Negotiation over Tasks in Hybrid Human-Agent Teams for Simulation-Based Training

David Traum,* Jeff Rickel,† 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, rickel@isi.edu, gratch@ict.usc.edu, marsella@isi.edu

ABSTRACT

The effectiveness of simulation-based training for individual tasks – such as piloting skills – is well established, but its use for team training raises challenging technical issues. Ideally, human users could gain valuable leadership experience by interacting with synthetic teammates in realistic and potentially stressful scenarios. However, creating human-like teammates that can support flexible, natural interactions with humans and other synthetic agents requires integrating a wide variety of capabilities, including models of teamwork, models of human negotiation, and the ability to participate in face-to-face spoken conversations in virtual worlds. We have developed such virtual humans by integrating and extending prior work in these areas, and we have applied our virtual humans to an example peacekeeping training scenario to guide and evaluate our research. Our models allow agents to reason about authority and responsibility for individual actions in a team task and, as appropriate, to carry out actions, give and accept orders, monitor task execution, and negotiate options. Negotiation is guided by the agents' dynamic assessment of alternative actions given the current scenario conditions, with the aim of guiding the human user towards an ability to make similar assessments.

Categories and Subject Descriptors

I.2 [**Computing Methodologies**]: Artificial Intelligence

General Terms

Design

Keywords

negotiation, conversational agents, animated agents

1. INTRODUCTION

While simulation-based training has become a common, effective method for teaching individual skills, its application to team training is a new and exciting development. From their historical roots in pilot training, simulation-based learning environments have been developed for a broad range of training scenarios and skills, sometimes including intelligent agents that can demonstrate tasks or give timely feedback to students [26, 27, 30, 38]. For team training, such learning environments could additionally include intelligent agents serving as the trainee's teammates, providing valuable practice in coordinating team actions, leadership, and, in suitably constructed learning environments, decision-making and team coordination under stress.

However, while techniques for building intelligent agents that can guide trainees on individual tasks are well understood, the capabilities needed to support synthetic teammates are far more challenging. The training scenario can set the overall goals and provide strong constraints on the trainee's options, but the trainee must have the freedom to make decisions, choose among reasonable alternatives, and interact naturally with teammates, all with a sufficient illusion of freedom to avoid destroying the realism of the simulation. The agents must deal with a range of teamwork issues, including authority, responsibility, coordinated actions, hierarchical organizational relationships, and group decision making; the agents must be proactive and responsive partners for the trainee, not simply puppets under his control. Finally, the agents must interact with the trainee through the normal modes of team coordination and communication, typically including face-to-face spoken dialogue and the ability to track each other's actions in a shared environment.

Prior work has laid the foundation for many of these capabilities in isolation, but no previous intelligent agents have integrated them to support such a team training environment. Models of teamwork [15, 16, 19, 32] address many of the issues in representing and reasoning about team tasks, but have not addressed the complex human interface issues that arise in hybrid human-agent teams. Conversely, work in computational linguistics and embodied conversational agents [4, 17, 36] has addressed many of these human interface issues, but has not previously addressed the required range of teamwork issues. Prior systems for team training that provide synthetic teammates have typically simplified both the complex teamwork issues and the difficult human interface issues, resulting in systems in which the trainee has a severely limited ability to interact with his teammates [3, 9, 12, 28].

In this paper, we describe an implemented intelligent agent architecture that provides a realistic model of teamwork while also supporting face-to-face spoken dialogue for realistic and flexible interactions among human and agent teammates. Our architecture integrates and extends a variety of prior work, giving our agents a wide range of capabilities, including the ability to reason about au-

AAMAS'03, July 14–18, 2003, Melbourne, Australia.

thority and responsibility for individual actions in a team task and, as appropriate, to carry out actions, give and accept orders, monitor task execution, and negotiate options. Negotiation is motivated by the agents' dynamic assessment of alternative actions given the current scenario conditions, with the aim of guiding the human user towards an ability to make similar assessments.

