Fight, Flight, or Negotiate: Believable Strategies for Conversing under Crisis

David Traum[1], William Swartout[1], Stacy Marsella[2], and Jonathan Gratch[1]

[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. This paper describes a model of conversation strategies implemented in virtual humans designed to help people learn negotiation skills. We motivate and discuss these strategies and their use to allow a virtual human to engage in complex adversarial negotiation with a human trainee. Choice of strategy depends on both the personality of the agent and assessment of the likelihood that the negotiation can be beneficial. Execution of strategies can be performed by choosing specific dialogue behaviors such as whether and how to respond to a proposal. Current assessment of the value of the topic, the utility of the strategy, and affiliation toward the other conversants can be used to dynamically change strategies throughout the course of a conversation. Examples will be given from the SASO-ST project, in which a trainee learns to negotiate by interacting with virtual humans who employ these strategies.

1 Introduction

How can we teach negotiating skills effectively? Effective negotiating skills are critical for many fields, such as commerce, diplomacy and the military. While general principles for effective negotiation can be taught in a classroom setting, becoming an effective negotiator requires practice, usually in a role-playing situation where a teacher or mentor plays the part of one of the opposing party in the negotiation. While this approach can be very effective, it is also expensive in terms of the human resources it requires. In this paper, we describe advances we have made in the technology of virtual humans with the aim of allowing them to act as role-players in a negotiation practice. While a negotiation can be viewed as a rational process of weighing costs and benefits, anyone who has haggled with a salesman over the purchase price of a new car knows that there are significant emotional and non-rational aspects. If virtual humans are to be effective role-players, they must incorporate these aspects as well.

Our work on virtual humans is part of the overall research agenda of creating embodied conversational agents (see collected papers in [1]) that can engage in spoken language interaction with humans, although our emphasis in this paper

on modeling human-like negotiation behavior is unique. This emphasis also sets us apart from the efforts in the multi-agent community on negotiation where the emphasis is in modeling largely agent-agent negotiations as a means to achieve better or more profitable coordination and cooperation (e.g., [2]). The research we describe here extends virtual human models such as those deployed in the MRE project [3, 4] by endowing the virtual humans with strategies for negotiation, endowing them with the ability to model the emotions that arise during a negotiation, and providing facilities for them to communicate verbally and non-verbally during a negotiation dialogue.

The next section describes the initial domain we have chosen to illustrate this research. Section 3 discusses an approach to adversarial communication based on analyses of negotiation in social sciences. Section 4 presents a first synthesis of this work in terms of strategies for virtual humans. Section 5 describes the extensions we have made to the virtual humans from the MRE project to incorporate these strategies and support adversarial negotiation. Section 6 concludes with a discussion of future work.

2 Domain Testbed: Stabilization and Support Operations

Fig. 1. SASO-ST VR clinic and virtual human doctor

Whether it is Kosovo, East Timor, or Iraq, one lesson that has emerged from attempts at "peacemaking" is that negotiation skills are needed across all levels of civilian and government organizations involved. To be successful in these operations, a local military commander must be able to interact with the local populace to find out information, negotiate solutions, and resolve minor problems

before they become major. To have a lasting positive effect, interactions between military and locals must be carried out in a way that generates goodwill and trust. We have selected this general class of operations as a testbed for our work on negotiation.

More specifically, we are developing a training scenario in which a local military commander (who has a rank of captain) must negotiate with a medical relief organization. A virtual human plays the role of a doctor running a clinic. A human trainee plays the role of the captain, and is supposed to negotiate with the doctor to get him to move the clinic, which could be damaged by a planned military operation. Ideally, the captain will convince the doctor without resorting to force or threats and without revealing information about the planned operation. Figure 1 shows the trainee's view of the doctor in his office inside the clinic. The success of the negotiation will depend on the trainee's ability to follow good negotiating techniques, when confronted with different types of behavior from the virtual doctor.

