Virtual Humans for non-team interaction training

David Traum,[*] William Swartout,[*] Jonathan Gratch[*] and Stacy Marsella[†]
[*]USC Institute for Creative Technologies, 13274 Fiji Way, Marina del Rey, CA 90292
[†]USC Information Sciences Institute, 4676 Admiralty Way, Marina del Rey, CA, 90292
traum@ict.usc.edu, swartout@isi.edu, gratch@ict.usc.edu, marsella@isi.edu

ABSTRACT

We describe a model of virtual humans to be used in training for non-team interactions, such as negotiating with people from other organizations. The virtual humans build on existing task, dialogue, and emotion models, with an added model of trust, which are used to understand and produce interactional moves. The model has been implemented within an agent in the SASO-ST system, and some example dialogues are given, illustrating the necessity for building social bonds.

1. INTRODUCTION

Some kind of social and affective relationships between agents are needed for all but the most individualistic kinds of interaction. For teams working together on a shared task, there are quite strong requirements on mutuality [7, 13]. Allwood [2], defined *Ideal Co-operation* between parties as a situation in which the parties

1. take each other into cognitive consideration
2. have a joint purpose
3. take each other into ethical consideration
4. trust each other to act in accordance with 1-3.

[3] discussed how cooperation can also be less than ideal when only some of these factors hold, or they hold only to a lesser extent. Teams do not exist a priori – generally they must be built by individuals who start with more neutral relationships. While some teams may be built artificially when agents engage in activities with specific team roles, or out of local perceived self-interest, in general bonds are needed to hold teams together. To do this, ethical consideration and trust must be built from a starting point in which such trust may not exist. Building such trust is a real issue in team-building, especially when there are conflicting goals or interests.

We claim that virtual humans can play an important role in helping train these skills of establishing bonds and teams. By building virtual humans that are not just humanoid in appearance and external behavior, but which also have internal models (including beliefs, goals, plans, and emotions) and ability to reason over these models and formulate appropriate strategies and behaviors on the basis of the models and perceptual input, virtual humans can behave appropriately for a range of social relationships, e.g., by taking other agents into cognitive and ethical consideration (more specifically, by fulfilling obligations or reasoning about politeness issues) and trusting other agents to do the same.

In previous work [22, 25], we described virtual humans that could engage as teammates and negotiate and carry out team tasks. While this model handled cases where strong social bonds were already assumed (including common end goals, a social institution with roles that the participants played, and strong trust in the teammates' abilities and veracity), it did not address how virtual humans might interact in the case where these bonds were lacking, and how to begin to form them through interaction.

In this paper, we describe the first attempts to extend this model to the more general case, where bonds may need to be developed during the interaction, and in which the virtual human's behavior may be very different depending on the nature and strength of the bonds. In the next section, we describe our initial testbed: a scenario within the SASO-ST project. In Section 3, we describe the virtual human model, including the task model, dialogue model, and emotion model, and how trust of the agent toward another is calculated. In section 4, we show two example interactions with this agent, showing how the dynamic trust model is developed during the interaction and how this can affect the agent's choice of utterance.

2. DOMAIN TESTBED: SUPPORT OPERATIONS

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 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 training interaction with non-teammates.

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

Figure 1: SASO-ST VR clinic and virtual human doctor

will depend on the trainee's ability to follow good negotiating techniques, when confronted with different types of behavior from the virtual doctor.

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 negotiation, since there is usually not much point in negotiating with someone you expect to lie, be ill-disposed toward you, or not keep their side of a bargain.

3. VIRTUAL HUMAN MODEL

We take as our starting point the virtual humans implemented as part of the MRE project [22]. 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 are extensions of the Steve agent [21], and include sophisticated models of emotion reasoning [11, 17, 18], dialogue reasoning [28, 24] and a model of team negotiation [25]. In this section, we briefly describe the building blocks that were used to support the trust-building interactions.

