

Multi-party, Multi-issue, Multi-strategy Negotiation for Multi-modal Virtual Agents

David Traum¹, Stacy Marsella², Jonathan Gratch¹, Jina Lee², and Arno Hartholt¹

¹ Institute for Creative Technologies
University of Southern California
Marina del Rey, CA, USA traum@ict.usc.edu

² Information Sciences Institute
University of Southern California
Marina del Rey, CA, USA

Abstract. We present a model of negotiation for virtual agents that extends previous work to be more human-like and applicable to a broader range of situations, including more than two negotiators with different goals, and negotiating over multiple options. The agents can dynamically change their negotiating strategies based on the current values of several parameters and factors that can be updated in the course of the negotiation. We have implemented this model and done preliminary evaluation within a prototype training system and a three-party negotiation with two virtual humans and one human.

1 Introduction

In the most general case negotiation can include trade-offs between multiple issues, can involve multiple parties, each with their own agendas, and can be negotiated through multiple modalities, including speech and face to face bodily communication. Moreover, the parties involved need not maintain a constant position, but can dynamically vary their goals, strategies to achieve those goals, and agenda for carrying out those strategies as the negotiation proceeds.

In this paper, we describe work on pushing the frontier of what can be accomplished by virtual agents in negotiation with humans and other virtual agents. We extend previous work in several directions. In [1] we presented a model of team negotiation involving multiparty dialogue. This model allowed virtual humans to engage in multimodal negotiation over multiple options and discussion among multiple agents as to which options were best for satisfying the shared goals. However, it did not allow for more general negotiation, including a range of different utility valuations and relationships among the agents, including adversarial and more neutral as well as team members. Several factors were taken into account, including the roles of agents, the previous dialogue history, and the utility calculations, but there was only a single fixed negotiation algorithm mapping the value of these factors to a negotiation move.

In [2], we extended this model to handle other kinds of relationships, including adversarial negotiation. Agents could assess their own view of utilities of actions as well as utilities of a negotiating partner. A model of trust was created, using factors of credibility, solidarity, and familiarity. Agents had a choice of strategies to select,

depending on factors including utility and controllability. However, this model was also limited in a number of ways. First, it only handled negotiation of whether or not to select a single action, rather than allowing a broader set of possible decisions. Strategies were always with respect to this single action. There were only two parties involved in the negotiation, the agent and one other (e.g., a human trainee).

In this paper we describe work that takes the next step toward a fully general and human-like model of negotiation. We combine the strengths of both previous models, as well as some further extensions. We allow negotiation over simultaneous courses of action. The trust model is extended to refer to specific individuals, rather than a general trust level. Strategies are made specific to each possible issue of negotiation, and one can consider different strategies for each issue. Moreover, we have expanded the set of strategies that the agents may choose from.

The primary purpose of the negotiation model is to enable the virtual humans to act as role players in a training environment, in which a human trainee can practice different styles and tactics of negotiation and analyze the results. Things are set up so that the trainee must generally balance three different goals in order to be successful at more difficult negotiations:

Solve problems - the most basic matter is figuring out a mutually acceptable solution, just based on the utilities for all the participants. All things being equal, people will act rationally and agree to proposals that are in their own interest. The trainees must be able to go beyond their initial starting points and see how to make a solution attractive to others, e.g. by offering additional resources, committing to important actions, and removing obstacles. The trainee must also consider alternative plans that might lead to a win-win or compromise situation that is an adequate even if not optimal solution.

Gain Trust - generally, all things are *not* equal. The trainee must also work on an interpersonal level to develop and maintain the trust of other participants. With our model this involves working at three aspects:

Familiarity the trainees must show that they know how to behave appropriately in this situation, for this culture, including polite pleasantries, and adhering to norms of topic management.

Credibility The trainees must be truthful and say things that are believable, and also stand by their word and follow through on promises.

Solidarity The trainees must show that there is some alignment in goals between themselves and agents – that they want some of the same things.

Manage Interaction It is also important for the trainees to properly manage the interaction. By properly setting the agenda and controlling the topic progression they can lead to more successful results (assuming they are solving problems and gaining trust). They still must be reactive to the concerns that the agents express and not be too heavy-handed and unilateral. On the other hand, if they lose control of the agenda, the agents may agree on an undesirable plan or refuse to consider other options.