3 Adversarial Negotiation

One of the central ways to characterize negotiation under adversarial conditions is with respect to the tension between competition and cooperation. Negotiators may have different goals, perceive themselves in conflict over those goals but may also perceive the need to cooperate to some degree to achieve their goals. In this view, one can characterize the state of a negotiation process from the perspective of the competitive/cooperative orientation of the parties to the negotiation and the strategies they employ in light of those orientations. Specifically, one oft-made distinction is between integrative and distributive [5] situations. If a negotiation is a win-lose game where there is a fixed value to be distributed, then it is called distributive. There will be a winner and a loser. In contrast, an integrative situation is one where both sides can potentially win, a win-win situation where negotiation could add value and be of benefit to both sides. These basic distinctions presume some commitment to engage in negotiation. However, an individual may simply believe that there is no possible benefit or even need to negotiate. This individual may have an orientation to simply avoid the negotiation or deny the need for it, what is termed avoidance (e.g., [6]). We thus start with three basic orientations toward a negotiation: avoidance, distributive, and integrative. Whenever an agent seriously considers a negotiation situation it will choose one of these three orientations

Negotiators may perceive a situation as one to be avoided, or as a distributive or integrative situation regardless of whether this reflects the true situation. Changing the perceptions of other agents is often one of the main tasks in a successful negotiation. Based on current perceptions, people tend to use a range of dialog tactics consistent with their orientations [7, 6]. Avoidance tactics include shifting the focus of conversation and delays. Distributive tactics can include various defensive moves such as stating prior commitments that bind the negotiator or arguments that support the negotiator's position. Distributive tactics

can also be more offensive, such as threats, criticisms, insults, etc. Integrative tactics are more cooperative with negotiators actually attempting to see issues from the other's perspective. Tactics can be arguments that support the other's position, acceptances of offers, offers of support, etc. Note at a finer grain of analysis, the tactics employed have both instrumental and affective components. For example, distributive tactics, besides trying to gain competitive advantage, tend to be associated with angry or intimidating behavior whereas the integrative tactics try to promote a positive affective climate [7].

Negotiators will often shift orientations during the course of a negotiation. Several factors have been identified as being critical to moving towards an integrative orientation, including acts of reciprocity, establishing trust, reinforcing shared goals, etc. (e.g., [8]).

4 Negotiation Strategies for Virtual Humans

One of our first steps toward implementing a virtual doctor character was to analyze how people act in that role. To this end, we have been conducting a series of role-play sessions, in which one person plays the role of the captain while another plays the role of doctor. Each is given a short set of instructions with different background information, goals, and resources for the negotiation, but given freedom as to how to conduct the negotiation and react to their partner. In these dialogues we can see examples of each of the orientations described in the previous section. For example in (1), the doctor displays an avoidance orientation, and is able to divert the topic of the conversation from the move to the military's role in upcoming operations for over 10 turns (only the first few are shown here). In (2), we see a doctor illustrating the distributive orientation, contesting the basic facts and goals rather than working together on common issues. In (3), we see an example of integrative orientation, the doctor having accepted the danger of the current location and willing to meet the captain's goals if his own are also addressed.

(1) C: it's a temporary move, once the battle is over, you will be moved back.
D: Why don't you cancel your battle? Why don't you not kill these people.
C: We're not the ones deciding the battle.
D: You're the ones here. You're telling me this.

(2) C: We need to move as soon as possible. There are insurgents in the area. This is very unsafe, you're putting yourself and your patients in danger.
D: Why? I don't want to move. I have all these patients here. They won't move, if I move who would who could save them?
C: Sir, Everyone is in danger! If we stay here there's ...
D: I'm not in danger

(3) C: insurgents will not hesitate to harm civilians if that's their path that they need to take. They won't hesitate to harm doctors, a doctor or even injured patients if they feel that's the the means to their end.

D: well
C: this is why you need to come to us.
D: I think we can make a deal. You can give me medical supply, and then we can go with you. I need supplies as soon as possible. As you can see, we are running out of supplies.