3.1 Task Model

The ability of our agents to negotiate and collaborate with humans and artificial agents on tasks in the virtual world stems from their understanding of those tasks but also the underlying motivations, beliefs and even emotions of other agents. Our models build on the causal representations developed for decision-theoretic planning (e.g., [16]) and augment them with methods that explicitly model commitments to beliefs and intentions [17]. 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 evaluating alternatives for achieving goals. Explicit representations of intentions and beliefs are critical for negotiation and also for assessing blame when negotiations fail [16].

Specifically, each task description is represented as a hierarchical task network. Interdependencies among steps are represented by ordering constraints (i.e., the effect of action A establishes a precondition of action B) and threat relations (i.e., the effect of action C threatens the establishment of a precondition of action D) Preferences over effects are represented by utility values that are either primitive (indicating that the effect has intrinsic worth for some entity) or derived (indicating that the effect has worth as a means towards some end, as in a subgoal). A probability calculus derives the likelihood of actions and effects. State predicates are tagged with a belief, indicating if the some entity has committed to a particular truth value concerning this predicate and a probability calculus represents the entity's measure of belief. Actions and states are also tagged with an intention, indicating an entity's's commitment to bringing about the indicated action or effect. Finally, to support negotiation, the representation encodes multiple, exclusive ways to achieve goals. These alternatives may differ in terms of their effects, likelihood, and utility, entities involved, etc.

This representation serves as a snapshot of an agent's mental state. Agent's continually revise this representation as the result of planning, execution, perception and natural language processing.

3.2 Dialogue Model

Our agents use a rich model of dialogue that is closely linked with the task model both for interpretation of utterances as well as for decisions about when the agent should speak and what to say. We follow the Trindi project approach to dialogue management [14]. The part of the context deemed relevant for dialogue modelling, termed *information state*, is maintained as a snapshot of the dialogue state. This state is then updated by dialogue moves, seen as abstract input and output descriptions for the dialogue modeling component. Moves are calculated at several layers, as described in [28]. the participants [26]. We focus here on the layer of social commitments, as this has the most direct connection to the task model and the emotion model. This includes *social commitments* — both obligations to act or restrictions on action, as well as commitments to factual information [27, 20]. There is also a *negotiation* layer, modeling how agents come to agree on these commitments [4, 23].

3.2.1 Obligations and Social Commitments

Core speech acts have functions related to influencing the topic under discussion and establishing and resolving the commitments and obligations of speakers and other conversational participants towards states and actions. Core speech acts have a content which is either a state, an action description or a question about one of these.

Each of the states and actions in the task model is annotated with semantic information that can be used to describe and recognize description of those states in natural language (and our speech-act based agent communication language). Speech recognition and natural language interpretation produces similar contents from spoken utterances. Dialogue processing then compares the NL representation to the relevant task model representations, and, if a sufficiently close match can be found with a task model state or action, that is seen as the referent.

The core speech acts that are currently modelled include **assert**, **info-request**, **order**, **request** and **suggest**. Unlike many accounts of the effects of these speech acts (e.g. [8, 1, 6, 10]), there are no direct effects on the beliefs, desires or intentions of the conversational participants. This allows for the possibility that participants are insincere in their utterances. Following [27], the direct effects involve social commitments, and one may then infer from these commitments the beliefs or intentions commonly associated with

these utterance types, given additional assumptions.

Assertions will have the effect of establishing a commitment by the speaker that the state holds, or that action happened, is happening, will happen, or should happen, depending on the tense and aspect of the utterance. **Info-requests** have a question as their contents. Questions are (possibly partial) propositions together with a designated *q-slot* indicating the part of the proposition asked about. Info-requests have as their effect an obligation to address the question. **Requests** have an action as content, and the effect is an obligation to address the request, e.g., to consider and give feedback on the request. **Orders**, which can only be performed by a superior to a subordinate in the social structure, have as their effect an obligation to perform the action that is its content. **Suggestions** do not impose obligations, but do focus the topic on the action.