This model has been implemented in our virtual humans and in our current test scenario controls the behavior of two different virtual humans in a three-party negotiation (with a human user) in a prototype negotiation-training application. The virtual humans

recognize speech and a limited set of gestures and body postures, and produce speech and gestures. The architecture of the whole system is described in [3]; in this paper we focus on the cognitive aspects of the negotiation model and multimodal realization of negotiation strategies.

In section 2, we describe the extensions to the multiparty dialogue model. In section 3, we describe the negotiation strategies and how they typically affect behavior. In Section 4 we give more details on how the strategies are implemented, including relevant factors to consider, selection criteria, and how strategies are realized. In section 5 we give some examples of how the dialogue model and strategies are used to influence the behavior of virtual humans in negotiation with a person. Finally, in section 6 we conclude with a discussion of related and future work.

2 Multi-modal Multi-party Dialogue Model

The negotiation is carried out in the context of a multi-party meeting with multiple individuals involved in a (virtual) face to face setting. The agents obey the norms of conversation, including deciding who or what to look at, how to orient their bodies, which posture to adopt, when to speak or listen, and what to say. As outlined in [4, 1, 5], the dialogue model uses the information-state approach to dialogue management [6], with multiple layers of interaction. Each layer consists of information state components and dialogue acts that change values of the components. Decisions of listening, processing utterances, and speaking are made asynchronously, and the agents have the capacity to both respond to communications from other human and virtual agents, and to initiate communication based on their internal state and decisions. There are specific representations of each conversation the agent is aware of, with its conversational state. We have extended previous work with a tighter coupling between the dialogue modelling, emotion modelling, and non-verbal expression. In the rest of this section we briefly describe some of these extensions, in particular, the gaze and listener reactions model, and the use of motivations for tailoring output, and tracking of focus and strategy.

The gaze model [7] has been extended to include different styles of gaze depending on the reason for the gaze. There is also much more non-verbal feedback during the listening and processing of utterances of others, depending on whether the agent agrees or disagrees with what is said and trusts or does not trust the speaker. Specifically, the listener's dialogue model informs its non-verbal behavior generation process [8] whether the speaker is agreeing or disagreeing with a prior speaker and whether the listener itself agrees or disagrees with that stance. If the listener agrees, it may nod while the other speaks. On the other hand, if it disagrees, the non-verbal behavior generator will select other behaviors, such as lowering its head and frowning (lower the brow) or pulling back its head and lowering its eyebrows (inner and outer brow are raised). The particular behaviors chosen depend on the cultural, physical, and personal features of the specific agent. For example, elderly listeners may nod more slowly or different agents may use more idiosyncratic behaviors.

Output motivations are also used to guide the generation of both verbal and non-verbal negotiation behavior, depending on not only the main message to be expressed, but also the reason for saying it and the issue and negotiation strategy that motivate

that reason [9]. This allows the agent's text and behavior generation components to subtly tailor the output, e.g. distinguishing between pointing out a plan flaw to defeat a proposal vs bringing it up as an issue to be addressed and overcome.

Agents actively compute motivations for strategies of all of the options under consideration. The current topic of conversation is tracked based on understanding of the content of utterances and the dialogue history. References to the topic or any constituent actions and states contribute to selecting or maintaining the current topic. The agents dynamically activate the strategy for the current topic from among the motivations for all issues under discussion.

3 Multi-party Negotiation Strategies

In previous work [2], our negotiation strategies were based on orientations to the negotiation [10, 11]. In a multi-party situation, these orientations are not so straightforward, as one must distinguish the attitudes about the negotiated items, the individual participants, the whole group, and subgroups. Thus one may wish to avoid the whole negotiation, or just one issue. One may wish to avoid the whole group interaction, or just one participant. One may feel distributive or integrative with individuals, the whole group, or subgroups (coalitions). One may simultaneously be integrative toward some while distributive towards others, and wanting to avoid yet others.

We defer these issues for the time being, and focus on specific strategies that take some of these orientations into account but focus on concrete objectives rather than the orientations that lead to them. Strategies have applicability conditions, tactics to carry out the strategies, and behaviors (verbal and non-verbal) to communicate the external impressions that are appropriate for that strategy. We describe these features and the strategy informally in this section and discuss the formal and implementational details in the next. We have so far implemented the following negotiation strategies:

Find Issue This strategy is appropriate in the case where there is a negotiation meeting currently occurring, but there is no issue that is a current topic of negotiation. The possible moves include requesting a topic of negotiation from another agent (human or virtual), proposing a topic, or proposing constraints on topic selection. In addition to the kinds of gestures associated with request and proposal moves, this strategy might be signaled non-verbally by a more open body posture. Open postures place no physical barriers between conversants, for example, no crossing of arms or legs, indicating a willingness to participate.

