research on multi-issue bargaining. Nonetheless, this research provides a foundation for the natural language understanding and dialog processes required for a virtual human negotiator.



224: I'll tell you what. I'll take this box of records 'cause it looks like it has the least.
 CRA: That doesn't seem fair though...
 224: Why not? [exasperated laugh]
 CRA: Well, you see, I have a buyer right now that is interested in old records.
 224: So do I.
 CRA: Your customers would probably love those lamps.
 224: My customers?

Fig. 2. A participant (#224) interacting with the Conflict Resolution Agent.

4 The Conflict Resolution Agent

The Conflict Resolution Agent (CRA), pictured in Figure 2, is a game-like environment that allows negotiation students to engage with different virtual human role-players across a variety of multi-issue bargaining problems. Our goal is to allow students to communicate with a fully automated agent through natural language and nonverbal expressions. By altering preference weights, BATNAS, and task-framing, students can be presented with a wide range of negotiation and dispute-resolution concepts such as integrative potential, anchoring, reciprocal information exchange, rights vs. interests, emotional intelligence and establishing rapport. This mirrors how bargaining games are used to teach negotiation in business schools.

Virtual humans can augment negotiation training in many of the same ways that automated tutoring research has benefited other “harder” skills, by allowing students more opportunities to experience the domain, tailoring their experience to match their

current skills, and providing targeted feedback. Currently, negotiation is taught by a mixture of lecture and in-class “simulations” where students role-play bargaining exercises with each other. Simulations are widely considered to have the greatest teaching value but are difficult to realize. In a typical semester-long negotiation course, students might participate in only a small number of simulated negotiations.² In-class simulations also have something of a blind-leading-the-blind flavor, with large variance in quality and many students failing to achieve the core principle underlying the exercise. The instructor then leads a discussion illustrating why certain students succeeded or failed. Agent technology can improve this process by allowing students to practice as much as they like with a partner programmed to more consistently evoke the intended behaviors and negotiation processes. Further, as all student behavior is being tracked, understood and recorded, agents or instructors could provide customized tutorial feedback and commentary on the student’s behavior. Virtual humans can also complement the growing interest in online courses. Of course, another aim of the project is to advance the capabilities of virtual humans more generally, and the techniques underlying CRA build upon and reinforce domain-independent techniques for virtual humans.

To this end, CRA is being developed through a series of iterations, beginning with a face-to-face data collection to serve as a baseline for comparison and to create a large corpus to inform the design of individual system components. Presently, we have completed the face-to-face data collection, designed the basic game environment, and completed several iterations on improving a wizard-controlled system (described below).

CRA is implemented with the publicly-available virtual human toolkit [29]. In the Wizard-of-Oz setup (WOz), CRA is semi-automated, with low-level functions carried out automatically, while two wizards make high-level decisions about the agent’s verbal and nonverbal behavior. Gestures and individual utterances are based on data collected during face-to-face negotiations between inexpert negotiators on variants of the task shown in Table 1. The WOz interface allows the agent to speak over 10,000 distinct utterances. Utterances are synthesized by the NeoSpeech text-to-speech system and gestures and expressions are generated automatically by NVBG [30] and realized using the SmartBody character animation system [31]. This low-level automation complements and facilitates the decision-making of the wizards. Details of the development and capabilities of the CRA WOz interface can be found in [32].

CRA realizes a physically-embodied version of the multi-issue bargaining task developed by Carnevale and described in [17]. As can be seen in Figure 2, issues are represented as different types of physical objects (e.g., crates of records, lamps, and paintings) and levels correspond to the number of each type of item the player receives. Participants communicate with CRA through spoken natural language (currently interpreted through the wizards) or by manipulating, gazing at, and/or gesturing at the physical objects. The intent behind the physical objects is to elicit multimodal behavior and create multiple communication channels to facilitate the understanding of participant intent. For example, the participant can make an offer via language (“Would you like

² Personal communication with Professor Peter Kim, instructor of the negotiation course at the University of Southern California’s Marshall School of Business.

