The Benefits of Virtual Humans for Teaching Negotiation

Jonathan Gratch, David DeVault, and Gale Lucas

University of Southern California, Institute for Creative Technologies, Playa Vista, CA USA
{gratch,devault}@ict.usc.edu

Abstract. This article examines the potential for teaching negotiation with virtual humans. Many people find negotiations to be aversive. We conjecture that students may be more comfortable practicing negotiation skills with an agent than with another person. We test this using the Conflict Resolution Agent, a semi-automated virtual human that negotiates with people via natural language. In a between-participants design, we independently manipulated two pedagogically-relevant factors while participants engaged in repeated negotiations with the agent: perceived agency (participants either believed they were negotiating with a computer program or another person) and pedagogical feedback (participants received instructional advice or no advice between negotiations). Findings indicate that novice negotiators were more comfortable negotiating with a computer program (they self-reported more comfort and punished their opponent less often) and expended more effort on the exercise following instructional feedback (both in time spent and in self-reported effort). These findings lend support to the notion of using virtual humans to teach interpersonal skills.

1 Introduction

Most people hate to negotiate and this aversion has real economic costs [1]. Not surprisingly, negotiation expertise is a highly-valued commodity. Negotiation is often taught in professional schools, as part of a business or law degree. For example, as part of a Master in business administration, students might take a semester-long course on negotiation concepts. For those seeking a more cursory introduction, consulting companies offer intensive short courses. For example, Vantage Partners, a spinoff of the Harvard Business School, offers 3-day tutorials to corporate executives. Regardless of the length of instruction, negotiation is taught via a mixture of instruction (typically classroom lectures) and hands-on experience (typically where students pair-off and engage in a simulated negotiation with each other). In business schools, these simulations are often run by dedicated staff trained to be experts in experiential learning techniques.

All of this is big business. It has been estimated that billions of dollars are spent on teaching negotiation [2]. Professional schools and consulting companies charge high fees for these services. Even the creation of simulated negotiations is a money making operation. Instructors submit their teaching cases to repositories, such as the Kellogg Schools Dispute Resolution Resource Center (DRRC) and instructors are expected to pay to use these cases in the classroom. As a result, professional negotiation skills are mainly limited to the elite.

Virtual human technology has the potential to address the challenges and expense in teaching negotiation [3-5]. In this article, we examine the experiential aspect of negotiation training. Currently, students experience negotiations by practicing their skills on each other, playing simulated negotiations such as those maintained by the DRRC. Being novices, these negotiations have something of the flavor of the blind leading the blind. Especially in introductory cases, the majority of the students fail to incorporate the key teaching points into their negotiation behavior. Rather, the professor or professional facilitator will walk around the classroom, find the few students that performed well, and lead a classroom discussion on why such outcomes occurred.

We argue virtual humans, in particular, can serve as an especially valuable tool for augmenting experiential learning. Virtual humans can serve as both automated role-players and automated tutors; allowing students to practice with more proficient (computerized) partners and then receive targeted feedback on their own performance, much as is done in cognitive tutoring in more conventional domains. But further, we conjecture that virtual humans, by the very nature of being artificial, can help mitigate the anxiety people often feel in negotiations, and thereby enhance their efficiency in learning.

We first review some key negotiation skills and motivate why virtual humans may be well-suited to teaching them. We introduce these within the context of the multi-issue bargaining task, an abstract characterization of negotiation often adopted for teaching these skills. We next review some previous approaches to using automation for teaching these skills, then present the Conflict Resolution Agent, a virtual human negotiator that we will use in our study. We finally present experimental results supporting our conjecture that virtual humans are uniquely beneficial to teaching negotiation and conclude with some final thoughts.

2 Negotiation Skills

Multi-issue Bargaining. We expect virtual humans will be most effective for the introductory modules of a semester-long course or for the short intensive instruction offered by consulting firms. In these more introductory settings, simulated negotiations often follow a very stylized form that is more amenable to automation. Specifically, we focus on one useful and common abstraction of negotiation known as the multi-issue bargaining task [6], which has become a *de facto* standard for introductory negotiation simulations, as well as research on negotiation in both the social and computer sciences [e.g., see 2, 7, 8]. 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.

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.

Negotiation concepts. One important set of negotiation concepts relates to the relative importance each party assigns to different issues. The payoff structure in Table 1 is used to teach the concept of *integrative potential* and serves to define an *integrative* (or win-win) negotiation 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. In contrast, if both parties have the same priorities, this creates a *distributive* (or zero-sum) negotiation as any gain in value to one side would result in an equal loss to the other side.