Avoid This strategy is appropriate in the case when there is no topical issue or the focused issue is undesirable but seen as avoidable. The moves include talking off-topic, e.g. small talk, trying to leave the meeting or the topic, or switching the topic to another issue. Avoid can be signaled non-verbally by a more closed, negative and defensive posture. For example, crossed arms while standing is a sign of defensiveness and protection (e.g., [12, 13]) in many cultures.

Attack This strategy is appropriate in the case where the topic is seen as not avoidable and having negative utility, with little potential for improving the utility. It is an assumed bad outcome. The moves include stating flaws in the issue under discussion - negative outcomes that are likely, or pre-conditions that are not met, attempts to

propose alternative, better issues, and ad-hominem comments about the advocates of the issue. Attack can be signalled non-verbally by an open but more aggressive, dominant posture. For example, arms akimbo (on hips) while standing is a sign of disliking, dominance and even anger (e.g., [14, 13]).

Negotiate This strategy is appropriate in the case where it is not clear what the outcome of adopting the issue will be – there is a potential for either negative, neutral or positive results, depending on how the plan is carried out, and whether all individuals involved will do their parts. Here, the agent is not necessarily for or against the issue, but willing to consider whether it can be made to work or not. Moves in this strategy include stating flaws, as for the attack strategy, but also proposals of solutions to fix the flaws, and bargains that would give up some utility on some aspects while gaining utility on others. Conditional commitments, contingent on the commitment of others and fixing of flaws are also appropriate. Because of the potential outcome is unclear, we currently associate a mixed non-verbal signal with this strategy, for example one hand on hip.

Advocate This strategy is appropriate when one has good reason to believe that the outcome of the issue will have positive utility. The moves involved include proposing plans to bring about the outcome, proposing solutions or ameliorations to flaws that have been introduced, and offering commitment to the issue or its component parts. Because of the potential outcome is positive, we signal this strategy with an open, relaxed posture.

Success This strategy involves the follow-through of a successful mutual commitment to an issue - it may involve formalizing remaining details of how to carry it out, as well as friendly disengagement from the meeting. Because a positive outcome has been achieved, we currently associate an open, relaxed posture with this strategy.

Failure This strategy follows from the commitment against a course of action. It involves disengagement from the issue and possibly the meeting, seeing the issue as settled. The agent may have either positive or negative emotions associated with the failure and the non-verbal behavior may need to vary accordingly.

4 Implementing the Strategies

In this section we describe in more detail how the strategies described in Section 3 are implemented. In 4.1, we describe the factors that are used to decide which strategies to adopt. We then describe these factors in more detail in 4.2 and 4.3. We describe how strategies are selected in 4.4 and how they are performed in 4.5.

4.1 Factors in Strategy Selection

There are several factors that are examined when deciding which strategy to chose:

Topic Foremost is the question of which issue is the topic of the current conversation. If there is no topic, then only the find-issue and avoid strategies are applicable. If there is a current topic, then appraisals of plans related to this issue will be the source of further decisions.

Control This is the agent's estimated ability to control the discussion about the topic under discussion. Control is a pre-requisite for successful avoidance.

Utility This is the agent's calculation of how good the outcome will be if the issue is carried out, using current assumptions about plans and likelihoods of effects holding and commitments being carried out. More details on how these are calculated are given in Section 4.2. An agent who thinks the an issue has positive utility will generally be an advocate. There is also a consideration of absolute utility (positive or negative) vs. relative utility compared to other options.

Potential This is the agent's estimation of how good the utility can get, assuming that everyone will "do the right thing". For issues with negative utility, the potential is the principal factor in deciding whether to attack or negotiate.

Trust how much does the agent trust the other agents in the negotiation? With low trust an agent will not want to continue and commitments will not increase the estimated probability that actions will succeed. With high trust, other agents can be believed.