We have developed *strategies* for each of these orientations. Our virtual humans can use the strategies to adjust their behavior toward the orientations described above. A strategy consists of several aspects including: **entry conditions**, which indicate when adoption is appropriate; **exit conditions**, which indicate when the strategy should be dropped (often in favor of more appropriate strategies); **associated moves**, which can be performed as tactics to implement the strategy; and **influences** of the strategy on behavior and reasoning. These aspects result from the underlying emotion and dialogue models of the virtual humans.

The EMA (**EM**otion and **A**daptation) model of emotion [9] describes how coping strategies arise as cognitive and physical responses to important events, based on the appraisal [10] of perceptions related to goals and beliefs. Appraisal characterizes events in terms of variables that guide the selection of an appropriate response (e.g., is this desirable? can it be avoided?), but the event need not be physical. Negotiation strategies can thus be seen as types of coping strategies in which the event in question is the negotiation itself, and moves are the types of dialogue actions an agent will perform as part of a negotiation.

The avoidance orientation arises from an appraisal that the negotiation is undesirable but avoidable. The main motivation is to try to escape from the negotiation. When this appraisal is active, the agent chooses an **avoidance** strategy. Exit conditions will be the negation of either of the entry conditions — when the agent believes either that the negotiation has some utility or that it is not avoidable, the agent will abandon the avoidance strategy. The avoidance strategy involves attempts to change the topic of a conversation or get out of it entirely. When applying the avoidance strategy an agent will refrain from commenting on the object of negotiation, even to refute claims.

When in distributive mode, the agent will attempt to "win" rather than "lose" the negotiation. This can be associated with several strategies, depending on the type of decisions to be made and the range of possible alternatives. An *attack* strategy is appropriate when the appraisal is that a negotiation is not avoidable and the proposal is undesirable. Other strategies are also appropriate for a distributive orientation, including defense against a threat rather than attack, or making unreasonable demands in the hope the other party will drop the negotiation. We defer this for future work. One should drop an attack strategy when either the negotiation becomes desirable, or it becomes more profitable to avoid (or defend) than attack. The attack strategy involves pointing out the reasons why a proposal is flawed, or ad hominem attacks on the negotiator.

An integrative orientation leads to attempts to satisfy the goals of each of the participants. The **negotiate** strategy is appropriate when an agent thinks there is a possible value to the negotiation — e.g., there is a higher expected utility

from the expected outcomes than would be the case without the negotiation. This strategy is dropped either when the perceived utility of continuing to negotiate drops below a threshold, or when the negotiation has been completed. Moves in the negotiation strategy involve problem solving and bargaining, much in the manner of the team negotiation in [4].

The success of a negotiation is also mediated by factors that influence the perceived trust between parties, including a belief in shared goals, credibility and interdependence. The doctor is unlikely to be swayed by an offer of aid if he does not believe the captain can and will fulfill his commitments. Trust issues are pervasive throughout the strategies, though building trust will be crucial in allowing the adoption of integrative strategies, since there can be little point in negotiating with someone you expect to lie, be ill-disposed toward you, or not keep their side of a bargain.

Implementing the strategies in a virtual human leads to much more realistic negotiation behavior, allowing our virtual human to engage in many of the types of behavior seen in the role play exercises. For example, the dialogue in Figure 2 shows a sample interaction with our virtual doctor. This is just one of many possible interactions, depending on the choices of the human captain, as well as several aspects (some probabilistic) influencing the choice of moves and strategy transitions of the virtual human doctor. We can see several distinct phases of this dialogue, relating to different negotiation strategies. The initial segment (turns 1-7) includes initial greetings and establishing the topic for the conversation - the captain wants to move the clinic. In turns 8-12, the doctor engages in the avoidance strategy, trying to avoid this topic by bringing up other issues, such as his need for supplies, and the general problems of conflict. In turns 14-20, the doctor has adopted the attack strategy, and points out problems with the proposed move. In turns 22-25, the doctor is in the negotiate strategy, and an actual bargain is struck. Finally, turns 26-30 show a closing phase in which the doctor disengages from the conversation, while the captain tries to establish good relations for future interaction. Application of these strategies influences not just the choice of dialogue move, but the whole body posture of the doctor and use of gestures and expressions as well. In the next section we give more details on how these strategies are implemented in virtual humans.