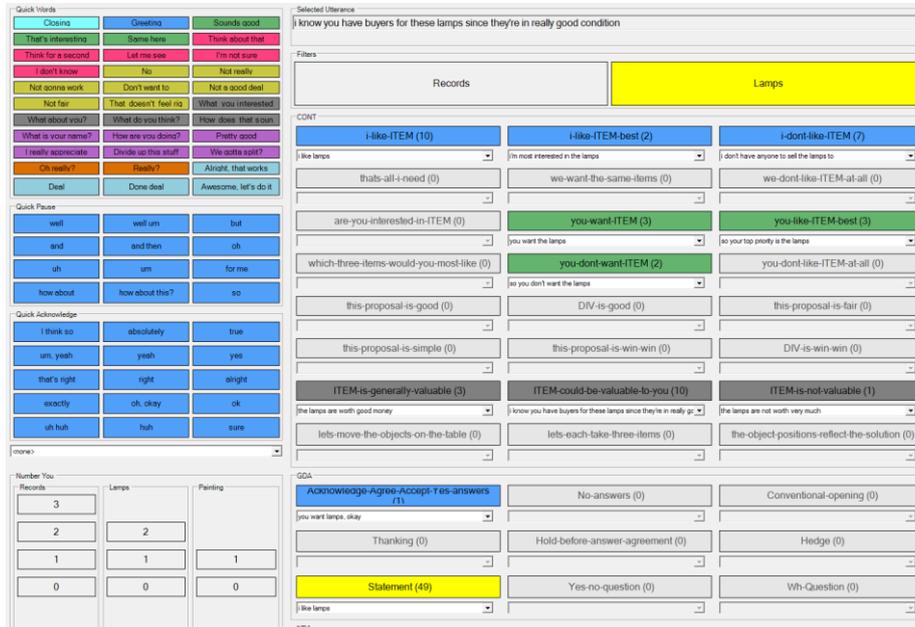


Fig. 3. Partial screenshot of wizard verbal interface. Items at upper-left correspond to short, high-frequency utterances. The right half of the interface organizes longer and less frequent utterances by speech acts. Filters at the top allow navigation by speech acts' specific topics.

the painting?"), moving the objects, or both. The agent can respond in kind, making offers either via speech or by manipulating the objects.

One of the challenges in designing the wizard interface is allowing wizards to rapidly access a large number of possible utterances quickly enough to approximate the pace of normal human dialog [32]. The wizards select amongst utterances using a filtering system which acts as a decision tree (see Figure 3). Utterances can be filtered by the class of speech act and by the negotiation items mentioned in the intended utterance. The wizards also use pause fillers (e.g., "uh", "um") and gaze behaviors (e.g., look away or at objects) to hold the turn until appropriate responses can be selected, or use "hot keys" to trigger common responses. A separate nonverbal interface allows wizards to select appropriate postures and gestures, and to move the items under discussion.

One reason for using human wizards is that it allows us to rapidly experiment with different dialog and negotiation policies. Although wizards use their own judgment to select verbal and nonverbal behavior, these choices are scripted to facilitate subsequent automation. This scripting is based on our analysis of the face-to-face dialogues and also to realize specific negotiation concepts. For example, the concept of reciprocal information exchange states that a negotiator should be reluctant to reveal his or her own preferences unless the other party reciprocates. Thus, wizard behavior is scripted to only reveal minor preference information initially and subsequently match the participant's own disclosure. Another concept is integrative potential (i.e., discovering win-win solutions). Many students fail to discover integrative solutions because they

assume fixed-pie bias (they assume negotiations are distributive); this is because they don't ask the right questions about the other party's interests and they are often too soft when negotiating with another student, and thereby don't effectively communicate their own priorities to the other party. Students also sometimes confuse the number or level of an item with its value. To facilitate the discovery of efficient solutions, wizards are scripted to query the other party about their preferences, and make proposals that are more efficient with respect to these preferences. Wizards are also scripted to be somewhat tough, to be reluctant to concede their high-value items, and to attend to the value of proposed deals rather than fixate on the number of items.