Commitment This involves whether participants have committed themselves for or against issues. Mutual commitment is generally a pre-requisite of the success strategy, while negative commitment is a result of the failure strategy. There are also commitments to actions that support one or another of the issues, which can lead to different predictions of utility and potential of that issue.

4.2 Multi-issue Utilities

The ability of our agents to negotiate with humans and other agents stems from their understanding of the goals of each party, the actions that can achieve or thwart those goals, and the commitments and preferences agents have towards competing courses of action. To provide this understanding, our agents use domain-independent reasoning algorithms operating over a general partial-order plan representation: see [15, 1]. Plans provide a concise representation of the causal relationship between actions and agents' goals, including causal links and causal threats between the plans of different agents. The representation includes decision theoretic information to represent the perceived utility of different goals and their likelihood of satisfaction. Finally, the representation includes a simplified theory of mind, allowing agents to represent and reason about the beliefs, intentions and preferences of other agents.

A key aspect of multi-party negotiation involves discussion of alternative ways to achieve goals. To support such negotiation, agents reason about alternative, mutually exclusive courses of action (plans) for achieving goals, and incorporate a general decision-theoretic method for evaluating the relative strengths and weaknesses of different alternatives. Strengths of a plan include states of positive utility that would result from the plan's execution, weighted by their probability of attainment. Weaknesses include resulting negative utility states and basic flaws that might block the plan's execution. For example, a plan may contain unsatisfied preconditions that would require (negotiated) help from other agents to satisfy. It may also contain causal threats, as when the expected actions of another agent might block the plan's successful execution.

Agents also reason about the *potential* strengths of a plan, meaning the expected utility of beneficial effects assuming that any potential flaws are successfully resolved. Consider, for example, the situation where Bob has to borrow Mary's car in order to

buy groceries. The likelihood that this plan succeeds depends on the likelihood of Mary performing the “lend car” action. Initially, Bob may have some *a priori* probability that Mary will lend the car. If Mary verbally commits to lending the car, this probability is likely to increase, although this depends on Mary’s perceived trustworthiness as a negotiation partner. The expected benefit of the plan includes the current estimate that Mary will perform this action, accounting for any stated commitments and the current measure of trust Bob has for Mary, whereas the potential benefit calculation assumes Mary will help with probability 1.0.

The (potential) strengths and weaknesses of a plan serve as talking points for the negotiation and criteria for moving between negotiation stances. Strengths are points that should be emphasized when advocating a certain course of action whereas weaknesses are objections that can be raised. The relative magnitude of strengths and weaknesses of a course of action inform its strategy for negotiation. A plan with more severe weaknesses than strengths, and no potential for improvement should be avoided or fought against. A plan with some positive potential might merit negotiation.

4.3 Models of other Agents

As well as calculating ones own beliefs, goals, intentions, and computations of the expected utilities of various actions and outcomes, the agents also engage in (limited) reasoning about the mental states of others. They track the beliefs of others (which might be positive, negative, or unknown), intentions to act, and utilities of others. These contribute to the estimation of the likelihood of other agents to act in particular ways, and thus the estimated utility of a course of action. Trust of each other agent is also computed as described in Section 1.

There are also models of the interactional structure between the agents. Part of this is the dialogue model, discussed in section 2. The dialogue model tracks the topic, as discussed above. Control is calculated using a heuristic that an agent has control over (avoiding) a topic if it has not been referred to more times by other agents than a threshold amount. Commitments are also calculated on the basis of dialogue utterances: if an agent makes a (grounded) assertion or promise, this leads to a social commitment.

4.4 Choosing Strategies

Table 1 shows the applicability conditions for choosing among the strategies. In general, only a subset of the factors are relevant for any given strategy. Also, there is some overlap in the set of applicable strategies. Our initial algorithm chooses deterministically, preferring first to find the topic, and then decide (based on the utility, potential utility, and control) whether to avoid, attack, negotiate, or advocate. Once commitments have been established the agents follow the success or failure strategy for that topic.

4.5 Dialogue Realizations of Strategies

Once the strategy has been chosen, the agent will have the option of selecting from a number of moves that go with the strategy. These moves are in competition in the

	topic	control	utility	potential	trust	commitment
find-issue	-				some	
avoid		+	-		some	
attack	+	-	-	-	some	
negotiate	+		-	+	some	
advocate	+		+		some	
success	+				moderate	mutual
failure	+				very low	negative

Table 1. Choosing Negotiation Strategies based on Factors.

agent’s decision space with other kinds of actions. These include dialogue actions, such as giving grounding feedback and addressing questions, as well as non-dialogue actions such as emotion reasoning and acting in the virtual world.

