Towards An Autonomous Agent that Provides Automated Feedback on Students' Negotiation Skills

Emmanuel Johnson, Jonathan Gratch
Institute for Creative Technologies
University of Southern California
Los Angeles, California
{ejohnson, gratch}@ict.usc.edu

David DeVault
University of Southern California
Los Angeles, California
devault@usc.edu

ABSTRACT

Although negotiation is an integral part of daily life, most people are unskilled negotiators. To improve one's skill set, a range of costly options including self-study guides, courses, and training programs are offered by various companies and educational institutions. For those who can't afford costly training options, virtual role playing agents offer a low-cost alternative. To be effective, these systems must allow students to engage in experiential learning exercises and provide personalized feedback on the learner's performance. In this paper, we show how a number of negotiation principles can be formalized and quantified. We then establish the pedagogical relevance of several automatic metrics, and show that these metrics are significantly correlated with negotiation outcomes in a human-agent negotiation. This illustrates the realism and helps to validate these principles. It also shows the potential of technology being used to quantify feedback that is traditionally provided through more qualitative approaches. The metrics we describe can provide students with personalized feedback on the errors they make in a negotiation exercise and thereby support guided experiential learning.

CCS Concepts

•Human-centered computing → Human computer interaction (HCI); •Computing methodologies → Artificial intelligence;

Keywords

Negotiation, learning system, feedback system, preference elicitation, virtual humans

1. INTRODUCTION

Negotiation skills are crucial across a wide range of jobs [2, 9], especially in the military[8]. In almost every social and organizational setting, people must negotiate to achieve their goals. Yet many avoid negotiations out of fear or lack of skill and this contributes to income inequality, political gridlock and economic inefficiencies. For example, 93% of women avoid salary negotiations and this behavior serves as a major contributor to unequal pay for women [1]. In politics, poor negotiators have less legislative influence [17]. In the military, soldiers of all ranks are increasingly involved in negotiations, and poor skills can have geopolitical consequences [37]. In the courtroom, over 90% of cases are settled through negotiation before they ever reach trial, and systematic inequities in negotiation ability across lawyers can undermine civil society [11, 31].Yet students graduating from high school, universities and even MBA programs are under-prepared in key interpersonal competencies [30]. As a result, individuals and organizations spend vast sums in remedial training.

Negotiation is typically taught via a mix of classroom lectures and experiential learning (where students receive hands-on experience by applying this knowledge in simulated negotiations with other students or negotiation experts). These skills are taught in professional schools, as part of a business or law degree, and by private consulting firms. For example, as part of a master's degree program in business administration, students might take a semester-long course on negotiation concepts. For those seeking a more cursory introduction, universities and consulting firms offer intensive short courses. Evidence suggests that the experiential aspect of negotiation training is an especially important component of training [3, 28] and is particularly effective at enhancing student motivation and commitment [19]. Yet these simulations add considerably to the expense and logistical constraints in teaching negotiation. In business schools, simulated negotiations are often run by dedicated staff trained who are experts in experiential learning techniques. Within companies, they are overseen by university executive-education programs or high-priced consultants.

Virtual role playing agents offer an opportunity to dramatically reduce the cost and increase access to negotiation training, yet the role of technology in current training practice is surprisingly limited, particularly with regard to the experiential aspects of training. We have all probably experienced the rudimentary web-based lectures on sexual harassment or compliance training. Online negotiation courses use similar techniques to teach concepts, but methods to "gamify" experiential learning have proved elusive, although some rudimentary approaches are beginning to appear in the marketplace [22]. Research within the autonomous agent community has worked to improve this state-of-the-art.

Appears in: *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2017), S. Das, E. Durfee, K. Larson, M. Winikoff (eds.), May 8–12, 2017, São Paulo, Brazil.*

A negotiation can be roughly divided into three phases - preparation, execution and reflection. Although automation can assist with each, much of the research on computer-supported pedagogy has focused on the first two phases. For example, before commencing a negotiation, a negotiator needs to gather information about his or her opponent and decide on a game plan. Systems and techniques have been developed to facilitate these skills [14, 7, 24, 20, 4]. Other research has focused on automated negotiation opponents that allow negotiators to practice skills such as gaining information about their opponent's preferences, building rapport, and adjusting strategies based on opponents' bids [26, 14, 5, 29].

In this paper, we focus on reflection, the third phase of a negotiation. Post-negotiation feedback is one of the most crucial parts of the learning experience [18, 6]. It provides an assessment of a student's progress, relates their performance to fundamental principles of negotiation, and illustrates how a better understanding of these principles can improve their performance in the future. To this end, this paper makes several contributions. First, we review several fundamental principles that underlie effective negotiation performance. Second, we show how to automatically measure a student's adherence to these principles within a corpus of human-agent negotiations. Third, we demonstrate that (most of) these automatic measures successfully predict student performance (i.e., students that follow these principles, as quantified by our measures, achieve better objective outcomes in their negotiations). This provides strong support that these metrics can provide students with meaningful feedback in a way that is more personalized, more immediate, and less ambiguous than current experiential learning techniques.

