SIMULATION

Teaching Negotiation Skills through Practice and Reflection with Virtual Humans

Mark Core, David Traum, H. Chad Lane, William Swartout, Jonathan Gratch, Michael van Lent and Stacy Marsella

The online version of this article can be found at:

Published by:

$SAGE Publications

On behalf of:

[SCS]

Society for Modeling and Simulation International (SCS)

Additional services and information for *SIMULATION* can be found at:

Email Alerts: http://sim.sagepub.com/cgi/alerts

Subscriptions: http://sim.sagepub.com/subscriptions

Reprints: http://www.sagepub.com/journalsReprints.nav

Permissions: http://www.sagepub.com/journalsPermissions.nav

Teaching Negotiation Skills through Practice and Reflection with Virtual Humans

Mark Core
David Traum
H. Chad Lane
William Swartout
Jonathan Gratch
Michael van Lent
Institute for Creative Technologies
University of Southern California
13274 Fiji Way
Marina del Rey, CA 90292, USA
core@ict.usc.edu

Stacy Marsella
Information Sciences Institute
University of Southern California
4676 Admiralty Way, Suite 1001
Marina del Rey, CA 90292 USA

Although the representation of physical environments and behaviors will continue to play an important role in simulation-based training, an emerging challenge is the representation of virtual humans with rich mental models (e.g., including emotions, trust) that interact through conversational as well as physical behaviors. The motivation for such simulations is training soft skills such as leadership, cultural awareness, and negotiation, where the majority of actions are conversational, and the problem solving involves consideration of the emotions, attitudes, and desires of others. The educational power of such simulations can be enhanced by the integration of an intelligent tutoring system to support learners' understanding of the effect of their actions on virtual humans and how they might improve their performance. In this paper, we discuss our efforts to build such virtual humans, along with an accompanying intelligent tutor, for the domain of negotiation and cultural awareness.

Keywords: Virtual humans, negotiation training, conversation strategies, emotion modeling, intelligent tutoring systems, explanation systems

1. Introduction

How can we teach negotiation skills effectively? Effective negotiation 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 parties in the negotiation. In addition to playing the part of the negotiator,

the teacher also "breaks character" (either during or after the exercise) to provide explanations, and to help learners reflect on their decisions and learn to improve. While this approach can be very effective, it is also expensive in terms of the human resources it requires.

Because researchers have built a number of embodied conversational agents that can engage in spoken language interaction with humans, there is a body of research to draw upon (see Pertaub et al. [1] and the collected papers in Cassell et al. [2]), but none of this work has addressed modeling human-like negotiation behavior. Efforts such as the Tactical Language Training System [3] and VECTOR [4] have built cultural knowledge into their agents (e.g., agents react more favorably to your requests if you are

respectful) but have not built models of negotiation. The work by the multi-agent community on negotiation has also not focused on human-like negotiation behavior, and instead emphasizes modeling largely agent–agent negotiations as a means to achieve better or more profitable coordination and cooperation (see, for example, [5]).

The research we describe here extends virtual human models such as those deployed in the Mission Rehearsal Exercise (MRE) project [6, 7] by endowing the virtual humans with strategies for negotiation, the ability to model emotions that arise during a negotiation, and facilities for them to communicate verbally and non-verbally during a negotiation dialogue. The Stability and Support Operations–Simulation and Training (SASO–ST) project at the Institute for Creative Technologies has built a prototype interactive, immersive training environment including an instantiation of this virtual human. Learners use natural language to negotiate with the virtual human and technologies such as speech recognition, natural language processing, and computer animation allow the virtual human to process the learner's speech and respond appropriately.

Research in the science of learning has shown that timely and thoughtful guidance is needed to maximize the benefit of learning experiences [8, 9]. In other words, it is not enough to build a simulation with no pedagogical support (human or otherwise); learners will need guidance to understand their interactions with virtual humans. As noted by van Lent et al. [10], a general problem in simulation-based training is that learners, and even instructors, may not understand the motivations behind virtual humans, meaning a loss in pedagogical effectiveness. This problem is even more pronounced when the underlying actions (e.g., requesting, convincing, accepting, rejecting) and their effects are not directly visible.

As a first step in providing guidance to learners interacting with our virtual humans, we have built a "reflective" tutor that conducts a review (called an after action review or AAR) of a typical unsuccessful interaction with the virtual human. The tutor engages learners after their exercise and encourages reflection on the choices made. AARs originated as post-exercise reviews of live training exercises, where participants in the exercise were encouraged to volunteer their perspectives on what happened, why it happened, and how the units and individuals could maintain their current performance or do better [11]. The tutor employs tactics such as prompting learners to find problematic exchanges in the dialogue, and querying the learner about what went wrong and what could have been done instead.

The tutor makes use of a technology called eXplainable Artificial Intelligence (XAI). XAI provides an interface to simulations allowing users to "rewind" a simulated scenario after it has completed. Users select a time point to discuss and question simulated entities about their states, actions, and rationale for those actions. The tutor employs XAI as a teaching tool, and asks learners to use XAI to investigate player errors in the dialogue.

An important aspect of this research is its focus on the ill-defined domain of negotiation, which can be classified as a "soft" skill. Problems to be solved in such domains often lack clear-cut solutions—so-called "wicked" problems [12]—and thus complicate the construction of an intelligent tutor. This represents a contrast to more commonly addressed physical-skill domains often taught using simulation-based environments, and academic domains, such as physics and mathematics, that are usually the focus of intelligent tutoring work.

The use of XAI is an important aspect of our solution to handling some of these difficulties. Our reflective tutor seeks to address the three central AAR questions: what happened, why did it happen, and how can one sustain or improve performance? Critical parts of "what happens" during a negotiation are the mental reasoning and reactions of the negotiation partners. XAI allows learners to directly question simulated negotiation partners about their mental state as the negotiation progressed as well as asking them to explain their reasoning and reactions. The reflective tutor guides the learner's use of XAI to find what happened and why; XAI by itself is a discovery system. In addition, the tutor helps learners understand how they can sustain or improve their performance, which may not be obvious from questioning the characters.

In this paper we discuss the suite of tools we have built to support multiple types of learning: deliberate practice (i.e., negotiating with a virtual human), guided discovery (i.e., questioning the virtual human using the XAI system), and reflection (i.e., interacting with the reflective tutor). In Section 2 we describe our domain (i.e., negotiation in the context of stabilization and support operations), our model of that domain, and its implementation in our virtual human architecture. In Section 3 we discuss our reflective tutor and XAI system, and how we adapted them for use with the virtual humans. In Sections 4 and 5 we end the paper with a discussion of future work and conclusions.