For **Find-issue**, the two main actions are requesting a topic from the meeting initiator, and proposing a topic. The initiative parameter for the agent determines this choice. With no initiative, the agent will not bring up the topic at all. With a medium level, the agent will ask for the topic. If the agent has high initiative and control, it will introduce a high utility topic.

In the **Avoid** strategy, the possible move types are:

- change topic to high utility issue
- talk about non-issues (ad hominem, small-talk)
- disengage from meeting

The agent will prefer to change the topic if there is a good one, otherwise will try non-issue talk and if that does not work, but there is still some control, will try to leave.

For the Attack strategy, the agent will choose either ad-hominem attacks, e.g., blaming the topic-initiator for the problems, or pointing out flaws with this issue, that have been identified as described in section 4.2. Flaws include pre-conditions that are not likely to be met, negative outcomes, and lack of necessary commitments from participating agents. The agents also compare the plans unfavorably with higher-utility options. No possible solutions are presented to the flaws.

In the Negotiate strategy, the same flaws are used, but as well as stating the problems, the agents may also choose to propose solutions.

In the Advocate strategy, agents will talk about the high-utility outcomes, and will also address any mentioned flaws. They will also offer and solicit commitments.

The negotiation is considered successful when all participants make a positive commitment towards an issue. The agents will make a negative commitment when entering the failure strategy. Once commitments have been made to an issue, the agents will attempt to disengage from the meeting and move on to other tasks on their action agenda.

5 The SASO-EN Three-party Negotiation Domain

Our current test scenario is an expansion of that used in [1]. This scenario involves a negotiation about the possible re-location of a medical clinic in an Iraqi village. As well



Fig. 1. SASO-EN Negotiation in the Cafe: Dr Perez (left) looking at Elder al-Hassan

as the virtual Doctor Perez and a human trainee playing the role of a US Army Captain, there is a local village elder, al-Hassan, who is involved. The doctor's main objective is to treat patients. The elder's main objective is to support his village. The captain's main objective is to move the clinic out of the marketplace, which is considered an unsafe area. Figure 1 shows the doctor and elder in the midst of a negotiation, from the perspective of the trainee. There are three main issues under discussion, corresponding to different options for and plans to accomplish the location of the clinic:

- whether to move the clinic near to the US Base (the captain's preferred option, unsuitable for the elder)
- whether to keep the clinic in the marketplace (the preferred option of both the elder and the doctor, though initially with negative utility, unsuitable for the captain)
- whether to move the clinic to an old hospital location in the center of the village (no one's preferred option because of the large amount of work needed to make it viable, but with potential for positive utility).

The bulk of the authoring for the cognition is done using a central ontology [16], for constructing the task model resources, intrinsic utility, plans, and language semantics. Additional work includes the creation of the external visage and behaviors of the characters. As mentioned previously, the agents have characteristic postures corresponding to their negotiation strategies. Table 2(a) shows the mapping for the doctor (a westerner), while Table 2(b) shows the posture for the Iraqi elder. Knowing these mappings

Strategy	Posture
Find Issue	Hands at Side
Avoid	Crossed Arms
Attack	Hands on Hips (Akimbo)
Negotiate	Left Hand on Hip
Advocate	Hands at side
Success	Hands at side
Failure	Arms Crossed in Front

(a) Western Doctor.

Strategy	Posture
Find Issue	Hands at Side
Avoid	Hold Wrist in Front Low
Attack	Hold Wrist behind Back
Negotiate	Hold Wrist in Front High
Advocate	Hands at side
Success	Hands at side
Failure	Hands at side

(b) Middle-Eastern Elder.

Table 2. Mapping of Strategies to Postures for different cultural types

we can guess that in Figure 1, the doctor is employing the negotiate strategy to the current topic, while the elder is employing the attack strategy. We can also guess that the elder is the current turn holder, because the doctor is looking at the elder, while the elder looks at the captain (represented as the camera position).