5 Virtual Human Negotiation Implementation

We take as our starting point the virtual humans implemented as part of the MRE project [3]. These virtual humans are embedded in a dynamic virtual world, in which events can happen, agents can perform actions, and humans and virtual humans can speak to each other and communicate using verbal and non-verbal means. The virtual humans include sophisticated models of emotion reasoning [9], dialogue reasoning [11] and a model of team negotiation [4]. Agents use a rich model of dialogue closely linked with a task model and emotional appraisals and coping strategies for both interpretation of utterances as well as for decisions about when the agent should speak and what to say.

1 C Hello Doctor Perez.
2 D Hello.
3.1 C I am Captain Kirk.
3.2 C Very nice to meet you! [+*FAMILIARITY*]
4 D What do you want?
5 C I have orders to assist you in moving this clinic to a safer location.
6 D You want to move the clinic?
7 C Yes [-*SOLIDARITY*]

Strategy shift: Avoid negotiating about move

8.1 D This conflict is madness, it is killing people!
8.2 D We need proper supplies here!
9 C It is not safe here we can't protect you. [-*CREDIBILITY*]
10.1 D Protect me? Protect me from what?
10.2 D Are you going to attack?
11 C I can not disclose more information but it will be a problem to stay here.
12 D You are the problem your bombs are killing these people!
13 C Sir, we are trying to help these people too. [-*CONTROL* +*SOLIDARITY*]

Strategy shift: Disrupt negotiation about move

14.1 D The patients are here,
14.2 D to move the clinic is not possible we have many patients in critical condition.
15 C We can move you to a safer location.
16 D I can't move all these patients.
17 C We can organize transportation and assist you with moving. [+*SOLIDARITY*]
18 D I can't move all these patients.
19 C We can have locals move you. [+*SOLIDARITY*]
20 D There are no supplies there.
21 C We can provide you with supplies to help your patients. [+*SOLIDARITY*]

Strategy shift: Willingly negotiate about move

22.1 D I think I understand what you say.
22.2 D Well perhaps we can reach an agreement,
22.3 D but before we can think about moving, we need you to bring antibiotics and plasma to stabilize the patients. [+*INTERDEPENDENCE*]
23 C We can do that! [+*SOLIDARITY*]
24.1 D Well,...
24.2 D Very well Captain contact my assistant to make further arrangements.
25 C I'll see to it personally. [+*SOLIDARITY*]
26.1 D I understand your position.
26.2 D My patients need my attention now.
27 C Thank you Doctor!
28.1 D Well,....
28.2 D I must go now
29 C Goodbye.
30 D Good bye.

Fig. 2. Example negotiation dialogue between C, a captain (human trainee) and D, a doctor (virtual Human), showing strategy shifts and positive and negative effects on trust.

To negotiate and collaborate with humans and artificial agents, virtual humans must understand not only the task under discussion but also the underlying motivations, beliefs and even emotions of other agents. The virtual human models build on the causal representations developed for decision-theoretic planning and augment them with methods that explicitly model commitments to beliefs and intentions. Plan representations provide a concise representation of the causal relationship between events and states, key for assessing the relevance of events to an agent's goals and for assessing causal attributions. Plan representations also lie at the heart of many reasoning techniques (e.g., planning, explanation, natural language processing) and facilitate their integration. The decision-theoretic concepts of utility and probability are key for modeling non-determinism and for assessing the value of alternative negotiation choices. Explicit representations of intentions and beliefs are critical for negotiation and for assessing blame when negotiations fail [12].

These virtual humans thus provided a good starting point for implementation of the negotiation strategies described in the previous section. In the rest of this section we describe the enhancements to these virtual humans which were necessary to allow adversarial negotiations such as that shown in Figure 2. First, we talk about aspects of the task and emotion model, including meta-actions for negotiation itself, which allows explicit calculation of the costs and benefits of negotiating, and serves to inform the decisions for entering and exiting strategies. Next, we talk about the trust model, which is both dynamic through the course of a dialogue and influences cognitive and expressive behavior. Then we examine extensions to the dialogue model to use strategies in choice of move and body posture. Finally we briefly describe a tool to look inside the mind of the virtual human and see the effects of specific utterances.