In the next section, we introduce a set of key principles of negotiation and discuss the state of the art in automated techniques for teaching these skills. Section 3 describes our approach for measuring how well a student follows these principles as they engage in simulated negotiations. Section 4 shows that, indeed, students are better negotiators when they follow these principles. Section 5 illustrates some techniques for visualizing the extent to which a student succeeds or fails in following these principles. We end with discussion of our contributions and limitations of the current work.

2. BACKGROUND

Negotiation is traditionally taught through a mixture of lectures (which introduce fundamental negotiation concepts) and experiential exercises (which allow students to apply these concepts in realistic situations). By design, exercises are crafted to vividly illustrate negotiation concepts, and student success in these exercises typically depends on their ability to translate book knowledge into effective action. Indeed, student's behavior in these exercises becomes grist for classroom discussion and much of the learning, arguably, occurs as instructors highlight the extent to which specific students succeeded or failed in this regard. Autonomous agent technology is playing an increasingly important role in facilitating this form of instruction. Here we review this recent progress. We first review some key negotiation principles. We then highlight three ways that agent technology facilitates experiential learning: in helping with planning in advance of a negotiation, in providing automated negotiation partners to enable practice of key skills, and in providing feedback on student performance.

2.1 Principles of Good Negotiators

Although there is no definitive set of negotiation principles, in this article we draw on the ontology from Harold Kelley's classic 1966 paper on teaching negotiation [23]. In this work, Kelley examined the behavior of students throughout a semester-long negotiation course and articulated principles that distinguished novices from experts. Despite its age, this paper was surprisingly prescient and these principles underlie much of contemporary teaching and have been validated by an extensive body of subsequent empirical work which we will cite as we introduce each principle. Kelley's objective was to understand student behaviors when negotiating under a moderately high incentive condition (student final grades were based on how well they negotiated with their opponents), and how these behaviors changed through a series of negotiations throughout the semester. From this experience, Kelley proposed the following principles that predicted student performance.

(1) Avoid early commitment: In the beginning of a negotiation, negotiators are uncertain about each other's preferences and thus should not prematurely commit to any part of the final deal. Expert negotiators know that to decrease uncertainty they must engage in some form of information exchange extending over a number of rounds. Negotiators expect the information that one's opponent is holding back will be revealed during the course of the negotiation. To avoid early commitment, Kelley introduced several more specific principles. First, negotiators should start with high demands, as this will avoid unnecessary concessions and force the opponent to respond in ways that may reveal their own goals [36]. Second, negotiations should make full use of the available time, as this allows time pressure to help motivate concession-making [32]. Third, negotiators should negotiate multiple aspects of the deal simultaneously rather than one issue at a time, as this allows tradeoffs across issues that may foster more efficient solutions [12, 33].

(2) Make efficient concessions: Negotiators should aim to make as few concessions as needed to reach an agreement and aspire to gain as much value in return for each concession. Concession-making, or at least the illusion of concession-making, is also essential to motivate the opponent to reciprocate. First, negotiators should actively query their opponent about their interests and goals, as this provides a better understanding of the opponent's preferences [34]. By understanding what one's opponent wants, a negotiation can sometimes "grow the pie" and discover concessions that are valuable to the opponent but come at little or no cost to the negotiator. Second, negotiators should consider making multiple equivalent offers to their opponent [21]. (Two offers are considered equivalent if a negotiator is indifferent to which one he/she receives, yet they offer potentially different value to the opponent.) By observing the opponent's inclination towards these different offers, a negotiator can "triangulate" on what their opponent truly wants.

(3) Induce others to make concessions: Whereas the previous principle focuses on how to give up as little as pos-