To stimulate and evaluate our research, we have adapted our agents to negotiate within an implemented peacekeeping training scenario, in which a human user (Army lieutenant) cohabits a 3D graphical virtual environment with animated virtual humans (a platoon of soldiers, and some civilians) and interacts with them through face-to-face spoken dialogue to deal with an unanticipated dilemma (Figure 1) involving a traffic accident causing potentially serious injuries, and a weapons inspection where another unit may require urgent assistance. The graphics are displayed on an 8-foot-tall screen that wraps around the user in a 150-degree arc with a 12-foot radius, and immersive audio software uses 10 audio channels and two subwoofer channels to envelop the user in spatialized sounds that include general ambience (such as crowd noise) and triggered effects (such as explosions or helicopter flyovers) [31]. The sergeant, medic and mother are intelligent agents with a fully integrated set of capabilities including speech recognition, natural language understanding, dialogue management, natural language generation, speech synthesis, human-like perceptual limitations, planning, emotions, and the ability to dynamically control their animated bodies, including synchronization of speech and gestures [29]; the sergeant and medic currently employ a fully implemented version of the models described in this paper to collaborate with the user (lieutenant), each other, and other characters. The user wears a microphone and communicates via unrestricted spoken dialogue. The remaining virtual humans (soldiers and civilians) react with pre-programmed responses to specific inputs; the soldiers can carry out orders from the sergeant or communicate in simple ways, and both soldiers and civilians react to events and direct communication addressed at them. The user's decisions influence the way the situation unfolds, culminating in a glowing news clip praising his actions or a scathing clip exposing decision flaws and describing their sad consequences.

Other papers have described the overall project [31], the aspects of virtual humans [29], prior versions of the team task modelling [27, 28], the dialogue model [36], and the emotional appraisal [14, 20]. In this paper, we show how multiple elements of the virtual human model combine to allow robust and flexible team-oriented behavior, including:

- engaging in dialogue to coordinate team agreement, including giving, receiving, and negotiating orders and requests
- using dialogue state, team roles, and dynamic plans to understand and assess different options and to decide when to give orders and execute actions

In Section 2, we describe the team task model, including the representation of actions (which can be tasks in a plan), states (which can be goals and preconditions of tasks), and team roles, including responsible and authorizing agents. The task model also includes an ability to represent under-determined multiple courses of action (COAs), and assessment among them, given current conditions. In Section 3, we summarize the dialogue model, with a focus on the aspects most relevant for engaging in team tasks: the core speech acts (which establish obligations and commitments toward task elements), and the grounding level, which establishes mutual belief about communicated elements. In Section 4, we describe the representation of task negotiation, and describe both how moves are recognized and how agents decide to make negotiation moves based

Figure 1: An interactive peacekeeping scenario featuring (left to right in foreground) a sergeant, a mother, and a medic.

on factors including the state of the dialogue, the state of the plan, and the social relationships among the teammates. In Section 5, we show how this reasoning is applied in a couple of dialogue examples from our domain. Finally, conclusions are presented in Section 6.

2. TASK MODELS FOR TEAM ACTION

2.1 Representing Team Tasks

The ability of our agents to collaborate with human and agent teammates on tasks in the virtual world stems from their understanding of those tasks. Agents must understand task goals and how to assess whether they are satisfied, the actions that can achieve those goals, how the team must coordinate the selection and execution of those actions, and how to adapt execution to unexpected events. To provide this understanding, our agents use domain-independent reasoning algorithms operating over a general, declarative representation of team tasks, and this representation is used to encode their domain-specific task knowledge for a given training scenario (or class of scenarios). Our representation and reasoning algorithms are based on our earlier work on virtual humans for team training [28] with two key extensions: a representation for authority relations and an ability to handle limited perception (i.e., goals whose status is unknown). Here we briefly review the prior work and these new extensions.

Our representation for team tasks uses a relatively standard plan representation. First, each task description includes a set of steps, each of which is either a primitive action (e.g., a physical or sensing action in the virtual world) or an abstract action (i.e., a task that must be further decomposed). Abstract actions give tasks a hierarchical structure. Second, there may be ordering constraints among the steps, which define a partial order. Third, the interdependencies among steps are represented as a set of causal links and threat relations [22]. Each causal link specifies that an effect of a step in the task achieves a particular goal that is a precondition for another step in the task (or for termination of the task). For example, in our military domain there is an action of marking a landing zone with smoke, which achieves the goal of allowing a helicopter pilot to visually identify the landing zone, which in turn is a precondition for landing it. Threat relations specify that an effect of a step threatens a causal link by undoing the goal before it is needed. For example, extinguishing the smoke before the helicopter arrives threatens the helicopter's ability to land at the desired location.

In addition to understanding the structure of tasks, agents must understand the roles of each team member. Each task step is associated with the team member that is responsible for performing it [28]. We have also extended our representation to include an optional association of each task step with the teammate who has authority over its execution; that is, the teammate responsible for a task step should not perform it until authorization is given by the specified teammate with authority. This extension to the representation was required to model the hierarchical organizational structure of some teams, such as in the military.