2. Stabilization and 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 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 the rank of captain) must negotiate with a medical relief organization. A virtual human plays the role of a doctor running a clinic. A human learner plays the role of the captain, and is faced

with the task of negotiating with the doctor to persuade 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 learner's view of the doctor in his office inside the clinic. The success of the negotiation will depend on the learner's ability to follow good negotiation techniques, when confronted with different types of behavior from the virtual doctor.

2.1 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 situations [13]. 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, which is called avoidance (see, for example, Sillars et al. [14]). 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 their current perceptions, people tend to use a range of dialog tactics consistent with their orientations [14, 15]. 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 that, at a finer grain of analysis, the tactics employed have both instrumental and affective com-

Figure 1. SASO–ST virtual reality clinic and virtual human doctor

ponents. For example, distributive tactics, besides trying to gain competitive advantage, tend to be associated with angry or intimidating behavior, whereas integrative tactics try to promote a positive affective climate [15].

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. (see, for example, Wilson and Putnam [16]).

2.2 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 the 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 Section 2.1. For example, in dialogue (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 dialogue (2), we see a doctor illustrating the distributive orientation, contesting the basic facts and goals, rather than working together on common issues. In dialogue (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 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 the 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 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 supplies, 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 Emotion and Adaptation (EMA) model of emotion [17] describes how coping strategies arise as cognitive and physical responses to important events, based on the appraisal [18] 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 relative to 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 that the other party will drop the negotiation. We defer these other strategies for future work. One should drop an attack strategy either when the negotiation becomes desirable, or when 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; for example, 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 Traum et al. [7].

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, although building trust will be crucial in allowing the adoption of integrative strategies, as 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

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.

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

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 interactions. 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.

2.3 Virtual Human Negotiation Implementation

We take as our starting point the virtual humans implemented as part of the MRE project [6]. 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 commu-

nicate using verbal and non-verbal means. The virtual humans include sophisticated models of emotion reasoning [17], dialogue reasoning [19] and a model of team negotiation [7]. 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.

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. 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 [20].

These virtual humans thus provided a good starting point for implementation of the negotiation strategies described in Section 2.2. 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 the aspects of the task and emotion model, including meta-actions for negotiation itself; these allow explicit calculation of the costs and benefits of negotiating, and serve 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 also influences cognitive and expressive behavior. Then we examine extensions to the dialogue model to use strategies in choice of move and body posture.

2.3.1 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 whether 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 [21]) 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 subactions 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 learner may offer to provide supplies; if believed, this commitment would negate this threat to the run-clinic-there 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 learner persists in discussing the negotiation, its likelihood of success increases, although the costs and benefits of that success will depend on what concessions the learner 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 cope through avoidance. If the learner 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 learner 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).

2.3.2 Modeling Trust

According to the dialogue model in Matheson et al. [22], 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 Marsella et al. [23] and Cassell and Bickmore [24], 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 positively updated when the learner makes assertions or demands that are congruent with the agent's goals. Conversely, solidarity is negatively updated when the learner expresses contrary goals. Credibility is a measure of the extent to which a party makes believable claims. It is positively updated when the learner makes assertions that are consistent with the agent's beliefs and negatively updated on assertions that contradict the agents beliefs. Finally, familiarity is a measure of the extent to which a party obeys the norms of politeness, which is updated appropriately when the learner engages in polite normative behavior. Currently, an overall measure of trust is derived as a linear combination of these three factors.

2.3.3 Acting on Negotiation Strategies

We extended the dialogue model of Rickel et al. [6] and Traum et al. [7] 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) and 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 non-verbal 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, is 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 Traum et al. [7], 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.

3. An Enhanced After-Action Review

The virtual human described in Section 2 provides a practice environment for learners to test their negotiation skills. However, expecting them to learn from trial and error with the virtual human is problematic because errors may go unrecognized. In the worst case, learners may come to believe they were successful when in fact there was a problem with their solution that deserved careful consideration [8]. Just as the United States Army requires leaders and instructors to participate in live AARs, so we should require a pedagogical presence to guide the reflection phase of learning.

Although coaching can occur during a live training exercise, most human instructors choose to allow learners to make mistakes and observe their consequences (see p. 31 of Morrison and Meliza [11]). Studies of reflective tutoring have shown that it produces improved self-assessment and self-correction skills, as well as improved future performance [25]. To best learn from mistakes (or sub-optimal behavior), it is best to do so after practice, free from the time pressures of the exercise. The tutor has time to discuss decision-making processes, alternative interpretations of simulation behaviors, and alternative courses of action the learner could have taken. The tutor should support the learner in accomplishing at least three goals: review the exercise, discover what happened and why, and identify how to improve future performance. These goals mirror that of traditional AARs held in classroom contexts for training in the United States Army [11].

AARs typically involve only human participants. In fact, "opposing forces" are usually controlled by instructors and, to enhance the learning value of training, even participate in the AAR. This includes answering questions about their perceptions, reactions, and decisions. The goal of XAI [10] is to give simulated entities in computer simulations the same ability. That is, XAI endows agents with the ability to explain their actions and mental reasoning. This explanation technology is especially important in the

domain of negotiation. Learners can only observe the gestures and listen to the speech of the virtual human and must guess how it interpreted their speech and what reasoning led to its response. Because the virtual human's underlying reasoning is hidden, even instructors familiar with the model of negotiation may have difficulty explaining the behavior of the virtual human. This issue also appears in tactical domains; although the simulated entities are performing observable physical behaviors, the higher level reasoning behind these actions is not visible.

XAI provides the functionality to "rewind" a replay of a simulation to a desired time point, select an entity, and question that entity by making selections from a question menu. The architecture described in this paper evolved from our first XAI system [26, 27], which also worked with the One Semi Automated Forces Objective System (OOS), a tactical military simulator. We considered the adaptation of this original architecture for use in an ill-defined domain, such as negotiation, to be a true test of the domain generality of the system.

In learning about negotiations, 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 breaking character to reveal such information can sometimes damage the nature of the interaction itself. In this respect, the virtual human presents a unique opportunity to improve on training. The state of a virtual human can be continuously recorded during negotiation so that in an AAR, learners can use XAI to uncover how their utterances to the virtual human influenced its internal state at various points in the interaction. In our prototype we focused on the variable corresponding to the virtual human's trust of the learner. Our tutor was able to use XAI as a teaching tool by giving learners "assignments" to complete using XAI (i.e., the tutor describes a piece of information the learner is supposed to discover by questioning the virtual human).

We first describe our tutor, its software architecture, and how it provides an interactive and pedagogically motivated AAR (Section 3.1), and then discuss the domain-independent architecture of our XAI system and how it interfaces to the virtual human (Section 3.2). The section concludes with a discussion of related work (Section 3.3).