5.1 Appraising the Negotiation

The EMA model of emotion incorporates general procedures that recast the notion of emotional appraisal into an analysis of the causal relationship between actions and goals in an agent's working memory. For example, if an action of the Captain threatens one of the doctor's goals, this is undesirable and deserving of blame, resulting in a response of anger. Depending on if the Doctor can take actions to confront the threat, he may feel in control and engage in problem-focused coping, or resign himself to the threat.

Our view of negotiation orientation as a form of appraisal and coping can be represented within this existing model by simply encoding the negotiation process as just another plan (albeit a meta-plan [13]) within the task representation described above. The potential outcomes of this plan are appraised alongside the rest of the task network by the existing appraisal mechanisms, and coping strategies applied to this task are mapped into different dialogue moves. Thus, the negotiation about moving the clinic is represented as a single "negotiate(move-clinic)" action that is automatically added to the task model in response to the user opening a negotiation. This action has two meta-effects,

"cost" and "benefit" which represent the potential costs and benefits of moving the clinic to another location.

Two extensions are needed to derive the utility of these meta-effects and their likelihood of attainment. One extension to the model is that the utilities of these meta-effects are dynamically computed based on the current task and dialogue state. In particular, the costs and benefits are derived by appraising the individual sub-actions of the "move-clinic" plan. Any desirable effects with high intensity are viewed as benefits and any undesirable effects with high intensity are costs. Currently, these are simply added to compute an overall cost and benefit. The perceived cost and benefit may change through the course of the negotiation. For example, the doctor may believe there are no supplies in the new location (a necessary precondition of the important goal of treating victims), but the trainee may offer to provide supplies, and if believed, this commitment would negate this threat to the move-clinic plan. A second extension is to base the likelihood that the negotiation will succeed on properties of the dialogue state. Currently, we adopt a simple heuristic. If the trainee persists in discussing the negotiation, its likelihood of success increases, though the costs and benefits of that success will depend on what concessions. the trainee has made.

Appraisal and coping operate directly on this meta-action. If the costs exceed the benefits (appraised as undesirable) but the negotiation is unlikely to succeed (leading to an appraisal of high changeability), the doctor will respond with mild fear and copes through avoidance. If the trainee persists in discussing the move (leading to an appraisal of low changeability), without addressing the underlying costs and benefits, the doctor will respond with anger and cope by working against the negotiation (corresponding to the distributive orientation). If the trainee makes concessions that raise the perceived benefits of the move, the doctor will respond with hope and work towards the negotiation (corresponding to the integrative orientation).

5.2 Modeling Trust

According to the dialogue model in [14], the direct effect of an assertion is the introduction of a commitment, whether or not either party believes in the assertion. While this is sufficient for reasoning about the claims and responsibility for information, we need to go further and potentially change beliefs and intentions based on communicated information. Trust is used to decide whether to adopt a new belief based on the commitments of another.

Similar to [15] and [16] , trust is modeled as function of underling variables that are easily derived from our task and dialogue representations. *Solidarity* is a measure of the extent to which parties have shared goals. Solidarity is positively updated when the trainee makes assertions or demands that are congruent with the agent's goals. *Credibility* is a measure of the extent to which a party makes believable claims. It is positively updated when the trainee makes assertions that are consistent with the agent's beliefs. Finally, *familiarity* is a measure of the extent to which a party obeys norms of politeness. Currently, an overall measure of trust is derived as a linear combination of these three factors.