Given a top-level abstract task for the team to accomplish, each agent independently uses its task knowledge to construct a complete task model. Starting with the task description for the top-level task, the agent recursively expands any abstract step with its task description, until the agent has a fully decomposed, hierarchical task model. Agents may or may not be given identical task knowledge, and so may or may not construct identical task models; this can be used to model teammates with partial or erroneous knowledge.

An agent's task model represents its understanding of the task in general, independent of the current scenario conditions. To guide execution of the task and robustly handle unexpected events that require adaptive execution or replanning, agents use a partial-order planning algorithm over the task model; the algorithm is described in detail in [27], and its application to reasoning about team tasks is detailed in [28]. The task model specifies all the steps that might be required to complete the task; it can be viewed as a worst-case plan. Agents continually monitor the state of the virtual world via messages from the simulator [27] that are filtered to reflect perceptual limitations [29]. These perceptions allow the agents to update their representations of the status of goals in the task model as being satisfied, unsatisfied, or unknown if they cannot currently perceive the state of the goal. (The possibility of "unknown" extends our prior work.) The planning algorithm works backwards through the causal links in the task model to identify goals that are currently desired and task steps that are currently intended to establish those desired goals. Just as the status of a goal can be satisfied, unsatisfied, or unknown, the planning algorithm marks the "desired" property of goals and the "intended" property of steps as true, false, or unknown. The result of this planning algorithm specifies how the agent privately believes that the team can collectively complete the task, with some causal links specifying the interdependencies among team members' actions. Agents revise this private plan as needed, given current conditions as the scenario unfolds.

2.2 Alternative Courses of Action

A key aspect of collaborative planning is negotiating about alternative ways to achieve team goals. To support such negotiation, we have extended our earlier representation so that task models support reasoning about alternative, mutually exclusive courses of action (recipes) for achieving tasks, and we have added mechanisms for evaluating the relative strengths and weaknesses of different alternatives. These courses of action are self-contained hierarchical tasks in the sense defined above, and subject to the same dynamic task reasoning. For example, one might evacuate someone to a hospital by using either a medevac helicopter or an ambulance. Depending on the circumstances, only one option might be possible (e.g., the medevac may be unavailable or the injuries may be too severe for an ambulance), but if both are valid options, they must be ranked through some reasoned analysis of their relative costs and benefits.

Tasks associated with alternative courses of action (COAs) are treated differently from standard tasks in several ways. Standard tasks are marked as intended if they establish a desired goal. In contrast, if a task associated with a COA establishes a desired goal, it is only marked as relevant, not intended. If a COA is adopted, all relevant tasks associated with the COA are simultaneously marked as intended, and any tasks associated with an alternative COA are marked as not intended. Tasks associated with a COA must also be treated differently with regard to threat detection. As alternative COAs are mutually exclusive, tasks in one alternative cannot threaten or be threatened by tasks in an alternative COA. Finally, to evaluate alternative COAs, they are ranked on the basis of their expected utility, using the decision-theoretic computations specified in [14]. This includes the likelihood that they will achieve the intended abstract task, but also the likelihood of any desirable or undesirable side effects of executing the course of action. We use fuzzy boundaries to sort alternatives so that one is better than another only if it differs substantially in expected utility. In addition to computing an overall expected utility, we identify salient positive and negative aspects of alternative COAs using a psychologically inspired theory of how people assess the significance of events [14, 20]. These aspects include individual consequences (side effects) that are significant either because they have intrinsic worth or make significant progress towards some intrinsically desirable state, as well as threats to desired states. For example, a helicopter evacuation requires certain personnel to setup and secure a landing zone. If these resources are needed for some other task, this resource conflict would be appraised as a negative aspect of the COA.

3. A DIALOGUE MODEL FOR MULTIPLE PARTICIPANT INTERACTION

Much of the overt behavior of negotiation involves communication as part of a *dialogue*. 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. Our dialogue model supports multiple simultaneous conversations among potentially overlapping groups of interlocutors in a shared virtual world [36].

We follow the Trindi project approach to dialogue management [18]. 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. A complex environment such as the MRE situation obviously requires a fairly elaborate information state to achieve fairly general performance within such a domain. We try to manage this complexity by partitioning the information state and dialogue moves into a set of *layers*each dealing with a coherent aspect of dialogue that is somewhat distinct from other aspects.