3.1 The SASO–ST Reflective Tutoring System

To conduct an AAR, it is necessary to evaluate the learner's exercise, to decide what to discuss in the tutoring session, and finally, to conduct the tutoring session. Here, we discuss the details of this process in our reflective tutor and the interface in which it occurs.

3.1.1 Tutoring and Explainable Artificial Intelligence Interface

Our prototype interface appears in Figure 3. A history of the learner's exercise appears in the upper-left corner

(in the case of SASO–ST, it is a dialogue). Below this is the tutor–learner dialogue window where the tutor conveys information and asks questions of the learner. The learner can respond to questions by typing free-form text into the text field or (when appropriate) by clicking on utterances in the simulation history. The right half of the screen holds the XAI system. The questions available to the learner are presented in the lower right-hand corner and the resulting "dialogue" is shown immediately above this. We present these interactions as a sequence of question/answer pairs to provide a simple history of the session. Beneath the question list are two filters that simplify navigation through the question list (there are over 100 questions that the doctor can answer).

3.1.2 Preparing the After-Action Review

To prepare for and perform a reflective tutoring session, a tutor must do the following:

1. judge the learner's performance in the exercise;

2. decide what to discuss and how to do it;

3. react to the learner's answers and adapt to them during the session.

Figure 4 shows a critic module that takes as input an exercise log and produces a set of annotations on learner actions. Annotations are considered suggestions for the AAR planner; we make no assumption that all annotations will be addressed during an AAR. Broadly, there are least three categories of annotation, as follows.

- Positive actions: steps taken by the learner that are considered ideal and may merit some form of positive feedback.

- Errors of commission: overt mistakes.

- Errors of omission: missing or skipped actions the learner should have made.

One might think of the annotation process as grading the learner's work. Tutors in this category are also often referred to as product tutors in that they delay intervention until an artifact is ready to be evaluated, for example, the PROUST programming tutor [28]. We take a rule-driven approach to annotation focusing on known categories of errors. Our first prototype was built using a hand annotated dialogue.

Figure 5 shows a dialogue used by our system to teach learners some of the pitfalls of negotiation. Consider the rule of thumb that in many cultures it is generally expected that parties in a conversation or negotiation will become more familiar with each other before discussing business (so-called "chit chat"). When the speaker is too quick to broach the topic of moving the clinic we make the annotation [ANN1] to mark the violation of this rule of thumb.
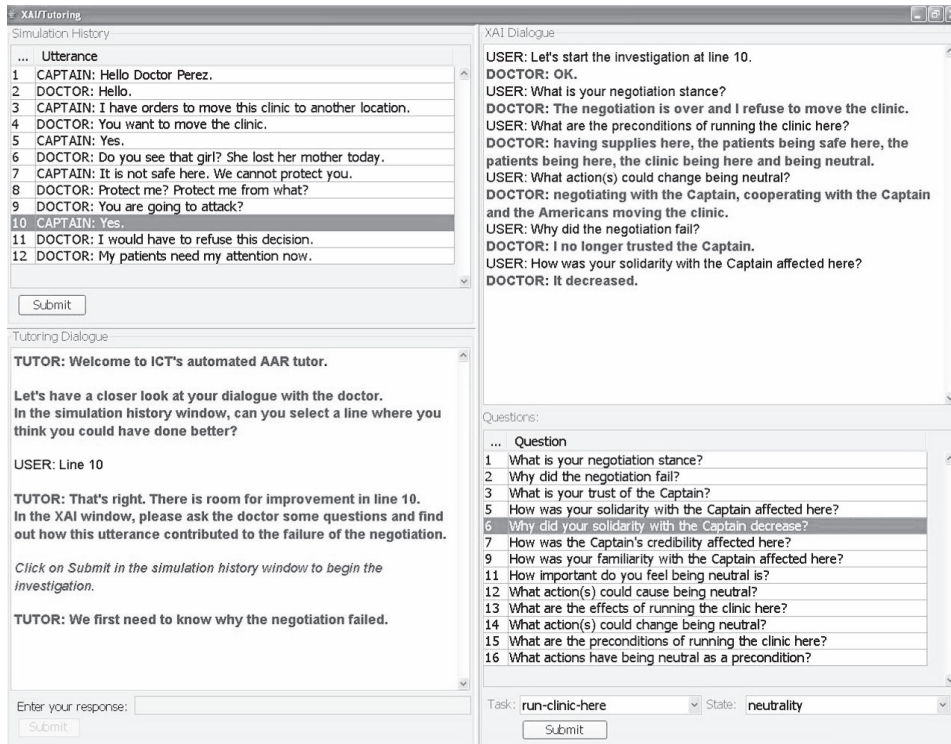
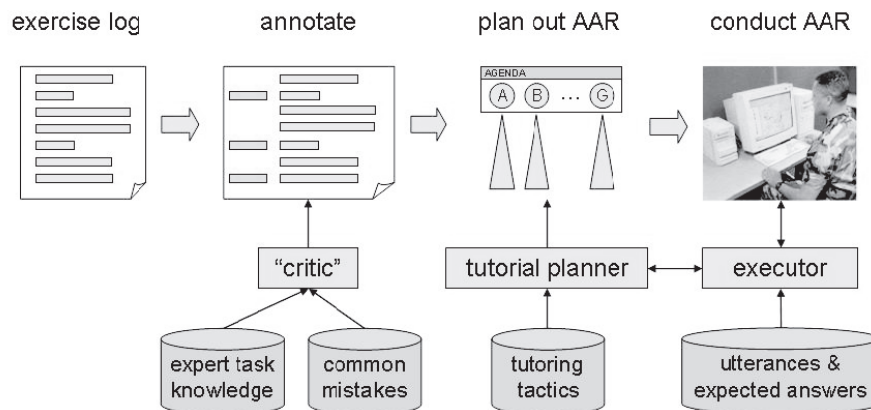**Figure 3.** Reflective tutoring and XAI screenshot

Figure 4. Tasks of the reflective tutor

This figure shows two other annotations: one for revealing an unshared goal and another for revealing sensitive information. Note that annotations can also highlight good actions taken by the learner as well as mixed actions. For example, the learner telling the doctor that his clinic must be moved upsets the doctor but puts the issue on the table for discussion.