Most students assume their opponent wants the same thing as them (i.e., they assume they are engaged in a distributive negotiation). Thus, integrative structures provides students the opportunity to *create* value by discovering integrative potential. They can only find this potential if they are willing to exchange information about their preferences, but students often fear to reveal too much, lest they be exploited by their opponent. Thus, integrative negotiations also provide the opportunity to teach ways to establish trust and safely exchange information. For example, one tactic is reciprocal information exchange, in which one provides a small amount of information and only provides more if the opponent reciprocates [9]. In contrast, distributive negotiations provide the students with the opportunities to learn tactics for how to *claim* value, as they can only improve their position by overpowering their opponent. These can including making threats or staking out strong positions [10].

Another important negotiation concept is the Best Alternative to a Negotiated Agreement (BATNA) for each player. This represents how much a party would receive

Table 1: An example 3-issue integrative bargaining problem

Side A Payoff

Records		Lamps		Painting	
Level	Value	Level	Value	Level	Value
0	$0	0	$0	0	$0
1	$20	1	$10	1	$5
2	$40	2	$20		
3	$60				

Side B Payoff

Records		Lamps		Painting	
Level	Value	Level	Value	Level	Value
0	$0	0	$0	0	$0
1	$10	1	$30	1	$0
2	$20	2	$60		
3	$30				

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. By focusing on their own and their opponent's BATNAs, students can better understand their bargaining power and how to claim value. For example, when claiming value, it can be effective to mislead one's opponent about the size of one's own BATNA.

Virtual humans can implement and reason about these different concepts. For example, an agent can be programmed to engage in reciprocal information exchange, allowing the student to explore and discover this concept. These techniques could also facilitate tutorial feedback. In this vein, Nazari showed that automated techniques can classify if a student is communicating distributive or integrative preferences [11] and an automated tutor could contrast this objective communication with a student's subjective beliefs about what he or she communicated to their opponent.

Negotiation anxiety. Teaching negotiation requires imparting several different types of skills. Up to this point, we have been discussing cognitive skills. These including recognizing the structure of the negotiation (integrative versus distributive), identifying each player's bargaining power, and deciding which tactics to use depending on these factors. But many people find negotiation aversive. People often experience negative affect or anxiety in negotiations and this can undermine their cognitive skills [12, 13] and lead to poorer negotiation outcomes [1]. This can be especially true when negotiations are distributive and students must focus on claiming value at the expense of their opponent. Thus, students of negotiation are confronted with the dual challenge of learning cognitive skills (negotiation concepts) while simultaneously learning to manage and regulate their emotions.

One of the best ways to reduce negative affect is to improve and automatize cognitive skills. Negotiations are cognitively challenging and can create high-cognitive load, but this cognitive load can make them more susceptible to emotional influences [14] and lower the cognitive resources available for emotion regulation [15, 16]. More broadly, negative affect can make it more difficult for negotiators to explore solutions and create value [17]. Thus, if students have the opportunity to practice cognitive skills in a safe and positive environment, they may learn to more quickly become comfortable with cognitive aspects of negotiation and thereby free up resources to regulate their emotions.

We conjecture that virtual humans can reduce negotiation negative affect and negotiation anxiety and promote cognitive learning. Previous research has suggested that people feel less fear and anxiety when they practice interpersonal skills with virtual humans [18, 19]. We predict that these findings will extend to the context of negotiation.

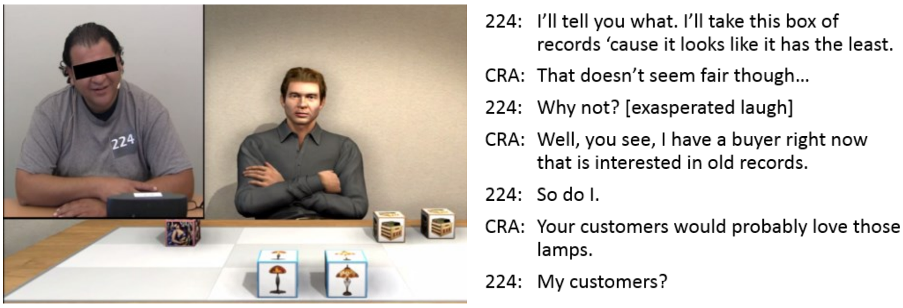

Fig. 1. A participant interacting with the Conflict Resolution Agent.