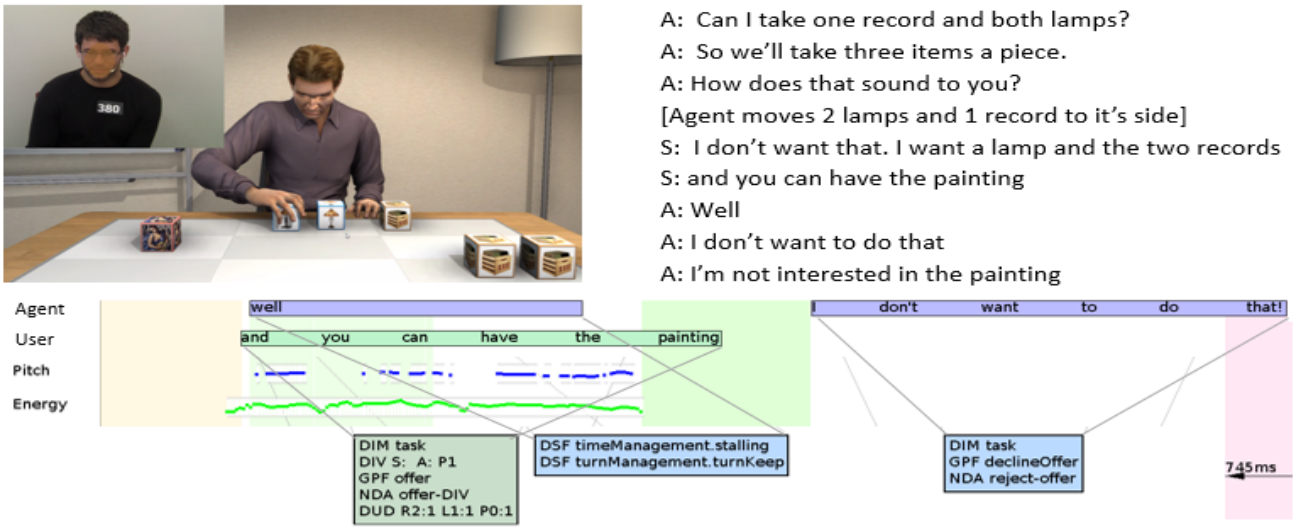

Figure 1: Conflict Resolution Agent with example dialogue acts. (S) indicate utterances made by a user/student while (A) indicate utterances by the agent.

sible when making a concession, this third general principle emphasizes the importance of motivating one's opponent to make concessions in the first place. At its heart, this general principle aims to convince the opponent that the negotiator has reached his or her limit (i.e., that the opponent's current offer is no better than what the negotiator could get elsewhere, referred to as the negotiator's BATNA) [35]. Kelley introduces two techniques to achieve this aim. First, the negotiator could make "negative concessions." Normally, negotiators make greater concessions over time but a negative concession occurs when a negotiator becomes tougher over time (signaling they are becoming impatient). Second, negotiators should reject an opponent's initial offers (under the presumption that the opponent is starting with ambitious initial offers).

(4) Shaping perception of value: In an attempt to gain more utility in a negotiation, negotiators tend to provide information to their opponent which alters the opponent's view of the negotiator's preferences. There are two tactics commonly used to accomplish this goal. First, negotiators tend to misrepresent their preferences [16]. Negotiators may signal to an opponent that they prefer one item while truly in favor of another. Another common tactic is to overplay the value of the concessions one is making. A negotiator may lead an opponent to believe that the concessions they have made are a lot more costly than they are in hopes that their opponent will make much larger concessions.

(5) Assemble information in advance of negotiating: Prior to the negotiation, negotiators want to gain a better understanding of what their opponent wants, reflect on how to approach the negotiation, and plan out their negotiation strategy. To do so, expert negotiators might gather information on their opponent and try to reason about a potential preference model of their opponent. From this they may be better able to determine how to proceed and what sequence of bids to make as well as questions to ask in the hope of maximizing their utility.

2.2 Technology For Preparation

Kelley's fifth point is the importance of private preparation (beginning stage of negotiation). In this stage, negotiators are trying to understand the domain (understanding the issues being discussed and the interests their opponent may not have explicitly stated), how they currently fair with their opponent, and what a reasonable concession might be. As negotiations are complex and sometimes emotional events, agents such as the pocket negotiator [20] have been proposed to serve as a support system to the negotiator. This system helps negotiators better understand the negotiation space and further decrease the cognitive load associated with complex negotiations. By doing so, negotiators are more confident and have a better grasp of the possible conflicts and strategies that underlie their negotiation.

2.3 Technology For Practice

A variety of autonomous negotiation agents have been developed to negotiate in a collection of environments with both people and other agents. Some examples of agent-to-agent negotiators include but are not limited to GENIUS, KBAgent, and QOAgent [38]. There has been a growing interest in negotiating agents that can interact with humans. Out of this line of work, the GENIUS platform has been updated to include a natural language understanding component [38], and a number of other agents have been developed including negoChat [29], BiLat [24], the conflict resolution agent (CRA) [14] and most recently IAGO [27]. These agents allow participants to engage in a form of experiential learning by simulating dyadic negotiation. [26] conducted a study in which the performance of negotiators who trained with a human and those who trained with an agent were compared to determine the effectiveness of each method of training. Participants assumed either the role of a job candidate or employer and engaged in a salary negotiation task. This study found that, even without feedback, autonomous agent systems improve negotiation ability. There is also evidence that people are more comfortable learning from automated agents than other people [14].

2.4 Technology For Feedback

Experiential Learning Theory states that learning is best viewed as a process rather than an outcome. This process is enhanced when learners are able to engage their cognitive abilities, emotion, perception, and behavior [25] and can reflect upon their subsequent actions. Thus, feedback plays a valuable role in the learning process. It informs students of their performance, but also identifies areas of improvement. One advantage of autonomous agents is that they can provide targeted feedback based on objective measures. Feedback, in a traditional classroom setting, typically involves the teacher sharing with students their performance relative to others in the course and the principles of expert negotiators. The limitations of this approach are that feedback is ambiguous, not personalized, and delayed in time. An automated system can provide more instantaneous feedback that is personalized to the specific needs of the student. Most of the current research on negotiation feedback systems has focused on providing feedback before or during the negotiation and very little work focuses on providing feedback at the conclusion of a negotiation. Based on the negotiation principles defined by [16], we have built an autonomous feedback system that can provide targeted feedback at the conclusion of a negotiation. We discuss this system in detail in the next section.