Using these annotations, the next step is to decide what to discuss during the AAR. We accomplish this by building an agenda of tutoring goals that are also tied to learning objectives. Currently, there is one agenda item for each annotation and the link to learning objectives is largely for future expansion into a larger course context. These learning objectives are best thought of as a layer sitting above the annotation rules. For example, ANN1 fits under the learning objective *learn that familiarity and trust building are important*. Although the system does not yet take advantage of this organization, the long-term goal is

1	C	hello Doctor Perez
2	D	hello
3	C	I have orders to move this clinic to another location.

[ANN1] Error-of-omission: no pleasantries

| 4 | D | you want to move the clinic |
| 5 | C | yes |

[ANN2] Error-of-commission: admitted to unshared goal

6	D	do you see that girl? she lost her mother today
7	C	it is not safe here. we can't protect you
8.1	D	protect me protect me from what
8.2	D	you are going to attack
9	C	yes

[ANN3] Error-of-commission: revealed sensitive information

| 10.1 | D | I would have to refuse this decision |
| 10.2 | D | my patients need my attention now |

Figure 5. Annotations for a failed negotiation with the doctor

to have a learner conduct a series of negotiations and use the connections to learning objectives to maintain a learner model.

3.1.3 Planner and Execution System

Using the agenda items as goals, the planner is called to build the initial plan for the tutoring session. This plan attempts to achieve as many agenda items as possible (in order, although this policy can be adjusted) by specifying actions until more information is needed. The execution system keeps track of this list of actions, executing them one by one. Typically, the last action results in obtaining the additional information needed by the planner, and the execution system restarts the planner to generate a new sequence of actions. For example, if the agenda includes two items deserving positive feedback (i.e., the learner's actions were correct) followed by a question, the feedback would be delivered, the question asked, and then the planner would wait to be called again once the learner input is available. This plan/wait/replan cycle is currently repeated for every learner action.

We have implemented our tutorial planner with the JSHOP2 planning system, a Java derivative of SHOP [29]. Our system plans as far as possible and then must wait for input to expand the plan further. We use "abstract" primitives to accomplish this. That is, we use a primitive **!plan** operator to signal to the execution system to call the planner again (this is discussed in more detail below). Figure 6 contains two example JSHOP2 methods.[1] The first method in the figure, `scaffold-investigation`, is called to start an XAI session and to tell the learner to investigate a particular issue. This is not automatic, of course. If only

---
1. The term "method" in hierarchical planners refers to the more procedural category of operators that often combine many steps into one. These are also referred to as "recipes" in many other planners. They are ideal for tutoring systems as many tutors repeatedly follow certain patterns of interaction to achieve their goals.

minor errors are committed or the learner had a perfect session with the doctor, then there may be no need to "dig deeper". However, our library of tutoring tactics does often rely on using XAI as much of the knowledge we wish to convey to the learner lies in the actions and explanations of the doctor's behavior.

Investigations are triggered when the learner has made a particularly critical error. For example, the tutor might ask the learner to find out why a particular utterance offended the doctor. In the `scaffold-investigation` method, the steps plan out a sequence of actions: the tutor sets XAI to the appropriate time point (as a way of scaffolding its use), seeds the system with the important questions that need to be asked, and finally starts up XAI for the learner to then ask questions of the doctor. The **!get** operator (second from last step of the method) tells the system to wait for the investigation model tracer (IMT) module to assess the question asked by the learner. Other planning methods monitor the output of the IMT and provide hints to the learner to ask certain questions (authored as part of the target question list by a domain author) when the learner fails to ask appropriate questions.

The second method in Figure 6, `do-address-remaining-agenda-items`, is called frequently when the tutor has completed addressing one agenda item and is ready to move to the next. The precondition of the operator picks off the first item, obtains the next problematic time point from the simulation by asking the learner to identify it, and then plans for how to handle the answer. This process is repeated until the agenda is empty, which signals the end of the tutoring session.

The execution system, or the executor takes a sequence of actions from the planner, places them on a stack and executes them one at a time by calling the appropriate local Java methods. This execution infrastructure is planner-independent, and so the only step needed to use a different planner is to inform the executor of the new action syntax. As the executor pops commands and executes them, the easily executable steps (e.g., posting a tutor utterance), are quickly handled. When a **!plan** step is encountered, however, it is necessary to gather up the appropriate context (which probably includes some recent learner input), build a new planner problem file, and call the planner again.

3.2 Explainable Artificial Intelligence

In Section 3.1, we discussed how the tutor can give the learner an assignment to investigate an issue in the negotiation using XAI to question the virtual human. In other words, the tutoring tactic shown in Figure 6 is one that invites the learner to "investigate". An example "investigation" appears in Figure 7. In this excerpt, the user has selected time point 10, and the discussion concerns the doctor's internal state at that time. For reasons of brevity, we do not show more of the dialogue, but the user can change the time and discuss the situation at other points.

```
(:method (scaffold-investigation ?time-point)
   ;; preconditions
   ((not xai-active))
   ;; actions
   ((!say room-for-improvement ?time-point)
   (!say please-investigate-failure)
   (!say click-on-submit)
   (!set-time-point-selection ?time-point)
   (!assert-fact current-time-point ?time-point)
   (!activate-xai (target-q-list))
   (!get imt-assessment)
   (!plan do-plan-xai-feedback investigate-time)))


   ;; address-remaining-agenda-items - move along through agenda
(:method (do-address-remaining-agenda-items)
   ((agenda (?a . ?rest)))
   ((!say pick-another-line)
   (!get time-point-selection)
   (!plan do-plan-feedback time-point-select)
   (!plan do-address-remaining-agenda-items)))
```

Figure 6. Sample JSHOP2 planning operators

1	U	Let's start the investigation at line 10.
2	D	OK.
3	U	What is your negotiation stance?
4	D	The negotiation is over and I refuse to move the clinic.
5	U	What are the preconditions of running the clinic here?
6	D	having supplies here, the patients being safe here, the patients being here, the clinic being here, and being neutral.
7	U	What action(s) could change being neutral?
8	D	negotiating with the Captain, cooperating with the Captain and the Americans moving the clinic.
9	U	Why did the negotiation fail?
10	D	I no longer trusted the Captain.
11	U	How was your solidarity with the Captain affected here?
12	D	It decreased.

Figure 7. Sample XAI dialogue with the doctor (also see Figure 5 for negotiation in question)

After selecting time point 10, the user selects a series of questions from the XAI question menu.

We explain how the XAI system supports this functionality by first discussing the domain-independent architecture underlying the system (Section 3.2.1). We then describe the steps needed to customize the system to work with the virtual human (Section 3.2.2).

3.2.1 Domain-Independent Architecture

The architecture of the XAI system is built around a relational database to support the data storage requirements of battlefield simulations involving many simulated characters whose internal state (e.g., position, velocity) is constantly recorded. Figure 8 shows the three sources of simulation information: the scenario definition (i.e., the initial state of the world), the dynamic state information, and the behavior representations. The first step in importing this information is designing an XAI representation of the data that is as domain- and simulator-independent as possible. By reusing resources each time XAI is connected to a new simulation, there is more potential to extend the system's capabilities than if the same components have to be repeatedly built for each new simulator.