5.3 Acting on Negotiation Strategies

We extended the dialogue model of [3, 4] to take explicit account of strategies and their influence on dialogue behavior. This model already allowed both reactive responses (e.g., to answer a question, to ground an utterance, to respond to a proposal) or speaker initiatives (e.g., to suggest a necessary or desired action, to bring the dialogue back on track, according to an agenda of "to be discussed" items). This model did not address non-team negotiation; the integrative approach was assumed and there was no possibility of avoiding a negotiation or trying for an outcome other than what was good for the whole team. We have extended the model to include explicit strategies, as described above, which govern how agenda items will be discussed. Strategies govern choice of topic and dialogue acts, base body posture, and verbal and non-verbal (e.g. words and gestures) realizations of acts.

The avoidance strategy is implemented by reversing the usual topical coherence guidelines of sticking with one topic until it is resolved before bringing up a new agenda item. When avoiding a topic, rather than direct grounding or negotiation, agenda items which are not central to the topic itself are raised. The doctor's nonverbal behavior also changes, including a posture shift to a crossed arm stance, as shown in Figure 1.

The attack strategy does focus on the topic itself, but only on the reasons why it might be bad. Each of these (potential) reasons, as calculated by the task model, are added to the agenda, prioritized by the importance of the objection. When the speaker no longer thinks they are objections, they will be removed from the agenda. There is also a preference to bring up new objections rather than repeat old ones (subject to the relative importance). If the attack strategy is used when there are no objections in the task model, the speaker will instead question the motivations for the action. When applying the attack strategy, the doctor assumes an aggressive stance, with arms on hips at rest position.

The negotiate strategy follows the model from [4], with the focus of negotiation to make sure that subactions of a plan to achieve a shared goal are committed to by the relevant agents, and maximizing utility for the speaker, perhaps through bargaining. When following the negotiate strategy, the doctor's posture is more open, with arms casually to the side, when at rest.

Some of the same topics may be brought up in both the attack and negotiate strategies, for example, the deficiencies of a plan. Generally there will be a difference in focus, however — in the attack strategy the focus is on why this is a reason not to act, while in the negotiate strategy, the focus is on the concern as a mutual problem to be addressed and solved.

5.4 Explaining Agent Negotiating Behavior

For really learning about negotiation it is very helpful to know not just what the other party did, but why. In real negotiations it is usually not possible to get "inside the head" of the negotiating partner, and even subsequent questions can sometimes damage the nature of the interaction itself. In this respect, virtual

9 C it is not safe here we cant protect you
DECREASES CREDIBILITY: captain asserted unbelieved state 'patients-unsafe-here'
10.1 D protect me protect me from what
'patients-unsafe-here' could be established by captain's act of 'planned-attack'
10.2 D are you going to attack
11 C i can not disclose more information but it will be a problem to stay here
12 D you are the problem your bombs are killing these people
13 C sir we are trying to help these people too
DECREASES CONTROL:captain persists in negotiating 'run-clinic-there'
INCREASES SOLIDARITY: captain committed to achieve desired state 'help-victims'
Strategy shift: Disrupt negotiation about run-clinic-there
(outcome seems negative but negotiation seems unavoidable)

Fig. 3. example trace from AAR tool

humans present a real opportunity to improve on training. We have implemented a trace facility that provides an annotated transcript of the dialogue, showing not just what the virtual human thought was said, but how it influenced his trust, beliefs, and strategy choice. This tool can be used in an "after action review" (AAR) to look in detail at the specific effects the trainee's negotiation tactics had. Figure 2 shows a very abbreviated version of this (for both space and clarity reasons). In Figure 3 we show the full trace for a small section of the dialogue. Here we can see the reason for decreases in credibility and control and increases in solidarity at these points as effects of the commitments the captain makes in relation to desires and beliefs of the doctor. Initially the doctor does not believe the assertion made in 9. However, he realizes that if the captain attacks, that would establish the unsafe condition, leading to the provocative question. Later on, we see that the captain's persistence in talking about moving leads to the abandonment of the avoidance strategy.

6 Current Directions and Future Work

Our current implementation allows a human to interact with the virtual doctor using speech and have many different negotiations of the sort illustrated in Figure 2. The success or failure of the negotiation depends on the use of good negotiating tactics. We are expanding the coverage in several directions to be able to handle fully spontaneous dialogue such as those from which (1),(2), and (3) were taken from. We also plan to evaluate the performance of the doctor virtual agent, in a manner similar to the evaluation done for the MRE system [17].