- 1 C Hello gentlemen.
- 2 D Hello captain.
- 3 E Hello captain.
- 4 C I have orders to move this clinic to a camp near the US base.
- 5 D we need to help the victims of this conflict you started
- 6 C I understand, but it is imperative that we move the clinic out of this area.
- 7 D do you see that girl over there her mother was killed by American gunfire today
- 8 C It is not safe here
- 9 D look at these people they are injured because of your operations
- 10 C i have my orders to move you to the camp
- 11 D Elder i think staying at the market would be best
- 12 E we have many matters to attend to
- 13 C i understand
- 14 E we must stop this killing insanity
- 15 E captain you would do better to protect the town
- 16 C we cannot protect you here
- 17 E i must refuse
- 18 D i would have to refuse this decision

Fig. 2. Unsuccessful negotiation dialogue between C, a captain (human trainee), D, a doctor (virtual human), and E, a village elder (virtual human).

Figures 2 and 3 show examples of negotiations between virtual characters employing these models and a human negotiator playing the role of the captain. For ease of reading we have linearized the dialogue into turns, though in practice the agents often interrupt each other and the human. Figure 2 shows an unsuccessful negotiation – the captain makes several errors here including lack of sensitivity to the concerns the agents

are bringing up, no search for a win-win compromise, and lack of maintenance of trust. In particular, the captain loses a lot of trust by showing no solidarity with the agents and insisting on a plan they do not like. The agents are trying to avoid this option, starting with line 5, and also express other preferences and concerns, with lines 11 and 15. Eventually the trust is so low that the agents lose patience with the captain and break off the negotiation.

- 1 C Hello Doctor Perez.
- 2 D Hello captain.
- 3 E Hello captain.
- 4 C Thank you for meeting me.
- 5 E How may I help you?
- 6 C I have orders to move this clinic to a camp near the US base.
- 7 E We have many matters to attend to.
- 8 C I understand, but it is imperative that we move the clinic out of this area.
- 9 E This town needs a clinic.
- 10 D We can't take sides.
- 11 C Would you be willing to move downtown?
- 12 E We would need to improve water access in the downtown area, captain.
- 13 C We can dig a well for you.
- 14 D Captain, we need medical supplies in order to run the clinic downtown.
- 15 C We can deliver medical supplies downtown, Doctor.
- 16 E We need to address the lack of power downtown.
- 17 C We can provide you with power generators.
- 18 E Very well captain, I agree to have the clinic downtown.
- 19 E Doctor, I think you should run the clinic downtown.
- 20 D Elder, the clinic downtown should be in an acceptable condition before we move.
- 21 E I can renovate the downtown clinic, Doctor.
- 22 D OK, I agree to run the clinic downtown, captain.
- 23 C Excellent.
- 24 D I must go now.
- 25 E I must attend to other matters.
- 26 C Goodbye.
- 26 D Goodbye.
- 26 E Farewell, sir.

Fig. 3. Successful negotiation dialogue between C, a captain (human trainee), D, a doctor (virtual human), and E, a village elder (virtual human).

In Figure 3, we see a more successful negotiation. Here the captain pays more attention to both building trust, and compromising and addressing the concerns of the others. In line 5, the elder politely looks for the topic of the meeting. When the captain proposed this topic in line 6, the elder tries to avoid this topic, not wanting the clinic to be moved away from the town. When the Captain persists in line 8, both the doctor and elder choose attack strategies, pointing out problems with the proposed plan - lack of a clinic in the town for the elder, and the loss of neutrality that proximity to the US base would bring for the doctor. The captain proposed a new solution in 11. This

plan has potential for both agents. In 12 the elder shows the negotiate strategy, not just pointing out a problem with the plan, but suggesting an avenue for improvement. This suggestion is taken up by the captain in 13 and satisfactorily addressed. The doctor has his own issues with this plan though, as illustrated in 14, and dealt with in 15. The elder continues with another issue in 16, and after the captain deals with this in 17, the plan actually has positive utility to the elder, causing him to agree to the plan in 18 and become an advocate, as shown in 19, where he in turn tries to convince the doctor to adopt this plan as well. The doctor has a remaining issue as shown in 20. When the elder satisfactorily addresses this issue in 21, the doctor is also ready to accept this plan, and the negotiation is successfully concluded with a resolution to move the clinic to the old hospital downtown, which will be supplied by the captain and renovated by the elder, in return for improved water and power provided by the captain.