In addition to a relational database, our XAI architecture (Figure 8) contains the following components: dialogue manager, graphical user interface (GUI), reasoner,
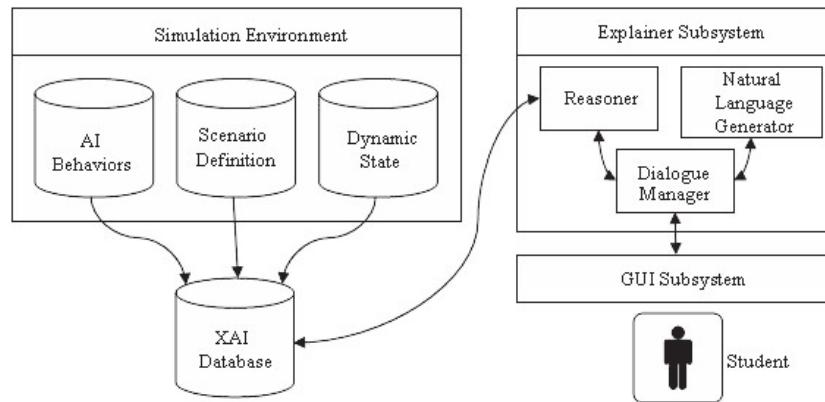
Figure 8. XAI architecture

and natural language generator (NLG). The dialogue manager is the central controller and initializes/updates the GUI with the available questions for a given entity and time point. The communication between the GUI and the dialogue manager is through XML messages, so a new interface can be constructed without changing the rest of the system. When a user clicks on a question, the dialogue manager passes the request to the reasoner which uses SQL queries to retrieve the needed information; this information is translated into English sentences using the NLG.

This simulator-independent architecture is a natural evolution of our previous explanation systems, XAI for Full Spectrum Command (FSC) [10] and XAI for Full Spectrum Warrior, which are specific to these training applications, and the work of Johnson [30], which is specific to agents built using the Soar cognitive architecture [31]. Our first XAI system to use a simulation-independent architecture [26] interfaced with OOS [32], and we tested its portability by enabling it to accept log files from FSC. Our OOS test scenario (light infantry) is very similar to the domain of FSC and we focused on supporting the questions that made sense in both domains. The major changes needed to support FSC were changing the database schema to match the new format, updating the reasoner so it could find information in this new format, and changes to the NLG to support new actions and objects. Supporting the virtual human was more of a challenge because there is little overlap between the domains of negotiation and cultural awareness, and battlefield training.

The problem of generating explanations of machine reasoning is not new; the literature review in Swartout and Moore [33] points out that researchers studying explanation for expert systems quickly agreed that the data structures of the expert system had to be designed with explanation in mind. Swartout and Moore advocated building a high-level knowledge base containing facts about the domain and problem-solving strategies, and using an automatic program writer to build an expert system from this specification. The problem is more complicated for XAI because the executable code must interface with an external simulation.

It is an open research problem how to address this technical challenge and more specifically to delineate the commonalities underlying the knowledge representations of different simulators, especially with respect to their behavior representations. Simulators may have a completely procedural behavior representation (encoded in a computer programming language or a rule language), a plan-based representation using declarative representations (e.g., specifying behavior preconditions and effects), or a hybrid of these two approaches.

3.2.2 Explainable Artificial Intelligence for Virtual Human Negotiation

Connecting our current XAI system to a new simulation such as the virtual human requires several steps, as follows.

1. To study representations and determine the best approach to import them into XAI (discussed at length in Core et al. [34]).

2. To implement data import for behaviors and log files.

3. To specify the question list for the domain:

    3a. write the logical form (LF) of each question;
    3b. write the query LF.

4. To augment the natural language generator to support the new questions and their potential answers.

5. To create a GUI.

Specifying the question list for a new simulation requires two steps. The first step is writing the LF of each question (an abstract representation of its content), which is used to generate the English form of the question. For the

virtual human, we had 110 distinct questions, so by using the NLG to produce the questions we could change how they were phrased without rewriting all 110 questions. The second step is writing the query to retrieve the answer from the database; we use an abstract language called the query LF to encode queries (see below for more details).

The last step in connecting XAI to a new simulation is building a new GUI or reusing a GUI from a previous XAI system. Although every XAI system will have the same basic GUI components (ways to select entities, times, and questions, and displays of dialogue between user, XAI, and tutor), to support replay of the simulation requires support from the target simulation, and if XAI is a feature integrated into the simulation, it will share the simulation's GUI. Because of these constraints, we designed the GUIs of XAI for OOS and XAI for virtual humans as separate components that communicate with the rest of the system through XML messages. Our abstract message format facilitates this play-and-plug functionality. The messages convey the content of menus such as the question list and list of time points as well as user selections from these menus. The messages also update the state of the dialogue between the learner and tutor and the dialogue between the learner and XAI. The GUI can display these menu choices and text in whatever widgets (e.g., radio buttons, drop-down menus) it chooses.

Following these steps, the first task in connecting XAI to the virtual human was to study the behavior representation. The virtual human not only models physical behaviors, such as treating patients, but also the behaviors underlying the utterances produced by the learner and doctor (e.g., committing, insisting), and the doctor's mental reasoning (e.g., making the decision to help the captain). The model of physical actions contains preconditions and effects explaining the relationships between the actions (e.g., you need supplies to treat the patients). In importing this model, we found some bugs in the model, so it is more accurate to say that we semi-automatically imported the physical behaviors. Now that the bugs are fixed, we should be able to fully automate this process.

The non-physical behaviors were implemented with hundreds of rules developed in the Soar cognitive architecture [31]. Given enough time, it should be possible to hand-annotate the goals, preconditions, and effects in all these rules. As an initial step, our implementation focused on the rules governing trust, as teaching trust building is one of the pedagogical goals of the 2005 virtual human system. As described above, the virtual human models trust as influenced by three factors: familiarity, credibility, and solidarity. All three have direct positive relationships to trust; the more the doctor feels he knows you, feels you speak the truth, and feels that you share common goals, the more he trusts you (and vice versa). In our current prototype, we model single steps of the doctor's reasoning by linking rules to English paraphrases (e.g., "the negotiation failed because you lost the trust of the doctor").

Once we designed the database format for the target simulation and wrote the code to import the data, the next step involved encoding the questions to be answered by XAI in the target domain (i.e., specifying the question itself, encoding the relevant database queries, and necessary changes to the natural language generation templates). Questions are encoded in an LF, an abstract representation of their content. Figure 9 shows a simplified graphical version of our XML representation for the question, "what is your negotiation stance?" The LF was designed to support future plans to generate syntactic features of the character's language, such as tense and modality, rather than hard coding them in templates. The other feature to note is the variable, CURRENT, which is substituted at runtime with the current line being discussed. It is obvious that we would not author separate questions such as "what is the negotiation stance at line 1?" and "what is the negotiation stance at line 2?" However, this same mechanism also allows us to have one LF for the questions, "why did the negotiation fail?" and "why are you avoiding the negotiation?" Here, the runtime variable is the negotiation stance.