In addition to these *forward-looking* acts [9], there are also backward-looking acts, that point back toward previous dialogue acts or aspects of conversational structure. These will tend to relieve obligations e.g., by performing obliged actions or addressing other utterances.

3.2.2 Dialogue Processing

Language processing occurs in two distinct and interleavable "cycles", one for understanding language and updating the information state, and a second for producing language. This separation of input and output processing cycles allows the agent to have an arbitrary interleaving of contributions by itself and others rather than enforcing a rigid turn-alternation. Each communicative contribution is simultaneously interpreted at each layer, and may correspond to a number of acts at different layers. Generation usually starts from an intention to perform a main act, however any realized utterance will also correspond to a number of acts, some of which (e.g., turn-taking) may be as much a result of the timing of the performance with respect to other events as to the planned behavior.

3.3 Emotion Model

The computational model of emotion in our virtual humans is called EMA (**EM**otion and **A**daptation) [12]. Like many computational models of emotion, EMA is informed by psychological appraisal theory [15], but unlike most computational models, EMA provides a deep process model of the mechanisms underlying emotion, including the cognitive assessments that precede emotion, their consequences on cognition and behavior, and the way these consequences impact subsequent assessments. In terms of the antecedents of emotion, it broadens the goal-based reasoning underlying prior process models to cover other cognitive assessments implicated by appraisal theories. It is also the first computation model of coping, a process associated with emotion in appraisal theory and which is claimed to coordinate an organism's cognitive and physical response to emotional events. By modeling these deep processes explicitly, we are able to facilitate a tight integration of emotion with dialogue processing.

A central tenant in appraisal theory is that appraisal and coping center around a person's interpretation of their relationship with the environment. This interpretation is constructed by cognitive processes, summarized in terms of a series of appraisal variables (e.g., desirability, controllability, and blame) and altered by coping responses (e.g., approach, avoid or deny). To capture this interpretative process in computational terms, we have found it most natural to build on the task representations described above. It encodes the causal relationship between events and states, key for assessing the relevance of events to an agent's goals. The decision-theoretic concepts of utility and probability are key for modeling appraisal variables of desirability and likelihood. Explicit representations of intentions and beliefs are critical for properly reasoning about causal attributions, as these involve reasoning if the causal agent intended or foresaw the consequences of their actions [18]. Commitments to beliefs and intentions also play a role in modeling coping strategies.

We treat appraisal as a set of feature detectors that map features of this representation into appraisal variables. For example, an effect that threatens a desired goal is assessed as a potential undesirable event. Appraisal variables include:

- Perspective: from whose viewpoint is the event judged
- Desirability: what is the utility of the event if it comes to pass, from the perspective taken (e.g., does it causally advance or inhibit a state of some utility)
- Likelihood: how probable is the outcome of the event
- Causal attribution: who deserves credit or blame
- Temporal status: is this past, present, or future
- Controllability: can the outcome be altered by actions under control of the agent whose perspective is taken
- Changeability: can the outcome be altered by some other causal agent

Coping sends control signals to auxiliary reasoning modules (i.e., planning, belief updates, etc.) to overturn or maintain those features that yielded the appraisals. For example, coping may resign the agent to the threat by abandoning the desired goal. Strategies include:

- Action: select an action for execution
- Planning: form an intention to perform some act (the planner uses intentions to drive its plan generation)
- Seek instrumental support: ask someone that is in control of an outcome for help
- Procrastination: wait for an external event to change the current circumstances
- Positive reinterpretation: increase utility of positive side-effect of an act with a negative outcome
- Acceptance: drop a threatened intention
- Denial: lower the probability of a pending undesirable outcome
- Mental disengagement: lower utility of desired state
- Shift blame: shift responsibility for an action toward some other agent
- Seek/suppress information: form a positive or negative intention to monitor some pending or unknown state

Not every strategy applies to a given stressor (e.g., an agent cannot engage in problem directed coping if it is unaware of an action that impacts the situation), however multiple strategies can apply. EMA proposes these in parallel but adopts strategies sequentially. EMA adopts a small set of search control rules to resolve ties. In particular, EMA prefers problem-directed strategies if control is appraised as high (take action, plan, seek information), procrastination if changeability is high, and emotion-focus strategies if control and changeability is low.