3 Prior Work and the Conflict Resolution Agent

Researchers have looked at the potential of artificial intelligence technology to teach negotiation. Several systems have used automated techniques to help students prepare for a negotiation. For example, the pocket negotiator uses preference-elicitation techniques and visualizations of the Pareto frontier to help students better understand their preferences and limits [20]. ELECT BiLAT explore the potential of an embodied agent to teach negotiation. Students could practice a series of negotiations with virtual characters that uses menu based "conversation" and sophisticated decision-theoretic and theory-of-mind techniques to guide their behavior. Like the Pocket Negotiator, this pedagogy focuses students on the preparations leading up to a negotiation [21].

Other researchers have focused on teaching tactics that occur during the negotiation. Kraus and colleagues have shown that negotiating with a disembodied rational agent can help students learn better negotiation tactics [3]. SASO is perhaps the only negotiation system that supports conversational negotiation with an embodied agent [22]. It allows student-soldiers to negotiate with a local leader over how best to conduct a peacekeeping operation, however, it adopted a very different formalism of negotiation, building more on planning and shared-plans frameworks (e.g., [23]), 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.

Most recently, our group has proposed a conversational virtual human that performs the multi-issue bargaining task and we adopt this system to examine our hypotheses. The Conflict Resolution Agent (CRA), pictured in Figure 1 [24], is a game-like environment that allows negotiation students to engage with a variety of virtual human role-players across a variety of multi-issue bargaining problems. The current, wizard-of-Oz (WOz) system allows students to communicate through natural language and nonverbal expressions. CRA is implemented with the publicly-available Virtual Human Toolkit [25]. Low-level functions such as speech and gesture generation are carried out automatically, while two wizards make high-level decisions about the agent's verbal and nonverbal behavior. The WOz interface allows the agent to speak over 5000 distinct utterances. Utterances are synthesized by the NeoSpeech text-to-speech system and gestures and expressions are generated automatically by NVBG [26] and realized using the SmartBody character animation system [27]. 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 [28].

CRA realizes a physically-embodied version of the multi-issue bargaining task developed by Carnevale and described in [29]. As can be seen in Figure 1, 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 the painting?"), moving the objects, or both. The agent can respond in kind, making offers either via speech or by manipulating the objects.

4 Experiment

We devised an experiment to test two hypotheses concerning the pedagogical potential of CRA. Most importantly, we wanted to assess if negotiating with a computer program felt more comfortable and safe than negotiation with another person. Secondly, we wanted to assess if pedagogical feedback would help improve negotiation effort and performance.

Hypothesis 1: Participants will feel more comfortable negotiating with a tough computer opponent compared with a tough human opponent

We instructed wizards to adopt a tough negotiation stance to evoke negotiation anxiety. We then manipulated participants' belief as to whether the CRA agent was controlled by a human or by a computer (in all cases it was controlled by human wizards), and assessed how aversive they found the negotiation. This was measured subjectively via scales (an 8-item subjective comfort scale and an 8-item friendliness/cooperativeness scale) and objectively by giving participants an opportunity to punish their opponent by reneging on the final deal if they felt dissatisfied. We hypothesize participants will feel more comfortable when CRA is framed as an automated agent.

Hypothesis 2: Participants will try harder to achieve a favorable deal following pedagogical feedback

Participants engaged in two negotiations: first an integrative negotiation that emphasized cooperation and creating value, then a distributive negotiation that emphasized competition and claiming value. The first negotiation was to give all participants a common familiarity with the system before exploring our primary manipulations. After the first negotiation, we manipulated whether participants received pedagogical feedback (about the concept of BATNA and how they could use this information to increase their bargaining power, as described further below) or no feedback (the control condition). We then measured how forcefully they negotiated in the second negotiation through subjective and objective measures.[1] We hypothesize the feedback will increase the effort they invest in the exercise.

4.1 Design

Ninety three participants (52 female) were recruited from an on-line job service and randomly assigned to one of four experimental conditions (described below). Each completed two negotiations: a cooperative/integrative negotiation and a competitive/distributive negotiation. The integrative round matched the structure of Table 1. Participants played Side A, and agents Side B. In the second, distributive round, the agent played side A, and the participant received a payoff similar to side A with the exception that the painting had no value. Note that the actual items differed in round 2 (i.e., chairs, crates of china plates, and a clock), but the values were equivalent to the original items, thus for simplicity, we discuss only the original set of items. 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.