5 Data Collection and Evaluation

We evaluate CRA to demonstrate both the naturalness of the spoken language interaction and the extent to which wizards can guide inexperienced negotiators towards efficient solutions. The initial design of CRA is informed by a face-to-face negotiation dataset (113 same-sex dyads; 38 female). Each dyad was recruited from an on-line job service and randomly assigned to an integrative or distributive negotiation. The integrative task matched the structure of Table 1, except that the painting was only worth 5 for side A (not 100) but still worth 0 for side B. In the distributive task, both sides received the payoff of side A in Table 1, except that side A received 5 for the painting and side B received 0 (i.e., the painting was a compatible issue in both the integrative and distributive tasks). Participants could make money based on their performance. Rather than dollars, participants received lottery tickets based on the value of items they obtained. If they failed to reach agreement, their BATNA equaled the number of tickets they would have received for one of their highest-value items. Tickets were then entered into a \$100 lottery.

We collected several traits on participants (not relevant to this paper) and several subjective post-game measures. They rated their own preferences (used to verify they understood the task), their satisfaction with the outcome, and how cooperatively they behaved. They also were asked for their impressions of their negotiation partner: the partner's preferences, their satisfaction with the deal, their cooperativeness, whether they established rapport, and how easy it was to come to an agreement with them.

The face-to-face data serves as a baseline for comparing the performance of the virtual human, both in terms of outcomes (e.g., the quality and efficiency of agreements), process (e.g., is the quality and fluidity of agent speech similar to human negotiators), and subjective impressions. All dialogues were manually transcribed and annotated with semantic frames containing up to eight different key-values (see [32]). The frames encode information about generic dialogue acts (e.g., statement or question), negotiation-specific dialogue acts (e.g., make or reject an offer), propositional content templates (e.g., i-like-ITEM-best), offers, and valence. These frame annotations serve as the basis for organizing the WOz interface. Sessions were also automatically annotated with facial expressions and vocal quality.

We next ran three rounds of WOz interaction (15, 10 and 12 participants, respectively – 18 female), improving the interface after each iteration to enhance the fluidity



Fig. 4. Lottery tickets obtained by players and their partners

of interaction and increase the variety of utterances. The recruitment and design were identical to the face-to-face collection except that all participants played as side B of the integrative task, and all participants were led to believe they were interacting with an agent. As with face-to-face data, participants received lottery tickets based on the value of items they obtained. Male participants interacted with a male virtual human, pictured in Figure 2, while females interacted with a female virtual human controlled by the same wizard interface. The female virtual human uses the same utterances, general dialogue policy and gestures, but does differ in appearance and voice.

Wizards followed a script: they acted as if the participant preferences were unknown; the wizard avoided volunteering their own preferences unless participants used reciprocal information exchange; the wizard avoided making the first offer unless directly asked, at which point they would make a distributive offer (ask for 2 lamps and 1 record); if the participant insisted on a 3-3 split, the wizard responded “We should focus on the value not the number of items” but did not push the issue further. If the participant revealed that the painting had no value, but still insisted on obtaining it, the wizard would first ask why the participant wanted something of no value before acquiescing. The aim of the WOz script was to see if participants would discover integrative value.

We next describe a comparison of face-to-face and WOz data. As all WOz participants faced an integrative negotiation, we only compare against the integrative subset of face-to-face data: 66 same-sex dyads (12 female).

Outcomes: The wizard script was designed to help participants to discover the concept of integrative potential and this goal was achieved. From t-tests on the number of points both participants and their partners received in the negotiation, we see that participants realized more integrative potential (i.e., joint gain) when negotiating with CRA than when negotiating with another participant. As can be seen in Figures 4, both participants ($t(1,64) = -1.95, p = .05$) and their computer partner ($t(1,64) = -5.98, p < .001$) earned more points when negotiating with CRA than when both participants were humans. This benefit persisted across the three phases of development.