Strategies change behavior, but they also change the agent's interpretation of its circumstances, leading to re-appraisal. For example, simply intending to perform an act can improve the agent's appraised sense of control and generate positive emotions. In terms of behavior, coping strategies provide the input to the behavioral, task and language processes that actually execute these directives. For example, plan related coping will generate an intention to perform some action that will make an undesirable situation better which in turn leads to the planning system to generate and execute a valid plan to accomplish this act. Alternatively, coping strategies might abandon the goal, lower the goal's importance, or reassess who is to blame. This close connection between appraisal, coping and cognition provides the processes "hooks" that facilitate the influences between emotion and dialogue.

3.4 Modeling Trust

According to the dialogue model in [20], 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 [19] and [5], trust is modeled as a function of underlying variables that are easily derived from our task and dialogue representations. *Solidarity* is a measure of the extent to which parties have shared goals. It is derived from a running tally of how many times the trainee makes assertions or demands that are congruent with the agent's goals. *Credibility* is a measure of the extent a party makes believable claims. It is derived from a running tally of how many times 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.

1 C Hello Doctor Perez.
2 D Hello.
3 C I have orders to move this clinic to another location.
4 D You want to move the clinic?
5 C Yes
DECREASES SOLIDARITY: captain endorses undesired act 'run-clinic-there'
6.1 D Look at these people!
6.2 D we need to help them.
7 C It is not safe here, we cant protect you.
DECREASES CREDIBILITY: captain asserted unbelieved (but possible) state 'patients-unsafe-here'
'patients-unsafe-here' could be established by captain's act of 'planned-attack'
8.1 D Protect me? Protect me from what?
8.2 D Are you going to attack?
9 C Yes
DECREASES SOLIDARITY: captain committed to perform undesired act 'planned-attack'
10.1 D We need proper supplies here
Unsuccessfully conclude the negotation about run-clinic-there
10.2 D To move the clinic is not possible, we have many patients in critical condition.
11 C It is very important to move now!
12.1 D You are the threat, I need protection from you!
12.2 D I would have to refuse this decision.
12.3 D I must go now.

Figure 2: Unsuccessful negotiation dialogue between C, a captain (human trainee) and D, a doctor (virtual Human) showing positive and negative effects on trust.

4. EXAMPLE INTERACTIONS

Consider the dialogue in Figure 2. 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.

Here the captain acts as he might with a team member. After starting the conversation, the captain launches directly into describing his purpose, and then answers a loaded question straightforwardly. While this would have worked with a subordinate teammember, it has disastrous effects on the neutral doctor, bringing his trust level almost down to zero and failing to accomplish both his objectives.

In this dialogue, nothing was done by the captain to try to establish a better relationship with the doctor, or address the issue of differing objectives and beliefs. The first exchange after the greetings (utterances 3-5) lowers solidarity by showing different objectives, setting up more of an antagonistic than cooperative interaction. The doctor tries to avoid the topic, focusing instead on his patients, rather than the captain's stated goal. The captain tries to argue for his proposed course of action, but only makes things worse with utterance 7. First, he says something the doctor doesn't believe (that the clinic is in danger), lowering his credibility. The doctor is able to reason though that perhaps the captain knows of a reason why it will be unsafe, and challenges by asking if he is going to cause the danger. In 9, the captain answers sincerely, which is a mistake on two fronts. First, he reveals more about his mission than he should to an outsider, possibly endangering its success if word gets out to his enemies. Second, he shows even further divergence from the doctor's goals — attacking rather than helping the patients. After one more brief attempt to change the topic and get help for his own goals, the doctor gives up on the captain in (10.2), and tries to get out of the negotiation. The captain has failed in his objective and prospects are not good for future relations.