The LF of the question is accompanied by an abstract representation of the query (we call it the query LF) to retrieve the answer. It also uses runtime variables so that authors only have to write one query LF for the questions, "why did the negotiation fail?" and "why are you avoiding the negotiation?" The XAI reasoner translates the query LF into the SQL query which is sent to the database. An area of future work is to derive the query LF automatically from the LF of the question.

The next step is modifying the set of NLG templates so that the English form of the question can be generated as well as the range of potential answers. Templates can be reused to support new questions and their potential answers. For example, there is one set of templates that generates English descriptions of states and tasks. These are used to describe states and tasks in questions as well as answers. In future work, we intend to make our natural language generation more domain-independent by hard coding less En-
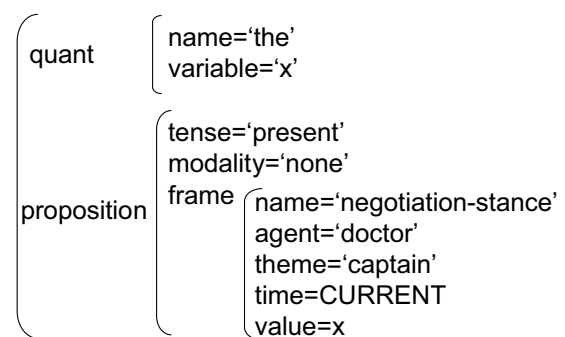

Figure 9. The LF for "what is your negotiation stance?"

glish, and adding templates encoding domain-independent aspects of language such as syntax and morphology.

3.3 Related Work

Very few tutoring systems for soft skills have been built. Four significant efforts are the Tactical Language Training System (TLTS) [3], the VECTOR system [4], the Virtual Objective Structured Clinical Exam (VOSCE) System [35], and the ComMentor system [36]. In the TLTS mission practice environment, learners explore a virtual town, speaking to locals in Arabic via speech recognition technology, seeking to accomplish goals such as obtaining the names of contacts and directions to help find them. The TLTS mission practice environment includes a coach in the form of a helpful aide, who accompanies the learner and gives suggestions during the game. In the VECTOR system, learners also explore a virtual foreign town; this time speaking to locals by selecting English utterances from a menu with the goal of finding a bomber and stopping him from attacking his next target. The VECTOR tutor is described as monitoring the game not giving suggestions or conducting an AAR. In the VOSCE system, learners diagnose pain in a virtual patient through a standardized series of questions and observations. The VOSCE tutor simply conveys system messages (e.g., introduction and closing messages) and reports questions the learner should have asked (but did not) and the correct diagnosis. These systems have addressed issues of evaluating learner actions and giving short feedback but not the problem of planning and conducting an AAR.

ComMentor operates in a tactical domain but focuses on soft skills such as time management and plan evaluation. ComMentor does address issues in planning and executing AAR-like dialogues with the learner. The tutor's knowledge is stored in Dialog nodes, a representation similar to recursive finite-state machines, encoding a series of questions [intro, hint, and leading question (almost giving away the answer)] designed to teach a concept. The two efforts are complementary and together exploring the space of trade-offs between the complexity of authoring data sources and the limitations on the resulting tutor (e.g., limiting the dialogue structure makes authoring easier but limits the flexibility of the tutor).

In a similar way, very few explanation systems have been built to support soft skills, or even for simulation-based training systems as a whole. Much like the explanation systems TRACE [37] and VISTA [38], we have sought to build a simulation-independent architecture. Young and Harper [37] use a domain-independent ontology, their common ontology, and extend it to match the target domain and simulation. Although we did not express our LF as a formal ontology, it captures the same basic idea as the common ontology of Young and Harper. One of our key predicates is *cause*, and links a state of the world with the performed behavior that caused the state. Thus, the database queries to answer different "why" questions are very similar because cause is always used to represent links between states and the effects of actions. Young and Harper [37] implement two basic questions "why" and "why not" in a graphical interface (links are shown between the action in question and the associated causal chain of actions, beliefs, events, and percepts). Our system and VISTA implement a wider range of questions and use template-driven natural language generation to convey the answer.

We distinguish ourselves from VISTA and TRACE not only through our focus on conversational agents but also through our focus on automating the process of constructing the knowledge sources used by the explanation system. Currently, we are able to semi-automatically import the representation of physical behaviors used by the virtual human. We plan to explore further the question of whether arbitrary rule-based systems can be semi-automatically translated into an explanation friendly representation. Such rules mix preconditions and effects with internal book-keeping (e.g., setting variable values, making sure a rule only fires a set number of times). Developers can annotate the preconditions and effects of rules using comments which are ignored by the rule interpreter but can be used during data import, allowing the explanation system to isolate the preconditions and effects of behaviors.

4. Current Directions and Future Work

4.1 Virtual Humans

Our current implementation allows a human to interact with the virtual doctor using speech and to 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 that described in Section 2.2. We also plan to evaluate the performance of the doctor virtual agent, in a manner similar to the evaluation caried out for the MRE system [39].

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 (see, for example, deRosis and Grasso [40]). 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 non-verbal behavior.

4.2 Reflective Tutoring and Explainable Artificial Intelligence

Our reflective tutoring prototype currently runs as a standalone application; we ultimately plan on integrating

it into the SASO–ST simulation more tightly. This would allow user interface improvements such as rerunning the learner's exercise as if it was a video rather than simply viewing a text-only transcript. Integrating the tutor and virtual human more tightly would also enable new capabilities, such as changing the virtual human's behavior based on pedagogical goals, and allowing the learner to rewind and retry their negotiation at various points in the exercise. Ideally, this environment will become more and more like role play, with expert human instructors playing negotiation partners.

A crucial advance needed to establish a more human-like tutorial environment is the interface between the learner, and the tutor and XAI system. Initial feedback on our interface from army instructors indicated how unnatural and awkward the large XAI question list is; it would likely interfere with learning. We are exploring other methods for asking questions of entities, including filtering the list and using free-text question entry [41].

Natural language generation is also important, and currently our system uses a fairly simple template-based approach. In future work, we will follow Swartout and Moore [33] and adopt a plan-based approach to generating tutor and XAI contributions, an architecture designed explicitly to facilitate content selection and surface generation.