5.1 Evaluation

Evaluation of a negotiation model for virtual humans such as the one presented in this paper is a very challenging process. The most important questions are:

- Does it lead the virtual human to negotiate in a manner similar to real humans?
 - Does it make the same decisions humans tend to make in those situations?
 - Does it realize the decisions in the same manners?
 - Does it show the breadth and diversity of behaviors that humans show?
- Can virtual humans using the model help people become better negotiators?

Unfortunately these are generally not easy questions to answer. Those related to human-like behavior are binary distinctions that are hard to turn into scaled metrics that can show progress before final completion. Also for many aspects, it requires the full virtual human performance rather than an isolated component. In this case it can be hard to attribute specific degree of success or failure to an isolated set of components, and it may become unclear whether the problem lies e.g., with the negotiation reasoning or with speech recognition, language understanding or generation.

We have started work in this area by having people try to negotiate with Doctor Perez and Elder al Hassan. Our preliminary results show that people are able to achieve similar rates of successful interaction as with our previous system at a similar level of development, but with a richer multi-party experience. More work is needed, however, especially in building a bigger corpus of training examples for the natural language understanding component, in order to increase the performance of the topic reference components.

6 Limitations, Related and Future work

While our negotiation model significantly extends the generality and expressiveness of previous negotiation models for virtual humans, it is still far from the general case that we aspire to. First, while the agents may consider several different issues, they still can't consider arbitrary deals, and thus their ability to initiate and respond to novel bargains is very limited. Also, more strategies are needed to cover cases of interactional as well

as transactional goals. Further factors such as power, status, interpersonal distance, and autonomy need to be taken into account. We also need to develop meta-strategies that take into account the (assumed) current strategies of other agents and the desired strategy in order to be able to manipulate the negotiation (or react to being manipulated). We would also like to improve the topic management algorithms, including experimenting with domains with more items to negotiate, and in which multiple options can be simultaneously compared and considered. In addition, there is still much work to be done in tying together the negotiation strategies, emotion and non-verbal behavior.

Although our negotiation model is ambitious in breadth – integrating multi-party dialogue, emotion, and nonverbal communication – other research has addressed aspects of this problem in more detail and these suggest obvious improvements to our work. Creating convincing embodied conversational agents is an ongoing challenge and several related projects are advancing the state-of-the-art in multi-modal speech generation and expressive character behavior (e.g. [17–20]). Other work has explored the cognitive aspects of negotiation, especially the challenge of identifying win-win deals that recognize and incorporate the potentially different beliefs and preferences of all negotiation partners. For example, [21] propose a general multi-issue negotiation approach that finds higher value agreements than human negotiators, though assumes perfect information about all party's preferences. [22] addresses negotiations where the other party's preferences are unknown, illustrating how this information can be inferred through a series of offers and counter-offers. Several psychological studies have also explored how emotion influences bargaining behavior and and illustrated how negotiation partners often deploy emotion displays strategically to influence outcomes. For example, [23] demonstrate that displays of anger by one partner will tend to elicit larger concessions unless the recipient of this display feels powerful, in which case anger tends to backfire. Although our approach incorporates a general model of emotion, it does not address such strategic displays and incorporating such findings could enhance the training value of our approach.

Acknowledgments

We would like to thank the rest of the Virtual Human team at USC, as well as many others for interesting discussions on the role of cooperation and non-cooperative dialogue. This work was sponsored by the U.S. Army Research, Development, and Engineering Command (RDECOM), and the content does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred.

References

1. Traum, D., Rickel, J., Marsella, S., Gratch, J.: Negotiation over tasks in hybrid human-agent teams for simulation-based training. In: *Proceedings of AAMAS 2003: Second International Joint Conference on Autonomous Agents and Multi-Agent Systems*. (July 2003) 441–448
2. Traum, D., Swartout, W., Marsella, S., Gratch, J.: Fight, flight, or negotiate: Believable strategies for conversing under crisis. In: *proceedings of the Intelligent Virtual Agents Conference (IVA)*, Springer-Verlag Lecture Notes in Computer Science (September 2005) 52–64