3. AUTONOMOUS FEEDBACK SYSTEM AND METRICS

Our main contribution is to show how to automatically provide students with feedback on their negotiation performance following a role-playing exercise, and specifically, on how well they adhered to the principles outlined in Section 2.1. We explore these techniques within the context of a particular role-playing agent the Conflict Resolution Agent(CRA) [15, 10], however, the methods could be easily adapted to most of the aforementioned negotiation agents, or even to annotated dialogues of student-on-student negotiations. We first describe CRA and how the details of a negotiation are annotated. We then describe how to derive metrics of student performance from such annotated traces of student behavior.

3.1 The Conflict Resolution Agent

CRA allows students to freely converse via natural language with simulated negotiation partners (see Figure 1). The system integrates speech recognition technology, incremental dialogue act understanding, and speech and gesture synthesis to support fairly natural interactions across a range of multi-issue negotiation scenarios. For the purpose of this paper, we focus on the dialogue act annotation which we use to derive our automated metrics.

Figure 1 illustrates a student engaged in the "Auction War" negotiation scenario. In this scenario, the student(S) and agent(A) play the role of antique dealers and must find an agreement on how to divide the contents of an abandoned storage locker containing records, lamps and a painting. Each player (i.e., the student and the agent) is given a relative set of preferences (e.g., the student wants records more than lamps) and a "reservation price" (sometimes called BATNA). A partial example of the negotiation's semantic representations is given in Figure 1, which shows the timing of a short sub dialogue as follows:

Table 1: Conflict Resolution Agent with example dialogue acts

Time(sec)	Speaker	Action		
35	User	Assert	R > 0	TRUTH
44	Agent	Assert	L > P	TRUTH
44	Agent	Assert	L > R	LIE
49	Agent	Assert	R > 0	TRUTH
52	Agent	Question	P > ?	
55	User	Assert	P = 0	LIE
78	User	Offer	R{3:0} L{0:2}P{0:0}	
84	Agent	Decline	declineOffer	
88	Agent	Assert	R > 0	TRUTH
96	Agent	Offer	R{1:2} L{2:0} P{0:1}	
123	User	Decline	declineOffer	
124	User	Offer	R{2:1} L{1:1} P{0:0}	
130	User	Offer	R{2:1} L{1:1} P{0:1}	
132	Agent	Decline	declineOffer	

User: And you can have the painting

CRA: Well

CRA: I don't want to do that

The student's utterance is interpreted as performing a dialogue act with a General Purpose Function (GPF) of making an offer (where the speaker suggests a (partial) division of the items). The content of this dialogue act is a specific division (DIV), "DIV S: A: P1", where the speaker receives nothing and the addressee receives one painting. The annotation also contains an updated deal-under-discussion (DUD), "DUD R2:1 L1:1 P0:1", which summarizes recent DIVs. In this case, the deal-under-discussion is that the speaker gets two records and one lamp, and the addressee gets one record, one lamp, and one painting. The agent responds with a GPF declineOffer, which is a rejection of the user's offer. These dialogue acts are modeled in the ISO standard dialogue act scheme (ISO/DIS-24617-2, 2010) which models dialogue act types using GPFs and Dimension-Specific Functions (DSFs).

In addition to offers and their rejection or acceptance, CRA models a variety of dialogue acts including asserting or requesting information about preferences (e.g., "Are the records more valuable than the painting?"), rapport-building actions, and general dialogue functions such as turn management and meta-statements about the negotiation (e.g., "We should be fair" or "We should negotiate"). Table 1 illustrates a condensed representation of the negotiation acts associated with the negotiation in Figure 1 (note that as CRA knows the actual preferences of both parties, it can annotate these statements with their veracity). For example, in Table 1, the user truthfully asserts they like the records 35 seconds into the negotiation whereas the agent incorrectly responds it like the lamps better than the records.

Table 2: Expert negotiator principles and associated metrics

Principle	Expert Behavior	Metric
Avoid Early Commitment	Use all the available time	Agreement Time: Time of final offer. Predict positive association with earnings
	Make high initial offer	Initial Claim: Fraction of maximum earnings possible claimed by negotiators initial offer. Predict positive association with earnings
	Negotiate multiple issues simultaneously	Single-Issue Offers: Fraction of offers made by negotiator involving a single issue. Predict negative association with earnings
Make Efficient Concessions	Ask about opponent's interests	Unasked Questions: Number of unasked questions required to fully infer opponent's preferences. Predict negative association with earnings
	Make "triangulating" offers	Triangulations: Number of times negotiator proposes multiple offers with equivalent value to him/herself. Count of offers that give equal value to proposer but different value to opponent. Predict positive association with earnings
Induce Opponent Concessions	Make negative concessions	Negative Concessions: Number of times negotiator takes back a previous concession. Predict positive association with earnings
	Reject initial offers	Num Rejections: Number of times negotiator rejects opponents offer. Predict positive association with earnings
Shaping Perceptions of Value	Emphasize the personal cost of concessions	Not yet measured, but involves a variety of behaviors that convey great cost of concessions, misrepresenting own value (e.g.,exaggeration of cost of unimportant issues), play up value of opponents' holding, appealing to fairness
	Undervalue opponent's concessions	
Do Your Homework	Assemble information about reasonable deals to expect	Not addressing this but could build on techniques from pocket negotiator

3.2 Quantifying Negotiation Principles

Based on the principles of expert negotiators presented in Section 2.1, we generated a list of quantifiable metrics, summarized in Table 2. Of the five basic principles outlined in Section 2.1, we address three - avoiding early commitment, making efficient concessions, and inducing opponents to make concessions. We selected these three because they seem to be the most straightforward to quantify and prior research has emphasized their predictive values in determining negotiation outcomes (e.g. [13, 27]).