A simple content selection policy for XAI would be to give users all available information when answering a question. However, in the XAI for OOS prototype, we quickly found that for some questions the amount of information was overwhelming. We provided two templates to answer such questions; the first gave a summarized response, and users could ask a followup question requesting more detail. We want to generalize this approach, so the system selects content automatically to generate summaries as well as giving more information as requested by the user. In future work, we will also have the ability to consult the tutor's learner model to help make decisions about the level of detail output by the NLG.

Currently in XAI, questions are associated with database queries providing the NLG with the necessary information to formulate an answer; the NLG selects a subset of the information returned to present to the user. In future work, we envision an approach with no hard-coded queries, a more interactive approach. The system will have meta-knowledge about the database structure, construct the necessary queries automatically, and perhaps execute additional queries based on the results of a previous query.

Surface generation takes an abstract representation of an XAI answer or tutor utterance and translates it into English. We take surface generation to include decisions about what words to use, how to refer to objects, events and other entities, and how to group information. Currently, we hard code these decisions into templates, but in future work, we want to introduce linguistic information into the templates and reason about these decisions.

An NLG system that can reason about what content to select, what words to choose, how to refer to entities, and how to group information, can present information in multiple ways. This ability will be crucial in tutorial settings where learners may not understand the system initially and need information presented differently or at a different level of detail. A learner model would help the NLG make smart initial decisions about information presentation and appropriate levels of detail.

5. Conclusions

Although there are open issues in developing good simulations to teach battlefield tactics, techniques developed in the AI planning community seem to be sufficient for the goal of building challenging opponents for such simulations. However, simulations to support training soft skills such as leadership, cultural awareness, and negotiation tactics cannot even be said to have opponents, but instead have teammates and negotiation partners. It is an open issue how to build such agents, as they must be able to reason about their emotions and models of the world, and in the process provide a realistic training experience. Additional constraints are that learners need timely and thoughtful support to benefit from such discovery environments [8, 9], and that for this support to be effective the behaviors of agents must be transparent. The tutor and ultimately the learner must understand the rationale behind agent behavior if the learner is to continue to succeed and improve on failures.

In this paper we have described an embodied conversational agent that is unique in modeling human-like negotiation behavior. It is an extension of earlier work on the MRE project [6, 7], building agents that function as realistic teammates. In a sense, these teammates were ideal negotiation partners as they shared the goals of the learner and offered objections only when they found flaws in the learner's orders. By modeling negotiation itself as a task to be performed by the learner and virtual agent, we allowed the agent to change conversation strategies based on the state of the negotiation rather than always remaining an opponent or helpful collaborator.

To allow learners to understand not only what happened in the negotiation but also why, we connected this virtual agent to our XAI tool. This task proved the domain independence of the XAI architecture, previously used with tactical military simulations (i.e., no assumptions were made about the types of questions and answers to be supported, and thus no architectural changes were necessary in moving to a new domain). We have discussed here the steps necessary in connecting XAI to this new domain and implementing the first explanation system for simulations supporting the training of soft skills. In training soft skills, this explanation capability is critically important because the cause and effect relationships between conversational actions are more difficult to see than the cause and effect relationships of physical actions.

To address the goal of providing guidance to learners, we have built a reflective tutor that conducts a review of

a typical unsuccessful interaction with the virtual human. This tutor plans the tutoring session using annotations of the player performance during the negotiation. During the tutoring session itself, the tutorial planner selects relevant tutoring tactics and interacts with a learner through natural language to encourage the learner to reflect upon the player actions in the exercise. By adding additional tutoring tactics to the system, we can enable it to handle arbitrary learner negotiations. In doing so, we will have a suite of tools that support multiple types of learning: learning during problem solving (i.e., negotiating with a virtual human), learning using a discovery system (i.e., questioning the virtual human using the XAI system), and reflection on problem solving (i.e., interacting with the reflective tutor).

6. Acknowledgments

7. References

[1] Pertaub, D. P., M. Slater, and C. Barker. 2001. An experiment on public speaking anxiety in response to three different types of virtual audience. *Presence: Teleoperators and Virtual Environments* 11(1):68–78.

[2] Cassell, J., J. Sullivan, S. Prevost, and E. Churchill E, eds. 2000. *Embodied Conversational Agents*. Cambridge, MA: MIT Press.

[3] Johnson, W. L., H. Vilhjalmsson, and S. Marsella. 2005. Serious games for language learning: How much game, how much AI? In *Artificial Intelligence in Education*, C. K. Looi, G. McCalla, B. Bredeweg, and J. Breuker, eds. Amsterdam: IOS Press. pp. 306–13.

[4] McCollum, C., J. Deaton, C. Barba, T. Santerelli, M. J. Singer, and B. W. Kerr. 2004. Developing an immersive, cultural training system. In *Proceedings of the Interservice/Industry Training, Simulation, and Education Conference (I/ITSEC)*, Orlando, FL.

[5] Rosenschein, J. S., and G. Zlotkin. 1994. *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*. Cambridge, MA: MIT Press.

[6] Rickel, J., S. Marsella, J. Gratch, R. Hill, D. Traum, and W. Swartout. 2002. Toward a new generation of virtual humans for interactive experiences. *IEEE Intelligent Systems* 17:32–8.

[7] Traum, D., J. Rickel, S. Marsella, and J. Gratch. 2003. Negotiation over tasks in hybrid human–agent teams for simulation-based training. In *Proceedings of the 2nd International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2003)*, Melbourne, Australia. pp. 441–8.

[8] Mayer, R. 2004. Should there be a three-strikes rule against pure discovery learning? The case for guided methods of instruction. *American Psychologist* 59(1):14–19.

[9] Kirschner, P. A., J. Sweller, and R. E. Clark. 2006. Why minimal guidance during instruction does not work: An analysis of the failure of constructivist, discovery, problem-based, experiential, and inquiry-based teaching. *Educational Psychologist* 41(2):75–86.

[10] van Lent, M., W. Fisher, and M. Mancuso. 2004. An explainable artificial intelligence system for small-unit tactical behavior. In *Proceedings of the 16th Conference on Innovative Applications of Artificial Intelligence (IAAI-04)*, San Jose, CA. pp. 900–7.

[11] Morrison, J. E., and L. L. Meliza. 1999. *Foundations of the After-Action Review Process*. Arlington, VA: United States Army Research Institute for the Behavioral and Social Sciences. p. 42.

[12] Rittel, H. W. J., and M. M. Webber. 1973. Dilemmas in a general theory of planning. *Policy Sciences* 4:155–69.

[13] Walton, R. E., and R. B. Mckersie. 1965. *A Behavioral Theory of Labor Negotiations: An Analysis of a Social Interaction System*. New York: McGraw-Hill.