Each layer is defined by information state components, a set of relevant dialogue acts, and then several classes of rules relating the two and enabling dialogue performance. Several layers are used in the current system. The *contact* layer [2, 5, 10] concerns whether and how other individuals can be accessible for communication. Modalities include visual, voice (shout, normal, whisper), and radio. The *attention* layer concerns the object or process that agents attend to [25]. Contact is a prerequisite for attention. The *Conversation* layer models the separate dialogue episodes that go on during an interaction. Each conversation consists of a number of sub-layers, each of which may have a different information content for different conversations happening at the same time. The *participants* may be active speakers, addressees, or overhearers [5]. The *turn* indicates the (active) participant with the right to communicate (using the primary channel) [25, 35]. The *initiative* indicates the participant who is controlling the direction of the conversation [37]. The *grounding* component of a conversation tracks how information is added to the common ground of the participants [33]. The conversation structure also includes a *topic* that governs relevance, and *rhetorical* connections between individual content units. Once material is grounded, even as it still relates to the topic and rhetorical structure of an ongoing conversation, it is also added to the social fabric linking agents, which is not part of any individual conversation. This includes *social commitments* — both obligations to act or restrictions on action, as well as commitments to factual information [34, 21]. The negotiation layer will be described in the next section. More details on these layers, with a focus on how the acts can be realized using verbal and non-verbal means, can be found in [36]. We focus here on the aspects most central to negotiation: social commitments and grounding.

3.1 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 descriptions of those states using natural language (and our speech-act based agent communication language). For example, the action of the sergeant securing the assembly area (which can be accomplished by having the squad leaders each secure a quadrant) is represented as shown in (1). The resulting state of the assembly area being secure is represented as shown in (2).

(1) **agent** sgt
event secure
patient assembly-area
type act

(2) **object-id** assembly-area
attribute safety
value secure
polarity positive
type state

Speech recognition and natural language interpretation produce semantic representations in the same format. Dialogue processing then tries to match the input semantic representation to the 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. An item is seen as a potential match if, for every role (such as agent, patient, attribute, etc) that has a value in the input, if the task model representation has the same role, then the task model has the same value for that role (partial matches in which a role is missing from one side or the other, or cases in which input or task model has more than one value are also allowed).

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, 7, 13]), 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 [34], 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 content state holds, or that the content 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. For example, (3) shows an info-request by the LT to the Sgt with the content being a question about whether the assembly area is secure. 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.

(3) **action** info-req
actor lt
addressee sgt
type core-speech-act
content **q-slot** polarity
  type question
  prop **object-id** assembly-area
    attribute safety
    value secure
    time present
    type state

In addition to these *forward-looking* acts, there are also backward-looking acts, that point back toward previous dialogue acts or aspects of conversational structure [11]. Backward-looking acts tend to relieve obligations e.g., by performing obliged actions or addressing other utterances. These include acceptances of requests (which will create an obligation to perform the requested act) as well as rejections and other moves that won't include such obligations. We will return to these acts in the next section.

3.2 Grounding

Following [6, 33, 24], we treat *grounding* as occurring in discrete bundles of dialogue-introduced information that are added to the common ground together. Common Ground Units (CGUs) are modeled as information stores with state, which can be updated by the performance of the *grounding acts* from [35, 33]: **initiate**, **continue**, **repair**, **request-repair**, **display**, **acknowledge**, **request-acknowledge**, and **cancel**. Core speech acts are not seen as having their full effects on the social state until they are grounded. Thus, even an attempted order, if not understood as such, will not impose any obligation on the addressee (other than to perform a grounding act, if the utterance is perceived). Grounding acts will often be parts of utterances that include core speech acts, e.g., an answer or acceptance will ground the info-request or request that it relates to. If the agent does not understand an utterance or is unable to decide on a reference for an underspecified act, it may request-repair. Obligations and commitments that have not yet been grounded are still accessible to the agent as *potential obligations*, which can be used in deciding how to react.

3.3 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 performance of 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.

4. NEGOTIATING TEAM TASKS

Negotiation is a higher-level discourse function involving multiple distinct acts to reach a group consensus. There are two views toward negotiation: focus on a particular task, or focus on a decision – which may involve consideration of multiple alternative tasks. For both views, we model a negotiation as a sequence of basic building blocks called *stances*. Each stance represents a public representation of an agent toward a task. In this section we describe the basic stances themselves, as well as a specification of which set of stances are sufficient for team action. This is followed by a description of how stances are created as effects of dialogue acts, and finally how an agent decides which negotiation moves to perform.