3.2.1 Avoiding early commitment

For avoiding early commitment, we computed three variables corresponding to Kelley's sub-principles (initial claim, agreement time, and single-issue offers) for realizing this high-level principle. Initial claim measures the fraction of the total outcome space contained in an initial offer. To compute this variable, we look at the negotiator's first offer and compute the maximum utility of the outcome space. From there, the ratio of the value of first offer to the maximum utility is calculated. For example, in Table 1, at 78 sec, we see the user made an initial offer. The user offers to take 3 records and give their opponent two lamps. We compute the value of three records and report what fraction of the total value of the outcome space is captured by three records. This output is represented as initial claim. Initial claim has been shown to positively correlate with earnings [13].

Agreement time measures how long it takes negotiators to reach an agreement. We compute this variable by looking at the time of the last offer made. We see that the user made a last offer at 130 seconds. 130 seconds would be stored as the agreement time. Agreement time has been shown to positively correlate with earnings. Another variable for single-issue offers measures the percentage of offers made involving only one issue. We check offers to see which offers contain only one issue and calculate the percentage of total offer space represented by single-issue offers. This has been shown to negatively correlate with earnings [14, 23].

3.2.2 Making efficient concessions

For making efficient concessions, we measured unasked questions and triangulations. Unasked questions are the number of questions a user can ask to gain more knowledge about the opponent's preferences. This is computed by examining each assert statement from the opponent, deciding what it tells us about our opponent and generating a list of possible questions we could have asked to gain more information about our opponent. For example, in Table 1, the agent makes a total of four assert statements: "L>P" which means "I like lamps more than paintings", "L>R" asserts "I like lamps more than records" and two "R>0" meaning "The record has value to me". From these statements, we know the agent likes the lamps more than the records and paintings. However, it is unclear as to whether or not the agent likes the painting or records more. There are two questions that could be asked to gain clarity on the opponent's preferences: "What do you like the least?" or "do you like records more than paintings?" (Could also be "do you like painting less than records?", we treat both the same). In this case, the variable unasked questions would be set to two.

Unasked questions has been shown to negatively correlated with earnings [23]. Triangulation quantifies the number of exploratory offers the user makes. This counts the number of times a user made multiple offers of equal value for themselves, yet different value for their opponent. We see at times 124 and 130 seconds, the user made two offers of equal value to themselves but two different values to the agent. In particular, at 130 seconds the user's second offer claims the same items for the user as before, but now offers the painting, which has no value for the user, to the agent. This would constitute a triangulation.

3.2.3 *Inducing opponent concessions*

Lastly, for inducing others to make concessions, we compute negative concessions and number of rejections. Negative concessions measures the number of times a negotiator makes a concession and later negates that concession by bidding with a higher utility. From Table 1, we see at time 78 seconds the user bid for three records. They later conceded by bidding for two records and one lamp at 124 seconds. If they were to bid again for items whose values are greater than two records and one lamp (e.g. 3 records), that would count as a negative concession. The number of rejections measures the number of times a user rejects an opponent's offer. This is computed by counting the number of declineOffer dialogue acts. In Table 1, we see that the agent declines two offers and the user declines one. Number of rejections has shown to positively correlate with earnings[14].

4. METRIC VALIDATION

Feedback is most valuable to students when it is visibly grounded in reality. For example, research has shown that making a high initial offer is important to gaining a good outcome [38], but this fact only has pedagogical value to the extent that a student can clearly see a meaningful impact of this factor on negotiation outcome in the classroom exercises. When a student sees that their own or other students' outcomes are significantly impacted by their initial offers, the lesson becomes vivid and meaningful. Thus, to test the validity of our metrics, we examined how well they predicted actual performance in a negotiation task. We obtained a dialogue corpus from CRA. We tested whether our metrics do indeed discriminate between good and bad negotiators.

4.1 Negotiation Corpus

We ran a study with the Conflict Resolution Agent and validated our metrics against this data. The corpus (described in [14]), consisted of 192 negotiations but limited our analysis to 159 as recording failure precluded complete annotations of some sessions. Each participant engaged in some variant of the "Auction Wars" negotiation (illustrated in Figure 1). Participants were incentivized to do well (they received tickets in a cash lottery based on their performance) and were given 15 minutes to reach a deal with the agent. The behavior of the agent was controlled by a human wizard following a pre-determined interaction policy.