Participants interacted with a male and a female virtual human controlled by the same wizard interface; order of presentation was counterbalanced and found to have no effect on the results presented below. The virtual humans use the same utterances, general dialogue policy and gestures, but differ in appearance and voice. For both virtual human agents, Wizards followed a script. In both rounds, 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. In *both* the integrative *and* distributive rounds, when directly asked their preferences, they would make a distributive offer (in the integrative case, asking for 2 lamps and 1 record; in the distributive case, asking for 2 records and a lamp). In the integrative round, participants always accepted this offer, as they received two of their highest value items (2 records). In the distributive round, such an offer was less attractive and participants negotiated to obtain a better deal for themselves; however, the wizard remained on script and did not budge on this offer.

Two factors were manipulated, resulting in a 2x2 between-subjects design. Participants were randomly assigned to framing condition, where they were either told that the agent was operated by a computer or a human as in [18]. They were also randomly assigned to the feedback condition, where they were given feedback after the first round about how they underperformed when their partner's BATNA is taken into account, or else were given no feedback. Specifically, in the feedback condition, after accepting the offer in the first round, it was pointed out to them that, although they received two

[1] Following standard practice (see [8]), wizards negotiate following a fixed script. This is to avoid the possibility of experimenter bias (e.g., if one participant seems more likeable than another). A disadvantage, however, is that all participants reach approximately the same final deal, making it difficult to judge the impact of pedagogical feedback. Thus we look at time on task and subjective effort to index if they are trying to apply the suggested advice.

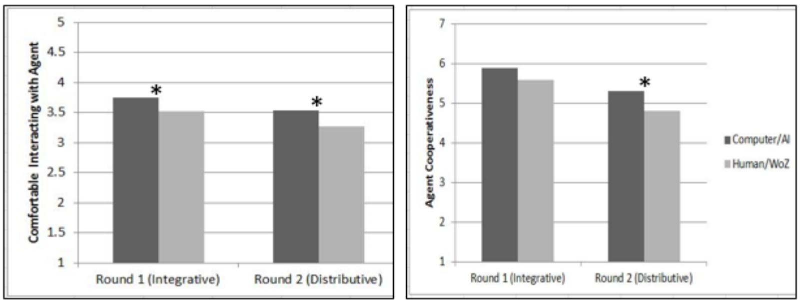

Fig. 2. Partner comfort (2a: left) and perceived cooperativeness (2b: right)

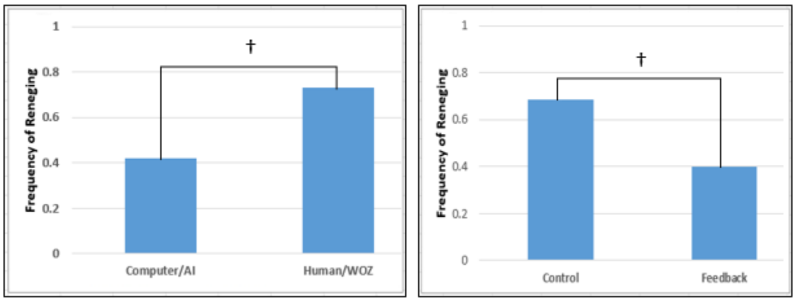

Fig. 3. Frequency of reneging by framing (3a: left) and feedback (3b: right)

of their highest value items (2 records), their partners' BATNA was only one lamp. Participants were encouraged to reflect on how they could have gotten more items in that first negotiation if they had considered how much better the deal was for their partner compared to the partners' BATNA. Those in the control group received no such feedback.

4.2 Results

Ratings of comfort. Participants were more comfortable dealing with a tough negotiator when framed as a computer. Participants responded to eight items on a 1 to 5 scale signaling their agreement that they were "comfortable interacting" with the agent, that "it felt natural to talk" to the agent, and that the agent was "easy to talk to," for example; we averaged these 8 items to index ratings of comfort. As can be seen in Figure 2a, participants felt more comfortable interacting with the agent when they believed it was a computer in both the first, integrative negotiation ($F(1,89) = 5.90$, $p = .02$) and the second, distributive negotiation ($F(1,89) = 5.60$, $p = .02$). There were no effects of or interactions with feedback condition, $Fs < 0.99$, $ps > .32$. This supports our first hypothesis that people are more comfortable dealing with a tough negotiator if they are negotiating with a computer.