We examined the outcomes to understand why people performed better with CRA than when negotiating with inexperienced human partners. Differences arise from two errors: 1) how participants dealt with the compatible issue and 2) confusion between the number and value of objects. Recall that the painting is a compatible issue, as it only has value for side A. Instead, side B often fought for the painting and got it far more often

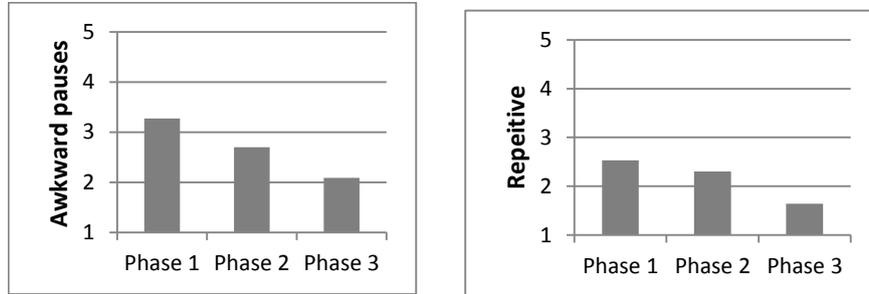


Fig. 5. Decreases in unnaturalness of the interaction across WOz iterations

with human (83%) vs. agent (53%) partners ($\chi^2(1, N=66) = 6.97, p = .008$). Participants also favored an equal division of the objects more often with human (73%) than agent (22%) partners ($\chi^2(1, N=66) = 17.24, p < .001$). For example, one participant said, “Me walking away with two things and you walking away with four things is not going to work” even though this was the most efficient deal. Indeed, side B would often fight for the painting only to give up on a higher value item in order to maintain a 3-3 split. They occurred less often with CRA, probably as the wizard-script called attention to difference between the number of items and their value. It is also possible that people are less willing to apply the norm of fairness to computer programs (e.g., see [33]). Further studies will disambiguate these factors.

Subjective impressions: We assessed several subjective impressions of the negotiation and one’s partner when negotiating with both human participants and CRA. These are summarized in Table 2.

Satisfaction: As shown in Table 2, participants were as satisfied when paired with CRA as when paired with human participants ($t(1,63) = -0.09, p = .93$). Furthermore, they perceived CRA to be more satisfied with the negotiation than they perceived their human partners ($t(1,63) = -2.19, p = .03$). This is likely due to the fact that CRA obtained better outcomes than the inexperienced human participants. Participants reported no greater satisfaction for themselves ($F(2,33) = 0.20, p = .82$) or for their partners ($F(2,33) = 1.27, p = .30$) as the agent was improved through development.

Cooperation: CRA was rated as cooperative as human partners on an 8-item scale ($t(1,63) = 0.19, p = .85$). They also rated their own level of cooperation no differently

when paired with CRA than with a human partner ($t(1,63) = -0.92, p = .36$). Participants also did not rate their partners ($F(2,33) = 1.29, p = .51$) or themselves ($F(2,33) = 1.84, p = .18$) differently as the agent was improved through each iteration of development.

Rapport: Participants felt the same level of rapport (11-item scale) when paired with CRA as when paired with other human participants ($t(1,63) = 0.08$,

Table 2: Subjective Impressions	Partner	
	Human	CRA
Satisfaction	6.28 (0.19)	6.31 (0.25)
Partner satisfaction	6.00 (0.19)	6.47 (0.12) *
Ratings of partner	5.66 (0.27)	5.60 (0.17)
Ratings of self	5.76 (0.31)	6.05 (0.13)
Rapport	5.25 (0.19)	5.23 (0.14)
Ease of agreement	5.73 (0.19)	5.83 (0.21)
Note. * = $p \leq .05$, ** = $p \leq .01$, *** = $p \leq .001$.		

$p = .94$). Likewise, they reported no less ease of reaching agreement when paired with our agent than with a human partner ($t(1,63) = -0.36, p = .72$). Participants also reported no greater rapport ($F(2,33) = 0.68, p = .51$) or ease of reaching agreement ($F(2,33) = 1.18, p = .32$) as the agent was improved through development.

Design goals: After each iteration of the WOz, we worked to improve the fluidity of the agent by reducing latency between speech acts as well as increasing the number of utterances that the agent could use during the negotiation. We tested these improvements by several subjective ratings. In addition to the survey questions described above, participants in the three rounds of WOz testing also rated the frequency of awkward pauses and how repetitive the agent seemed. We conducted t-tests on their ratings of how many awkward pauses and how much repetition there was in the conversation. As can be seen in Figures 5, over the phases of development, participants reported significantly fewer awkward pauses ($F(2,33) = 4.85, p = .01$) and viewed the agent as significantly less repetitive ($F(2,33) = 4.26, p = .02$).