The resulting dialogues were manually transcribed and annotated by two expert negotiators using the annotation scheme described in Section 3.1[1]. In addition to the annotated acts, each negotiation contains an objective measure of the student's performance. Before negotiating, each student received a payoff matrix which described the earnings they would receive (in terms of lottery tickets) based on how good a deal they could negotiate. Thus, the number of lottery tickets earned serves as our objective measure. Additionally, all participants were asked to guess how much value the agent assigned to each item to index how well they understood their opponent's interests.

[1]Although an automated version of the agent exists, we could only obtain manual annotations for this study. These are likely to be more accurate that the annotations possible with a fully-automated system.

Table 3: The relationship between negotiation metrics and lottery tickets within the CRA Corpus

Negotiation Metrics	Lottery tickets
Avoid Early Commitment:	
Initial Claim	.47**
Agreement Time	.59**
Single-Issue Offers	-.16*
Make Efficient Concessions:	
Unasked Questions	-.35**
Triangulation	.36**
Induce Opponent Concessions:	
Number of Rejections	.53**
Negative Concessions	.26**

Note: *p <.05 **p <.01

Figure 1 illustrates a partial example of one of these negotiations. In this case participant number 380 is negotiating over records, lamps and a painting. The participant wants the records the most and does not care about the painting, whereas the agent wants the lamps the most and the painting the least.

4.2 Analysis and Results

We examined the relationship between our metrics and outcomes by conducting Pearson's correlations between each metric and lottery tickets. The results, shown in Table 3, illustrate that our metrics are indeed indicative of a negotiation outcome and are aligned with the outcomes of the results from the BiLat project [7]. Concerning how participants avoid early commitment, we see that users with high initial offer tend to gain more lottery tickets ($r = .47$, $p < .001$). This supports the idea that participants who make strong initial claims end up getting more from the negotiation. Additionally, users with higher agreement time gained more tickets ($r = .59$, $p < .001$). Thus, users who utilized more time in a negotiation tended to get more. This could be due to more offers and information exchanged overtime or progressive concession by opponents. Moreover, users who made less single-issue offers gain more lottery tickets ($r = -.16$, $p = .05$). This means participants who made more multi-issue offers got more in a negotiation.

In regards to making efficient concessions, the more unasked questions a user had about opponent's preference, the less

lottery tickets they received (r = -.35, p < .001). This signals that those who knew more about their opponent were better able to position themselves to win the negotiation. This could be due to improved understanding of their opponent's preferences, but it is inconclusive. Also, we see that the more triangulation offers a user made, the more lottery tickets they received (r = .36, p < .001). Lastly, in terms of inducing opponent concessions, users with a higher number of rejections also claimed more lottery tickets (r = .53, p < .001). That is, users who were less likely to accept an opponent's offer performed better. Furthermore, users who made more negative concessions earned a greater number of lottery tickets (r = .26, p = .01). Thus, negotiators who increase their offers' value over time gain more than those who conceded or stayed the same.

5. VISUALIZING FEEDBACK METRICS

Up to this point we have shown that it is possible to quantify student adherence to negotiation principles and that lack of adherence undermines a student's ability to obtain good objective outcomes. One possibility is to simply provide this information to instructors. This could help them find good discussion examples, especially in a large or online course. However, we would also like to explore the potential of automated methods of feedback delivery.

Figure 2 illustrates some of the detailed metrics we can provide on student performance and how they emphasize pedagogical points. Figure 2a helps illustrate the principle of efficient concessions. The chart illustrates the percentage of information the user has obtained about their opponent. It also signals what students could have asked to gain more complete information. Figure 2b emphasizes the principle of avoiding early commitment. The graph visualizes the sequence of offers made by parties in a negotiation as a function of time (this illustrates the same negotiation that was shown in Figure 1 and Table 1). The blue line (marked User) indicates the offers made by the student. The student made an initial offer 78 seconds into the negotiation, followed by a concession at 124 seconds (the y-axis represents the fraction of value claimed by each party with 1.0 representing an offer in which the party's value is maximized). The green line (marked Agent) illustrates the sequence made by the virtual role player: e.g., it makes a counteroffer 96 seconds into the exercise. Finally, the red line (marked Class) shows the initial and final offers, on average, made by all students that participated in the exercise.

From this graph, we can see that the student started with a strong initial offer compared with the class (claiming 75% of their maximum possible value compared to only 30% on average) and made this offer at about the same time as the average student. We can also see that the student negotiated a better than average deal and also used far more of the available time before concluding the negotiation (the final offer occurred 270 seconds into the exercise whereas the average student concluded the negotiation after only 153 seconds). Indeed, the positive correlation between initial and final offers is typically observed in these classroom exercises and serves as an important teaching point. Finally, the fact that the combined value of the deals for each side exceeds 1.0 emphasizes that the student discovered integrative potential and created a win-win solution.