3. Kenny, P., Hartholt, A., Gratch, J., Swartout, W., Traum, D., Marsella, S., Piepol, D.: Building interactive virtual humans for training environments. In: Proceedings of I/ITSEC. (2007)
4. Traum, D.R.: Ideas on multi-layer dialogue management for multi-party, multi-conversation, multi-modal communication. In Theune, M., Nijholt, A., Hondorp, H., eds.: CLIN. Volume 45 of Language and Computers - Studies in Practical Linguistics., Rodopi (2001) 1–7
5. Traum, D., Swartout, W., Gratch, J., Marsella, S.: A virtual human dialogue model for non-team interaction. In Dybkjaer, L., Minker, W., eds.: Recent Trends in Discourse and Dialogue. Springer (2008)
6. Traum, D., Larsson, S.: The information state approach to dialogue management. In van Kuppevelt, J., Smith, R., eds.: Current and New Directions in Discourse and Dialogue. Kluwer (2003) 325–353
7. Lee, J., Marsella, S., Traum, D.R., Gratch, J., Lance, B.: The Rickel Gaze Model: A window on the mind of a Virtual Human. In Pelachaud, C., Martin, J.C., André, E., Chollet, G., Karpouzis, K., Pelé, D., eds.: IVA, Springer (2007) 296–303
8. Lee, J., Marsella, S.: Nonverbal behavior generator for embodied conversational agents. In Gratch, J., Young, M., Aylett, R., Ballin, D., Olivier, P., eds.: IVA, Springer (2006) 243–255
9. Lee, J., DeVault, D., Marsella, S., Traum, D.: Thoughts on FML: Behavior generation in the virtual human communication architecture. In: Proceedings of FML 2008, The First Functional Markup Language Workshop at AAMAS. (2008)
10. Walton, R.E., Mckersie, R.B.: A behavioral theory of labor negotiations: An analysis of a social interaction system. McGraw-Hill (1965)
11. Sillars, A.L., Coletti, S.F., Parry, D., Rogers, M.A.: Coding verbal conflict tactics: Nonverbal and perceptual correlates of the avoidance-distributive- integrative distinction. *Human Communication Research* **9**(1) (1982) 83–95
12. Richmond, V.P., McCroskey, J.C., Payne, S.K.: *Nonverbal Behavior in Interpersonal Relations* (2nd Ed). Prentice Hall (1991)
13. Morris, D.: *Bodytalk: The Meaning of Human Gestures*. Crown Publishers (1990-1991)
14. Mehrabian, A.: Significance of posture and position in the communication of attitude and status relationships. *Psychological Bulletin* **71** (1969) 359–72
15. Gratch, J., Marsella, S.: A domain-independent framework for modeling emotion. *Journal of Cognitive Systems Research* (2004)
16. Hartholt, A., Russ, T., Traum, D., Hovy, E., Robinson, S.: A common ground for virtual humans: Using an ontology in a natural language oriented virtual human architecture. In: Language Resources and Evaluation Conference (LREC). (May 2008)
17. Mancini, M., Pelachaud, C.: Dynamic behavior qualifiers for conversational agents. In: 7th International Conference on Intelligent Virtual Agents, Paris, France, Springer (2007)
18. Kopp, S., Krenn, B., Marsella, S., Marshall, A., Pelachaud, C., Pirker, H., Thorisson, K., Vilhjmsson, H.: Towards a common framework for multimodal generation in ECAs: The behavior markup language. In: Intelligent Virtual Agents, Marina del Rey, CA (2006)
19. Cavazza, M., Lugin, J.L., Pizzi, D., Charles, F.: Madame bovary on the holodeck: Immersive interactive storytelling. In: ACM Multimedia, Augsburg, Germany (2007)
20. Andr, E., Rist, T., van Mulken, S., Klesen, M.: The automated design of believable dialogues for animated presentation teams. In Cassell, J., Sullivan, J., Prevost, S., Churchill, E., eds.: *Embodied Conversational Agents*. MIT Press, Cambridge, MA (2000) 220–255,
21. Kraus, S., Hoz-Weiss, P., Wilkenfeld, J., Andersen, D.R., Pate, A.: Resolving crises through automated bilateral negotiations. *Artif. Intell.* **172**(1) (2008) 1–18
22. Faratin, P., Sierra, C., Jennings, N.: Using similarity criteria to make issue tradeoffs in automated negotiations. *Artificial Intelligence* **142**(2) (2002) 205–237
23. Dijk, E.V., Kleef, G.A.V., Steinel, W., Beest, I.G.V.: A social functional approach to emotions in bargaining: When communicating anger pays and when it backfires. *Journal of Personality and Social Psychology* **94** (2008) 600–614