4.1 Modelling Negotiation States

The current state of team negotiation on a task step is represented by a sequence of negotiation stances. Each stance is a tuple consisting of the information shown in (4).

(4)
- the agent who holds the stance
- the action that this stance is about
- the stance the agent holds toward the action
- the audience (a set of agents) that the agent has made the stance in front of
- the reason for holding the stance
- the time at which the stance was made

Stances are one of the following set, graded from most positive to most negative: {committed, endorsed, mentioned, not mentioned, disparaged, rejected}. The minimal stances needed to be confident that an action will be performed is at least endorsement from the authorizing agent, and commitment by the responsible agent. An endorsed or committed stance by the authority with the responsible agent in the audience is sufficient for the action to be seen as *authorized*. If the responsible agent has a committed stance to an authorized action, he will be expected to either perform the action, enlist and supervise the performance of others, or retract the stance and explain, if action becomes infeasible. The sequence of negotiation stances indicates the progression of the negotiation, who first proposed, and who finally accepted. Stances by the same agent generally move toward the extremes, with the most recent stance being the predominant one.

4.2 Recognizing Negotiation Stances

Negotiation stances on actions arise from core speech acts or negotiation acts referring to those actions. A suggestion or offer will lead to a mention stance. A request, order or promise leads to a committed stance. An assertion will lead to some stance depending on the modality of the content, e.g., as something that should be done (endorsed) or must be done (committed) or could be done (mentioned). As well as the core speech acts, there are some acts specifically aimed at negotiation. **Accept** acts produce a committed stance. **Reject** acts produce a rejected stance. **Counterproposals** produce two stances: a disparaged stance toward the original proposal, and an endorsed stance for the new proposal. **Explanations** produce either a disparaged or endorsed stance, depending on the relation of the explanation to the action.

4.3 Engaging in Negotiation

There are several factors involved in whether and how to proceed in a negotiation. First, there is the issue of initiating or responding. We use *initiative* in a conversation as a factor in deciding to start a negotiation. Depending on the level of initiative, the agent can decide to refrain from mentioning an act, or mention, endorse, or commit.

If someone else (e.g., another agent or a human trainee) starts a negotiation with an order or request, the agent has an obligation (either to perform the action or at least to address the request). In this case, the agent will respond in order to deal with this obligation. The style and timing of response depends on several aspects of the mental and interactional state of the agent, shown in (5).

(5)
- relevant party
- dialogue state
- plan state

We will take each of these in turn, and then describe how they motivate specific dialogue action. **Relevant party** is a relation between the authorizing agent, the responsible agent, and the agent considering the act. If the agent is the authorizing agent, then the agent is the relevant party. If not, then if the action is not authorized, the authorizing agent is the relevant party. If the action is authorized then the responsible agent is the relevant party. In case the authorizing agent or responsible party is not known (depending on whether authorized or not), the relevant party may be unknown.

Dialogue state is one of three values: *discussed, needs-discussion, or unmentioned* depending on who has already produced stances on this action. If no one has produced a stance, the value is *unmentioned*. If someone has made a request or order regarding the action to the agent, but the agent has not produced a stance, the state is *needs-discussion*. If both agents already have a stance, then the state is *discussed*. Minimal negotiation should reach the discussed stage, however this is not necessarily the end of a negotiation. In general a negotiation between teammates should proceed until both agree, either by both accepting or rejecting, or by one dropping a contrary stance. It is also possible, however, to "agree to disagree," ending the negotiation without having come to an agreement (assuming that it is possible to proceed with other actions).

Plan state is the most complex factor, relying on the relation of the act to the overall plan and execution environment. The values for the plan-state of an action are shown in (6). Not all of the defining conditions are mutually exclusive. E.g., the plan state of an action might be *bad* (because of lack of relevance) and also *goals-satisfied*, because its goals already hold. Likewise, the plan state of an action might be *good*, because it is intended, but *evaluate*, because the plan is in flux. Generally we prefer the most specific conditions (e.g., *goals-satisfied* over *bad*) and *evaluate* over other possibilities. These preferences are easy to modify, however, and we intend to experiment with different settings for different agents, according to emotion and personality type. E.g., a cautious agent may not want to commit to an action until all planning is completed, while a more reckless agent might always answer based on current intentions. A more clever agent may be able to narrow down which changes in the world impact which parts of the plan, and thus have more specific conditions for preferring *evaluate* vs. (considered) *good* or *bad*. In any case, the plan-state *conflict* exists when there are no explicit rules for deciding between applicable plan-state values.