(a) Opponent's Preference Knowledge

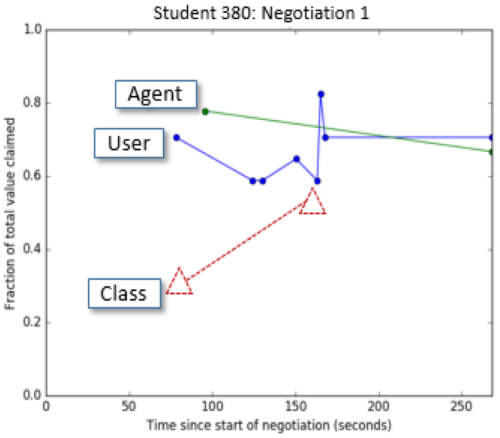

(b) User and Agent Concessions

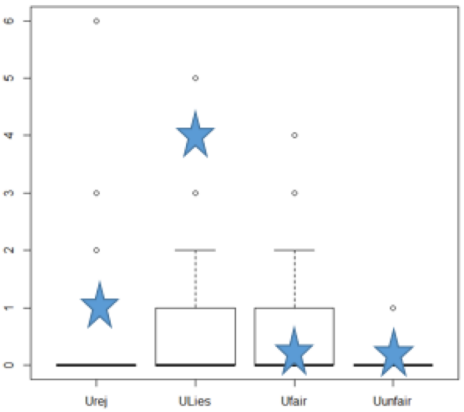

(c) User Relative Scores Across Metrics

Figure 2: Example visualization of students negotiation behavior

Finally, the box and whisker plot in Figure 3c is one efficient way to illustrate a student's performance vis-à-vis the remainder of the class. This figure shows how the entire class performed across a variety of metrics. Stars show the student how he or she behaved relative to the class. In this figure, it illustrates that the student lied far more often than his classmates. Moving forward, more research will be needed in finding the most intuitive and effective ways of conveying information about students' absolute and relative performance and how this information might best motivate and inform subsequent learning.

6. DISCUSSION AND CONCLUSIONS

In this paper, we have outlined some of the key principles taught to novice negotiators and shown how they can be quantified autonomously given manually annotated data. We have tested our quantitative metrics using experimental data and shown that they correlate with negotiation outcomes in ways that underscore the key negotiation principles. We have also illustrated how one may be able to use objective metrics to provide effective feedback to help students develop their negotiation abilities.

There are a few limitations in this work which we aim to address in subsequent research. First, we have focused only on the mechanics of a negotiation (offers, positions, information exchange), which encompass three of Kelley's five principles, rather than the more qualitative factors (how negotiators use misrepresentation or emotion to shape impressions, establishing rapport and relationship building overtime). We hope to expand our metrics to include Kelley's other principles. Second, our analysis shows how these metrics can be quantify and displayed to users, but we do not compare our approach for teaching negotiation principles with other traditional methods. Moving forward, we hope to compare the effective of our approach to more traditional negotiation training methods. Finally, this paper focuses on one method of providing feedback to the learner (on screen visualization). In the future, we would like to explore other modalities such as audio to determine the best methods of providing feedback to a learner.

7. ACKNOWLEDGMENTS

We are grateful for the input of the anonymous reviewers. This work is supported by the US Army Research Laboratory and a National Science Foundation Graduate Research Fellowship. The content does not necessarily reflect the position or the policy of any Government, and no official endorsement should be inferred.

REFERENCES

[1] L. Babcock and S. Laschever. *Women don't ask: Negotiation and the gender divide*. Princeton University Press, 2009.

[2] W. L. Bedwell, S. M. Fiore, and E. Salas. Developing the future workforce: An approach for integrating interpersonal skills into the mba classroom. *Academy of Management Learning & Education*, 13(2):171–186, 2014.

[3] M. E. Bredemeier and C. S. Greenblat. The educational effectiveness of simulation games: A synthesis of findings. *Simulation & Games*, 12(3):307–332, 1981.

[4] J. Broekens, M. Harbers, W.-P. Brinkman, C. M. Jonker, K. Van den Bosch, and J.-J. Meyer. *Virtual Reality Negotiation Training Increases Negotiation Knowledge and Skill*, pages 218–230. Springer Berlin Heidelberg, Berlin, Heidelberg, 2012.

[5] J. Broekens, C. M. Jonker, and J.-J. C. Meyer. Affective negotiation support systems. *Journal of Ambient Intelligence and Smart Environments*, 2(2):121–144, 2010.

[6] M. Conley Tyler and N. Cukier. Making it fresh: ideas for teaching negotiation skills. *ADR Bulletin*, 9(4):2, 2006.

[7] M. Core, D. Traum, H. C. Lane, W. Swartout, J. Gratch, M. Van Lent, and S. Marsella. Teaching negotiation skills through practice and reflection with virtual humans. *Simulation*, 82(11):685–701, Nov. 2006.

[8] N. R. Council et al. *Human behavior in military contexts*. National Academies Press, 2008.

[9] N. R. Council et al. *Exploring the intersection of science education and 21st century skills: A workshop summary*. National Academies Press, 2010.

[10] D. DeVault, J. Mell, and J. Gratch. Toward natural turn-taking in a virtual human negotiation agent. 2015.

[11] T. Eisenberg and C. Lanvers. What is the settlement rate and why should we care? *Journal of Empirical Legal Studies*, 6(1):111–146, 2009.