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 humans present a real opportunity to improve on training. As shown in Figure 2, 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. Here we can see the reasons for decreases in credibility and solidarity as effects of the commitments the captain makes in relation to desires and beliefs of the doctor.

Figure 3 shows a more successful interaction, where the captain tries to build bonds as well as accomplish his task. While the captain's behavior in this dialogue is not perfect either (the captain might have been better served spending more time up front establishing familiarity and solidarity and perhaps addressing the doctor's concerns first), it is a big improvement over the dialogue in Figure 2. Here the greetings in turn 3 add some familiarity, and the evasion in turn 11 does not do as much damage as the blan-

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 the transportation and assist you with the 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 the supplies you need 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.

Figure 3: 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.

ket statement of acting against the doctor's interest in the previous dialogue. Things are still not going very well, though, until the captain establishes some common goals with turn 13. With slightly higher trust, the doctor does not break off negotiation at this point, but rather raises a series of objections. By addressing each of the doctor's concerns: safety of patients, lack of supplies, lack of transport, and neutrality, the captain is able to bring him around to the point where the move is not an absolute negative, but is worthy of consideration, as part of a team plan. Finally, the two participants reach an agreement including giving needed supplies as part of the conditions of moving the clinic.

In a companion paper, we describe the negotiation strategies that the virtual doctor uses, based on his current feeling about the desirability and avoidability of the object of negotiation, and the degree of closeness with his interlocutor. We can see several distinct phases of the dialogue in Figure 3, 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 an *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 an *attack* strategy, and points out problems with the proposed move. In turns 22-25, the doctor adopts a more open negotiation 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. For example, when the doctor is feeling more distant and less trusting, he adopts the closed posture as shown in Figure 1. When he is more trusting and open to negotiation, the posture becomes more relaxed, as in Figure 4.

Figure 4: More relaxed and open doctor

5. CONCLUSIONS AND CURRENT WORK

In this paper we have described some aspects of our virtual humans with whom humans can create and maintain social bonds. This work extends previous virtual human models to allow fully uncooperative behavior as well as more helpful negotiation and teamwork. Such models allow configurability as to personality as well as initial closeness, and thus a wide range of possible interaction styles. The aim is to train people for interacting with others from a variety of backgrounds on contentious topics.

We have completed implementation of a basic doctor agent, capable of having a range of face to face spoken interactions such as those in Figures 2 and 3. We are currently testing the domain coverage compared to the kinds of things a person says when playing the doctor's role, and ability for a captain to successfully negotiate. While some expansion of task model and vocabulary are needed, Wizard of Oz tests with captains who have been in similar negotiations show a strong potential of this work for use as a training aide.

Acknowledgements

The work described in this paper was supported by the Department of the Army under contract number DAAD 19-99-D-0046. Any opinions, findings and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the Department of the Army.

6. REFERENCES

[1] J. F. Allen. Recognizing intentions from natural language utterances. In M. Brady and R. C. Berwick, editors, *Computational Models of Discourse*. MIT Press, 1983.

[2] J. Allwood. *Linguistic Communication as Action and Cooperation*. PhD thesis, Göteborg University, Department of Linguistics, 1976.

[3] J. Allwood, D. Traum, and K. Jokinen. Cooperation, dialog, and ethics. *In International Journal of Human Computer Studies*, 53(6):871–914, 2000.

[4] M. Baker. A model for negotiation in teaching-learning dialogues. *Journal of Artificial Intelligence in Education*, 5(2):199–254, 1994.

[5] J. Cassell and T. Bickmore. A relational agent: A model and implementation of building user trust. In *Proceedings of ACM CHI conference*, pages 396–403, New York, 2001. ACM Press.

[6] P. R. Cohen and H. J. Levesque. Rational interaction as the basis for communication. In P. R. Cohen, J. Morgan, and M. E. Pollack, editors, *Intentions in Communication*. MIT Press, 1990.