Ratings of agent's cooperativeness. Participants viewed their opponent as more cooperative if framed as a computer. Participants rated the agent on eight items using a

1 to 7 scale with bipolar anchors such as uncooperative/cooperative and unfriendly/friendly; we averaged these 8 items to index ratings of cooperativeness. As can be seen in Figure 2b, there was a trend for the participants to rate the agent as more cooperative when framed as a computer in the first, integrative negotiation ($F(1,89) = 1.98$, $p = .16$); they rated the agent to be significantly more cooperative when they believed it was operated by a computer in the second, distributive negotiation ($F(1,89) = 3.91$, $p = .05$). There were no effects of or interactions with feedback condition, Fs $<$ 0.36, ps $>$ .55. This provides additional support for our first hypothesis.

Behavioral measure of commitment to the final deal. Participants were less likely to punish a tough opponent when they were framed as a computer. At the end of the study, participants were told that the negotiator from the first (integrative) round found another storage unit with identical items to those from the second (distributive) round. Participants were offered the option of following through with their agreement with the second negotiator, or taking an identical offer with the first negotiator instead. As switching leaves the participant's profits unchanged but hurts the second negotiation partner, we interpreted this as a measure of dissatisfaction and/or anger with the second opponent. Although there was an error and data was only recorded for 34 participants, among that group there was a marginally significant effect of framing ($\chi^2(1) = 3.32$, $p = .07$). As can be seen in Figure 3a, participants who believed the agent was operated by a computer were less likely to renege on the second deal than those who believed it was operated by a human. We interpret this as an indication that participants found the negotiation less aversive and felt less desire to retaliate, lending further support for our first hypothesis. There was also a trend for feedback ($\chi^2(1) = 2.75$, $p = .10$), such that those in the feedback condition were less likely to renege on their second deal (as can be seen in Figure 3b). This later trend could suggest that participants were more satisfied with their negotiated outcome following such feedback. There was not a significant interaction between framing and feedback ($p > .14$).

Time spent negotiating. Participants tried harder to win when they received pedagogical feedback. Participants who were given feedback after the first round about how they underperformed when their partner's BATNA is taken into account spent significantly longer negotiating in the second round (466 seconds) than those in the no feedback control condition (373 seconds); ($F(1, 89) = 5.58$, $p = .02$). This provides behavioral support for our second hypothesis that feedback will increase effort. This effect was not qualified by framing condition ($F(1, 89) = 0.1$, $p = .91$). The human/computer framing did not significantly impact negotiation time. There were no effects or interactions on time spent negotiating during the first round (Fs $< .69$, ps $> .40$). On the second round, there was a trend for participants to spend longer negotiating with the agent when they believed it to be controlled by a human than when they believed it to be controlled by a computer ($F(1, 89) = 2.77$, $p = .10$).

Appraisal of negotiation. Participants also felt they tried harder when they received pedagogical feedback and felt their performance could be further improved in the future. Using a scale from 0 to 100, they were asked to rate how much effort they put into the negotiation. They also rated the extent to which they felt like they "revealed too much information during the negotiation" on a 1 to 5 scale. Participants who were given feedback after the first round reported expending marginally more effort in the second round (88.0 out of 100) than those in the no feedback control condition (81.4 out of 100); ($F(1, 89) = 3.55$, $p = .06$). Also, participants who received feedback were more aware that they may have still revealed too much information during the second negotiation (2.89 out of 5 vs. 2.40 out of 5); ($F(1, 89) = 6.54$, $p = .01$). All other effects did not approach significance ($Fs < 2.5$, $ps > .12$). These subjective impressions provide further support for our second hypothesis that feedback improves effort.

5 Discussion and limitations

Our findings show that people found it more comfortable to practice tough negotiations with a computer program. When framed as a computer, participants reported more comfort with the negotiation and found their partner more cooperative. When given the opportunity to punish their opponent (by abandoning their negotiated deal for a certain payoff from a different party), they took advantage of this opportunity more often with human opponents, again implying more discomfort with human versus computer opponents. These findings lend support to the notion that students will find virtual human negotiators less aversive than human role-players and this may translate into more motivation to practice.

Participants also invested more effort in the exercise when receiving pedagogical feedback, and this effect occurs regardless of whether the opponent was framed as real or computer. Specifically, participants spend more time negotiating, reported trying harder, and realized they could have improved their performance further when they were explained the concept of BATNA and negotiation power. There was also a trend for less punishment with feedback, suggesting more satisfaction with their negotiated outcome. Again, this supports the potential benefits of virtual human role-players.