Negotiation is a complex human interaction. Although we have made significant progress in modeling negotiation, much work remains and there are several directions we plan to take our research next in order to extend our models. The social science literature has identified a wide range of dialog moves/tactics that negotiators use and we are interested in extending our work to incorporate these

moves. We also want to extend the reasoning capabilities to handle other issues in constructing arguments and conflict resolution, e.g. [18]. Another key interest for us is the role that cultural factors play in negotiation, specifically, the role that culture plays in the concerns of the negotiators, their tactics and nonverbal behavior.

Acknowledgments

We would like to thank other members of the SASO-ST project. The project described here has been sponsored by the U.S. Army Research, Development, and Engineering Command (RDECOM). 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. Cassell, J., Sullivan, J., Prevost, S., Churchill, E., eds.: Embodied Conversational Agents. MIT Press, Cambridge, MA (2000)
2. Rosenschein, J.S., Zlotkin, G.: Rules of Encounter: Designing Conventions for Automated Negotiation among Computers. MIT Press (1994)
3. Rickel, J., Marsella, S., Gratch, J., Hill, R., Traum, D., Swartout, W.: Toward a new generation of virtual humans for interactive experiences. IEEE Intelligent Systems **17** (2002)
4. Traum, D., Rickel, J., Marsella, S., Gratch, J.: Negotiation over tasks in hybrid human-agent teams for simulation-based training. In: In proceedings of AAMAS 2003: Second International Joint Conference on Autonomous Agents and Multi-Agent Systems. (2003) 441–448
5. Walton, R.E., Mckersie, R.B.: A behavioral theory of labor negotiations: An analysis of a social interaction system. McGraw-Hill (1965)
6. 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** (1982) 83–95
7. Putnam, L.L., Jones, T.S.: Reciprocity in negotiations: An analysis of bargaining interaction. Communications Monograph (1982)
8. Wilson, S.R., Putnam, L.L.: Interaction goals in negotiation. Communication Yearbook **13** (1990) 374–406
9. Gratch, J., Marsella, S.: A domain-independent framework for modeling emotion. Journal of Cognitive Systems Research (2004)
10. Scherer, K.R., Schorr, A., Jonstone, T., eds.: Appraisal Processes in Emotion. Oxford University Press (2001) Series in Affective Science.
11. Traum, D.R., Rickel, J.: Embodied agents for multi-party dialogue in immersive virtual worlds. In: Proceedings of the first International Joint conference on Autonomous Agents and Multiagent systems. (2002) 766–773
12. Mao, W., Gratch, J.: Social judgment in multiagent interactions. In: In proceedings of AAMAS 2004: Third International Joint Conference on Autonomous Agents and Multi-Agent Systems. (2004) 210–217
13. Litman, D.J., Allen, J.F.: A plan recognition model for subdialogues in conversation. Cognitive Science **11** (1987) 163–200

14. Matheson, C., Poesio, M., Traum, D.: Modelling grounding and discourse obligations using update rules. In: Proceedings of the First Conference of the North American Chapter of the Association for Computational Linguistics. (2000)
15. Marsella, S., Pynadath, D., Read, S.: Psychsim: Agent-based modeling of social interactions and influence. In: In proceedings of International Conference on Cognitive Modeling. (2004) 243–248
16. Cassell, J., Bickmore, T.: A relational agent: A model and implementation of building user trust. In: Proceedings of ACM CHI conference, New York, ACM Press (2001) 396–403
17. Traum, D.R., Robinson, S., Stephan, J.: Evaluation of multi-party virtual reality dialogue interaction. In: Proceedings of Fourth International Conference on Language Resources and Evaluation (LREC 2004). (2004) 1699–1702
18. de Rosis, F., Grasso, F., Castelfranchi, C., Poggi, I.: Modelling conflict resolution dialogues. In Müller, H., Dieng, R., eds.: Computational Conflicts - Conflict Modeling for Distributed Intelligent Systems. Springer (2000)