[7] P. R. Cohen and H. J. Levesque. Teamwork. *Nous*, 35, 1991.

[8] P. R. Cohen and C. R. Perrault. Elements of a plan-based theory of speech acts. *Cognitive Science*, 3(3):177–212, 1979.

[9] Discourse Resource Initiative. Standards for dialogue coding in natural language processing. Report no. 167, Dagstuhl-Seminar, 1997.

[10] FIPA. Fipa 97 specification part 2: Agent communication language. working paper available at http://drogo.cselt.stet.it/fipa/spec/fipa97/f8a21.zip, 1997.

[11] J. Gratch and S. Marsella. Tears and fears: Modeling emotions and emotional behaviors in synthetic agents. In *Proceedings of the Fifth International Conference on Autonomous Agents*, 2001.

[12] J. Gratch and S. Marsella. A domain-independent framework for modeling emotion. *Journal of Cognitive Systems Research*, 2004.

[13] B. J. Grosz and C. L. Sidner. Plans for discourse. In P. R. Cohen, J. Morgan, and M. E. Pollack, editors, *Intentions in Communication*. MIT Press, 1990.

[14] S. Larsson and D. Traum. Information state and dialogue management in the TRINDI dialogue move engine toolkit. *Natural Language Engineering*, 6:323–340, September 2000. Special Issue on Spoken Language Dialogue System Engineering.

[15] R. Lazarus. *Emotion and Adaptation*. Oxford Press, 1991.

[16] W. Mao and J. Gratch. Social judgment in multiagent interactions. In *In proceedings of AAMAS 2004: Third International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 210–217, 2004.

[17] S. Marsella and J. Gratch. A step towards irrationality: using emotion to change belief. In *In proceedings of AAMAS 2002: First International Joint Conference on Autonomous Agents and Multi-Agent Systems*, June 2002.

[18] S. Marsella and J. Gratch. Modeling coping behavior in virtual humans: Don't worry, be happy. In *In proceedings of AAMAS 2003: Second International Joint Conference on Autonomous Agents and Multi-Agent Systems*, July 2003.

[19] S. Marsella, D. Pynadath, and S. Read. Psychsim: Agent-based modeling of social interactions and influence. In *In proceedings of International Conference on Cognitive Modeling*, pages 243–248, 2004.

[20] C. Matheson, M. Poesio, and D. Traum. 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.

[21] J. Rickel and W. L. Johnson. Animated agents for procedural training in virtual reality: Perception, cognition, and motor control. *Applied Artificial Intelligence*, 13:343–382, 1999.

[22] J. Rickel, S. Marsella, J. Gratch, R. Hill, D. Traum, and W. Swartout. Toward a new generation of virtual humans for interactive experiences. *IEEE Intelligent Systems*, 17, 2002.

[23] C. L. Sidner. An artificial discourse language for collaborative negotiation. In *Proceedings of the Fourteenth National Conference of the American Association for Artificial Intelligence (AAAI-94)*, pages 814–819, 1994.

[24] D. Traum. Semantics and pragmatics of questions and answers for dialogue agents. In *proceedings of the International Workshop on Computational Semantics*, pages 380–394, January 2003.

[25] D. Traum, J. Rickel, S. Marsella, and J. Gratch. 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*, pages 441–448, July 2003.

[26] D. R. Traum. *A Computational Theory of Grounding in Natural Language Conversation*. PhD thesis, Department of Computer Science, University of Rochester, 1994. Also available as TR 545, Department of Computer Science, University of Rochester.

[27] D. R. Traum and J. F. Allen. Discourse obligations in dialogue processing. In *Proceedings of the 32^{nd} Annual Meeting of the Association for Computational Linguistics*, pages 1–8, 1994.

[28] D. R. Traum and J. Rickel. Embodied agents for multi-party dialogue in immersive virtual worlds. In *Proceedings of the first International Joint conference on Autonomous Agents and Multiagent systems*, pages 766–773, 2002.