6 Summary and conclusion

Negotiation requires a number of cognitive and interpersonal skills and the multi-issue bargaining tasks allow many of these skills to be independently evoked and practiced by the judicious choice of preference weights and BATNA assigned to each party in a negotiation. Because of the complexity and necessity of negotiation in everyday life, bargaining problems have been an active focus of research across several research communities including multi-agent systems, social science, and education. Here we argue that bargaining is an important challenge area as it evokes the three pillars of the intelligent virtual agent community: intelligence (e.g., theory of mind reasoning, emotion, and decision-making), language (i.e., the full range of natural language processing) and embodiment (e.g., nonverbal communication).

The project has already yielded some tangible results along these lines. For example, the data collected in the face-to-face negotiations has been used to develop and evaluate potential algorithms for inferring a user's preferences from their pattern of offers and stated preferences. We found that techniques developed to fully automate negotiation perform poorly on human data but that heuristic methods can do a fairly accurate job at estimating a negotiator's preferences. We also find evidence that techniques can detect if a negotiator is lying by identifying discrepancies in what they say they want vs. what they are willing to offer [34].

In summary, we reported on our efforts to build a Conflict Resolution Agent that can negotiate with people and allow students to practice their negotiation skills. The current Wizard-of-Oz system evokes behavior similar to behaviors observed in face-to-face negotiations and leads people towards better solutions than they discover when negotiating against novice human negotiators.

Acknowledgments: This research was supported by the US Army. The content does not necessarily reflect the position or the policy of any Government, and no official endorsement should be inferred.

References

1. Thompson, L.L. and R. Hastie, *Judgement tasks and biases in negotiation*. 1990: National Institute for Dispute Resolution.
2. Movius, H., *The Effectiveness of Negotiation Training*. Negotiation Journal, 2008. **24**(4): p. 509-531.
3. Lin, R., Y. Oshrat, and S. Kraus. *Investigating the benefits of automated negotiations in enhancing people's negotiation skills*. in *8th International Conference on Autonomous Agents and Multiagent Systems*. 2009.
4. Core, M., et al., *Teaching negotiation skills through practice and reflection with virtual humans*. Simulation, 2006. **82**(11): p. 685-701.
5. Broekens, J., et al. *Virtual reality negotiation training increases negotiation knowledge and skill*. in *Intelligent Virtual Agents*. 2012. Springer.
6. Simon, H., *The sciences of the Artificial*. 1969, Cambridge, MA: MIT Press.
7. Blascovich, J., et al., *Immersive virtual environment technology as a methodological tool for social psychology*. Psychological Inquiry, 2002. **13**: p. 103-124.
8. Carnevale, P.J. and D.G. Pruitt, *Negotiation and Mediation*. Annual Review of Psychology, 1992. **43**(1): p. 531-582.
9. Kelley, H.H. and D.P. Schenitzki, *Bargaining*. Experimental Social Psychology. New York: Holt, Rinehart, and Winston, 1972: p. 298-337.
10. Baarslag, T., et al., *Evaluating practical negotiating agents: Results and analysis of the 2011 international competition*. Artificial Intelligence, 2013. **198**: p. 73-103.
11. Van Kleef, G.A., C.K.W. De Dreu, and A.S.R. Manstead, *The interpersonal effects of anger and happiness in negotiations*. Journal of Personality and Social Psychology, 2004. **86**(1): p. 57-76.
12. O'Connor, K.M. and P.J. Carnevale, *A Nasty but Effective Negotiation Strategy: Misrepresentation of a Common-Value Issue*. Personality and Social Psychology Bulletin, 1997. **23**(5): p. 504-515.
13. Dehghani, M., et al., *Sacred values and conflict over Iran's nuclear program*. Judgment and Decision Making, 2010. **5**(7): p. 540-546.
14. O'Connor, K.M., *Groups and Solos in Context: The Effects of Accountability on Team Negotiation*. Organizational Behavior and Human Decision Processes, 1997. **72**(3): p. 384-407.
15. Bhatt, M.A., et al., *Neural signatures of strategic types in a two-person bargaining game*. Proceedings of the National Academy of Sciences, 2010. **107**(46): p. 19720-19725.
16. Nash Jr, J.F., *The bargaining problem*. Econometrica: Journal of the Econometric Society, 1950: p. 155-162.
17. Park, S., et al. *Mutual behaviors during dyadic negotiation: Automatic prediction of respondent reactions*. in *Affective Computing and Intelligent Interaction (ACII), 2013 Humaine Association Conference on*. 2013. IEEE.
18. Li, R., J. Curhan, and M.E. Hoque, *Predicting Video-Conferencing Conversation Outcomes Based on Modeling Facial Expression Synchronization*, in *The 11th IEEE International Conference on Automatic Face and Gesture Recognition*. 2015, IEEE: Ljubljana, Slovenia.