There were several limitations to the study. Our manipulation of the nature of the opponent (human versus computer) has some strengths but also limitations. By manipulating "mere belief" about the nature of the opponent we ensure appearance and behavior were controlled (i.e., participants interacted with the identical system but the system was framed as an interaction with a human or computer), however a more comprehensive study would have also included face-to-face interaction as it is possible that negotiating *via* a computer is significantly different than practicing with another student directly [19]. Additionally, we reported only high-level indices of negotiation performance (e.g., time and self-reported effort). Further analysis must be performed to examine how the experimental factors altered negotiation processes. For example, did people make tougher offers when receiving pedagogical feedback? Did they reason more carefully about their BATNA? Answering these questions will require detailed annotation and analysis of the content of the negotiations. Finally, and most importantly, we must verify that these positive findings translate into measureable benefits when negotiating with human opponents.

Acknowledgments: The paper benefited from the feedback of the anonymous reviewers. This research was supported by the Air Force Office of Scientific Research under grant FA9550-14-1-0364 and the U.S. Army. Statements and opinions expressed do not necessarily reflect the position or the policy of the United States Government, and no official endorsement should be inferred

References

1. Brooks, A.W. and M.E. Schweitzer, *Can Nervous Nelly negotiate? How anxiety causes negotiators to make low first offers, exit early, and earn less profit.* Organizational Behavior and Human Decision Processes, 2011. **115**(1): p. 43-54.
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. p. 218-230.
6. Kelley, H.H. and D.P. Schenitzki, *Bargaining.* Experimental Social Psychology. New York: Holt, Rinehart, and Winston, 1972: p. 298-337.
7. Baarslag, T., et al., *Evaluating practical negotiating agents: Results and analysis of the 2011 international competition.* Artificial Intelligence, 2013. **198**: p. 73-103.
8. 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.
9. Thompson, L.L., *Information exchange in negotiation.* Journal of Experimental Social Psychology, 1991. **27**(2): p. 161-179.
10. Van Kleef, G.A., et al., *Power and emotion in negotiation: power moderates the interpersonal effects of anger and happiness on concession making.* European Journal of Social Psychology, 2006. **36**(4): p. 557-581.
11. 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.
12. Pillutla, M.M. and J.K. Murnighan, *Unfairness, Anger, and Spite: Emotional Rejections of Ultimatum Offers.* Organizational Behavior and Human Decision Processes, 1996. **68**(3): p. 208-224.
13. Grecucci, A., et al., *Reappraising the Ultimatum: an fMRI Study of Emotion Regulation and Decision Making.* Cerebral Cortex, 2012.
14. Van Kleef, G.A., *Emotion in Conflict and Negotiation: Introducing the Emotions as Social Information (EASI) Model.* SSRN eLibrary, 2005.
15. Gross, J.J., *Emotion regulation: past, present, future.* Cognition and Emotion, 1999. **13**(5): p. 551-573.
16. Broekens, J., C.M. Jonker, and J.-J.C. Meyer, *Affective negotiation support systems.* Journal of Ambient Intelligence and Smart Environments, 2010. **2**(2): p. 121-144.
17. Carnevale, P.J. and A.M. Isen, *The influence of positive affect and visual access on the discovery of integrative solutions in bilateral negotiation.* Organizational Behavior and Human Decision Processes, 1986. **37**: p. 1-13.
18. Lucas, G.M., et al., *It's only a computer: Virtual humans increase willingness to disclose.* Computers in Human Behavior, 2014. **37**: p. 94-100.

19. de Melo, C., J. Gratch, and P.J. Carnevale, *The Effect of Agency on the Impact of Emotion Expressions on People's Decision Making*, in *Affective Computing and Intelligent Interaction (ACII), 2013 Humaine Association Conference on*. 2013. p. 546-551.
20. 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. p. 47-54.
21. 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.
22. 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.
23. Grosz, B. and S. Kraus, *Collaborative Plans for Complex Group Action.* Artificial Intelligence, 1996. **86**(2): p. 269-357.
24. Gratch, J., et al., *Negotiation as a Challenge Problem for Virtual Humans*, in *15th International Conference on Intelligent Virtual Agents*. 2015: Delft, The Netherlands.
25. Hartholt, A., et al. *All together now: Introducing the virtual human toolkit*. in *Intelligent Virtual Agents*. 2013. Springer.
26. 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.
27. Thiebaux, M., et al., *SmartBody: Behavior Realization for Embodied Conversational Agents*, in *International Conference on Autonomous Agents and Multi-Agent Systems*. 2008: Portugal.
28. 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.
29. 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.