[12] S. Fatima, M. Wooldridge, and N. R. Jennings. Optimal negotiation of multiple issues in incomplete information settings. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems - Volume 3*, AAMAS '04, pages 1080–1087, Washington, DC, USA, 2004. IEEE Computer Society.

[13] A. D. Galinsky and T. Mussweiler. First offers as anchors: the role of perspective-taking and negotiator focus. *Journal of personality and social psychology*, 81(4):657, 2001.

[14] J. Gratch, D. DeVault, and G. Lucas. The Benefits of Virtual Humans for Teaching Negotiation. In *Proceedings of the 16th International Conference on Intelligent Virtual Agents (IVA), 2016*, Los Angeles, CA, Sept. 2016. Springer.

[15] J. Gratch, D. DeVault, G. M. Lucas, and S. Marsella. *Negotiation as a Challenge Problem for Virtual Humans*, pages 201–215. Springer International Publishing, Cham, 2015.

[16] J. Gratch, Z. Nazari, and E. Johnson. The misrepresentation game: How to win at negotiation while seeming like a nice guy. In *Proceedings of the 2016 International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '16, pages 728–737, Richland, SC, 2016. International Foundation for Autonomous Agents and Multiagent Systems.

[17] R. L. Hall. Measuring legislative influence. *Legislative Studies Quarterly*, pages 205–231, 1992.

[18] J. Hattie and H. Timperley. The power of feedback. *Review of educational research*, 77(1):81–112, 2007.

[19] E. Henry, J. Fowlie, and M. Gordon. *Using Games and Simulations in the Classroom: A Practical Guide for Teachers*. Routledge Studies in the Modern World Economy. Taylor & Francis, 2013.

[20] K. V. Hindriks 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*, pages 47–54. ACM, 2008.

[21] V. Husted Medvec, G. J. Leonardelli, A. D. Galinsky, and A. Claussen-Schulz. Choice and achievement at the bargaining table: The distributive, integrative, and interpersonal advantages of making multiple equivalent simultaneous offers. 2005.

[22] K. M. Kapp. *The gamification of learning and instruction fieldbook: Ideas into practice*. John Wiley & Sons, 2013.

[23] H. H. Kelley. A classroom study of the dilemmas in interpersonal negotiations. *Strategic interaction and conflict*, 49:73, 1966.

[24] J. M. Kim, R. W. Hill Jr, P. J. Durlach, H. C. Lane, E. Forbell, M. Core, S. Marsella, D. Pynadath, and J. Hart. Bilat: A game-based environment for practicing negotiation in a cultural context. *International Journal of Artificial Intelligence in Education*, 19(3):289–308, 2009.

[25] A. Y. Kolb and D. A. Kolb. Experiential learning theory. In *Encyclopedia of the Sciences of Learning*, pages 1215–1219. Springer, 2012.

[26] R. Lin, Y. Oshrat, and S. Kraus. Investigating the benefits of automated negotiations in enhancing people's negotiation skills. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*, pages 345–352. International Foundation for Autonomous Agents and Multiagent Systems, 2009.

[27] J. Mell and J. Gratch. Iago: Interactive arbitration guide online. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, pages 1510–1512. International Foundation for Autonomous Agents and Multiagent Systems, 2016.

[28] H. Movius. The effectiveness of negotiation training. *Negotiation Journal*, 24(4):509–531, 2008.

[29] A. Rosenfeld, I. Zuckerman, E. Segal-Halevi, O. Drein, and S. Kraus. Negochat: a chat-based negotiation agent. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, pages 525–532. International Foundation for Autonomous Agents and Multiagent Systems, 2014.

[30] R. S. Rubin and E. C. Dierdorff. How relevant is the mba? assessing the alignment of required curricula and required managerial competencies. *Academy of Management Learning & Education*, 8(2):208–224, 2009.

[31] H. V. Samborn. Vanishing trial, the. *ABAJ*, 88:24, 2002.

[32] A. F. Stuhlmacher, T. L. Gillespie, and M. V. Champagne. The impact of time pressure in negotiation: A metaâĂŘanalysis. *International Journal of Conflict Management*, 9(2):97–116, 02 1998.

[33] L. Thompson. An examination of naive and experienced negotiators. 1990.

[34] L. L. Thompson. Information exchange in negotiation. *Journal of Experimental Social Psychology*, 27(2):161 – 179, 1991.

[35] S. B. White and M. A. Neale. Reservation prices, resistance points, and batnas: Determining the parameters of acceptable negotiated outcomes. *Negotiation Journal*, 7(4):379–388, 1991.

[36] G. Whyte and J. K. Sebenius. The effect of multiple anchors on anchoring in individual and group judgment. *Organizational Behavior and Human Decision Processes*, 69(1):75 – 85, 1997.

[37] W. Wunderle. How to negotiate in the middle east. *Military Review*, 2007.

[38] I. Zuckerman, A. Rosenfeld, S. Kraus, and E. Segal-Halevi. Towards automated negotiation agents that use chat interfaces. In *The sixth international workshop on agent-based complex automated negotiations (ACAN)*, 2013.