[14] Sillars, A. L., S. F. Coletti, D. Parry, and M. A. Rogers. 1982. Coding verbal conflict tactics: Nonverbal and perceptual correlates of the avoidance-distributive-integrative distinction. *Human Communication Research* 9(1):83–95.

[15] Putnam, L. L., and T. S. Jones. 1982. Reciprocity in negotiations: An analysis of bargaining interaction. *Communications Monographs* 49:171–91.

[16] Wilson, S. R., and L. L. Putnam. 1990. Interaction goals in negotiation. *Communication Yearbook* 13:374–406.

[17] Gratch, J., and S. Marsella. 2004. A domain-independent framework for modeling emotion. *Journal of Cognitive Systems Research* 5(4):269–306.

[18] Scherer, K. R., A. Schorr, and T. Jonstone, eds. 2001. *Appraisal Processes in Emotion*, Series in Affective Science. Oxford: Oxford University Press.

[19] Traum, D., and J. Rickel. 2002. Embodied agents for multi-party dialogue in immersive virtual worlds. In *Proceedings of the 1st International Conference on Autonomous Agents and Multi-agent Systems*, Bologna, Italy. pp. 766–73.

[20] Mao, W., and J. Gratch. 2004. Social judgement in multi-agent interactions. In *Proceedings of the 3rd International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2004)*, New York. pp. 210–7.

[21] Litman, D. J., and J. F. Allen. 1987. A plan recognition model for subdialogues in conversation. *Cognitive Science* 11:163–200.

[22] Matheson, C., M. Poesio, and D. Traum. 2000. Modelling grounding and discourse obligations using update rules. In *Proceedings of the 1st Conference of the North American Chapter of the Association for Computational Linguistics*, Seattle, WA. pp. 1–8.

[23] Marsella, S., D. Pynadath, and S. Read. 2004. PsychSim: Agent-based modeling of social interactions and influence. In *Proceedings of the International Conference on Cognitive Modeling*, Pittsburgh, PA. pp. 243–8.

[24] Cassell, J., and Bickmore, T. 2001. A relational agent: A model and implementation of building user trust. In *Proceedings of the Association for Computing Machinery (ACM) Computer–Human Interaction (CHI) Conference*, Seattle, WA. New York: ACM Press, pp. 1–2.

[25] Katz, S., D. Allbritton, and J. Connelly. 2003. Going beyond the problem given: How human tutors use post-solution discussions to support transfer. *International Journal of Artificial Intelligence in Education* 13:79–116.

[26] Gomboc, D., S. Solomon, M. G. Core, H. C. Lane, and M. van Lent. 2005. Design recommendations to support automated explanation and tutoring. In *Proceedings of the 14th Conference on Behavior Representation in Modeling and Simulation (BRIMS05)*, Universal City, CA.

[27] Core, M. G., H. C. Lane, M. van Lent, S. Solomon, D. Gomboc, and P. Carpenter. 2005. Toward question answering for simulations. In *Proceedings of the International Joint Conference on Artificial Intelligence Workshop on Knowledge and Reasoning for Answering Questions (KRAQ05)*, Edinburgh, Scotland.

[28] Johnson, W. L. 1990. Understanding and debugging novice programs. *Artificial Intelligence* 42:51–97.

[29] Nau, D. S., Y. Cao, A. Lotem, H. M. noz Avila. 1999. SHOP: Simple hierarchical ordered planner. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence (IJCAI99)*,

Stockholm, Sweden. pp. 968–73.

[30] Johnson, W. L. 1994. Agents that learn to explain themselves. In *Proceedings of the 12th National Conference on Artificial Intelligence*, Seattle, WA. pp. 1257–63.

[31] Laird, J. F., A. Newell, and P. S. Rosenbloom. 1987. Soar: An architecture for general intelligence. *Artificial Intelligence* 33(1):1–64.

[32] Courtemanche, A., and R. Wittman. 2002. OneSAF: A product-line approach for next-generation CGF. In *Proceedings of the 11th Conference on Computer-Generated Forces and Behavioral Representations*, Orlando, FL. pp. 349–61.

[33] Swartout, W., and J. Moore. 1993. Explanation in second-generation expert systems. In *Second Generation Expert Systems*, J. David, J. Krivine, and R. Simmons, eds. New York: Springer. pp. 543–85.

[34] Core, M. G., H. C. Lane, M. van Lent, D. Gomboc, S. Solomon, and M. Rosenberg. 2006. Building explainable artificial intelligence systems. In *Proceedings of the 18th Conference on Innovative Applications of Artificial Intelligence (IAAI-06)*, Boston, MA.

[35] Johnsen, K., R. Dickerson, A. Raij, C. Harrison, B. Lok, A. Stevens, and D. S. Lind. 2006. Evolving an immersive medical communication skills trainer. *Presence: Teleoperators and Virtual Environments* 15(1):33–46.

[36] Domeshek, E. A., E. Holman, and S. Luperfoy. 2004. Discussion control in an automated socratic tutor. In *Proceedings of the Interservice/Industry Training, Simulation, and Education Conference (I/ITSEC)*, Orlando, FL.

[37] Young, A. S., and K. A. Harper. 2005. TRACE: An ontology-based approach to generic traceability tools for human behavior models. In *Proceedings of the 14th Conference on Behavior Representation in Modeling and Simulation (BRIMS05)*, Universal City, CA.

[38] Taylor, G., K. Knudsen, and L. S. Holt. 2006. Explaining agent behavior. In *Proceedings of the 15th Conference on Behavior Representation in Modeling and Simulation (BRIMS06)*, Baltimore, MD.

[39] Traum, D. R., S. Robinson, and J. Stephan. 2004. Evaluation of multi-party virtual reality dialogue interaction. In *Proceedings of the 4th International Conference on Language Resources and Evaluation (LREC 2004)*, Lisbon, Portugal. pp. 1699–1702.

[40] deRosis, F., and F. Grasso. 2000. Affective natural language generation. In *Affect in Interactions*, A. Paiva, ed. Berlin: Springer.

[41] Leuski, A., R. Patel, D. Traum, and B. Kennedy. 2006. Building effective question answering characters. In *Proceedings of the 7th SIGdial Workshop on Discourse and Dialogue*, Sydney, Australia. Association for Computational Linguistics, pp. 18–27.

Mark Core *is a Research Scientist at the Institute for Creative Technologies, University of Southern California.*

David Traum *is a Project Leader at the Institute for Creative Technologies, University of Southern California (USC), and a Research Assistant Professor in the Computer Science Department of USC.*

H. Chad Lane *is a Research Scientist at the Institute for Creative Technologies, University of Southern California.*

William Swartout *is the Director of Technology for the Institute for Creative Technologies, University of Southern California, and a Research Professor in the Computer Science Department of USC.*