19. Khooshabeh, P., et al., *Negotiation strategies with incongruent facial expressions of emotion cause cardiovascular threat*, in *35th annual meeting of the Cognitive Science Society*. 2013: Berlin, Germany.
20. de Melo, C., P.J. Carnevale, and J. Gratch, *The Effect of Expression of Anger and Happiness in Computer Agents on Negotiations with Humans*, in *the Tenth International Conference on Autonomous Agents and Multiagent Systems*. 2011: Taipai, Taiwan.
21. Mell, J., G. Lucas, and J. Gratch, *An Effective Conversation Tactic for Creating Value over Repeated Negotiations*, in *14th International Conference on Autonomous Agents and Multiagent Systems*. 2015: Istanbul, Turkey.
22. DeVault, D., K. Sagae, and D. Traum, *Incremental interpretation and prediction of utterance meaning for interactive dialogue*. *Dialogue & Discourse*, 2011. **2**(1): p. 143-170.
23. Curhan, J.R. and H.A. Elfenbein, *What do people value when they negotiate? Mapping the domain of subjective value in negotiation*. *Journal of Personality & Social Psychology*, 2006. **91**(3): p. 493-512.
24. Rosenfeld, A., et al. *NegoChat: a chat-based negotiation agent*. in *13th international conference on Autonomous agents and multi-agent systems*. 2014. Paris, France.
25. Hindriks, K.V. and C.M. Jonker. *Creating human-machine synergy in negotiation support systems: Towards the pocket negotiator*. in *Proceedings of the 1st International Working Conference on Human Factors and Computational Models in Negotiation*. 2008. ACM.
26. Kim, J.M., et al., *BiLAT: A game-based environment for practicing negotiation in a cultural context*. *International Journal of Artificial Intelligence in Education*, 2009. **19**(3): p. 289-308.
27. Traum, D., et al., *Multi-party, Multi-issue, Multi-strategy Negotiation for Multi-modal Virtual Agents*, in *8th International Conference on Intelligent Virtual Agents*. 2008, Springer: Tokyo, Japan.
28. Grosz, B. and S. Kraus, *Collaborative Plans for Complex Group Action*. *Artificial Intelligence*, 1996. **86**(2): p. 269-357.
29. Hartholt, A., et al. *All together now: Introducing the virtual human toolkit*. in *Intelligent Virtual Agents*. 2013. Springer.
30. Lee, J. and S. Marsella. *Nonverbal Behavior Generator for Embodied Conversational Agents*. in *6th International Conference on Intelligent Virtual Agents*. 2006. Marina del Rey, CA: Springer.
31. Thiebaut, M., et al., *SmartBody: Behavior Realization for Embodied Conversational Agents*, in *International Conference on Autonomous Agents and Multi-Agent Systems*. 2008: Portugal.
32. DeVault, D., 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*. 2015, AAAI Press: Stanford, CA.
33. de Melo, C., P.J. Carnevale, and J. Gratch, *Humans vs. Computers: The Effect of Perceived Agency on People's Decision Making*. 2013, University of Southern California.
34. Nazari, Z., G. Lucas, and J. Gratch, *Opponent Modeling for Virtual Human Negotiators*, in *15th International Conference on Intelligent Virtual Agents*. 2015, Springer: Delft, the Netherlands.