(6) **evaluate:** the world has changed in important ways since last re-planning, and the agent is unsure of current relevance and applicability of the action

good: the action is both intended and is a next-step (based on ordering constraints in the task model)

considered good: the action is not (yet) intended, but is relevant and also part of a best course of action

considered bad: the action is not intended, but is relevant and part of a course of action that is not the best choice

not-in-coa: intention is unknown, but the action is not part of a course of action

premature: the action is intended but it is not a next-step

goals satisfied: the action is not a next-step, and the goals that the action would achieve have already been satisfied

bad: the action is not intended or considered relevant

unknown: there is no step in the task model corresponding to this action

conflict: there are irreconcilable preferences for identifying the task

4.3.1 Planning negotiation responses

The combination of the three features *relevant-party, dialogue-state, and plan-state* guide the decision of what negotiation acts to perform in response to an initial action proposal (e.g., an order, request, or suggestion to do the act), according to the motivation sets shown in (7).

(7) **Accept:** relevant-party=me, plan-state $\in$ {good, considered-good, not-in-coa }, dialogue-state=needs-discussion

Reject: plan-state $\in$ { bad, considered-bad, unknown, conflict, goals-satisfied } , dialogue-state=needs-discussion

Counterpropose: plan-state $\in$ {considered-bad, premature }, dialogue-state=needs-discussion

Delay: Plan-state=evaluate

Redirect Relevant-party $\neq$ me

Accept (reluctantly) relevant-party=me, plan-state=considered-bad, dialogue-state=discussed

Express discussed dialogue-state=discussed

Express role-unknown Relevant-party = *unknown*

In some cases more than one of these sets of motivation conditions will apply, and the agent will need to decide which to do, and, if more than one, in which order. There are a number of guidelines that govern the final decision. Some are based on generic preference rules, such as, *when counterproposing because the action is premature, and several actions should be done first, choose the highest-level immediately performable action*. Likewise, when counterproposing because another action is better, propose the best action (not just one that is better than the alternative originally proposed). Other choices may depend on the social relationships. E.g., a repeated order by a superior will be reluctantly accepted (once an agent has tried to propose a better alternative). On the other hand, a repeated request from a subordinate may be rather motivation to express that the matter has already been discussed, rather than giving in when not really convinced, or repeating the rejection and explanation.

4.3.2 Argumentation sequences

Sometimes an agent may want to express a more complex rhetorical argument rather than just produce a simple response. A good example is when a trainee makes an order or request that is considered bad. One option might be to simply reject. Another option would be to make a counterproposal. Neither of these is completely satisfying, however. What is better is to give a justification of why the action is bad, which can serve multiple purposes, including helping to convince the trainee, as well as teaching the factors to look at in making decisions.

We are just beginning to look at the general issue of explanation and argumentation in dialogue, which must be sensitive not just to the range of physical, social and intentional factors at the time of planning the explanation, but also how the world and interactional state changes as the explanation is being performed. For example, if an interlocutor accepts a counterproposal right off it may not be as necessary to explain the rationale. Likewise, if the interlocutor gives an argument, either in favor of the new proposal or reinforcing the original proposal, the explanation needs to be sensitive to the argument [23].

Our current approach uses a very simple template for counterproposing. First, an initial assessment is given of the proposal. If this assessment is negative, then the agent will give a reason why. There may be multiple reasons to disprefer an action, so it can be tricky to find the best explanation of the real issues involved. The COA assessments mentioned in Section 2.2 are used both to decide which COAs are better than others, but also to decide what the most important reason for this preference is. In the case of intrinsic utility, either because of a COA leading to an undesired state, or blocking the establishment of a desired state, that effect is described. In the case of states that have their utility in terms of enabling actions which will eventually lead to desired states, the ultimate goals are mentioned as well.

5. EVALUATION

Currently we have evaluated the model only within the peacekeeping domain, but even within this single domain, there are many sorts of variation that can be handled by our agents. E.g.,

- the user could propose or order any action (from a total of 40 represented by the task model and hundreds more that are representable in language but do not have methods for performing them in the domain.
- the user could use many different linguistic means for proposing the same act, e.g., "secure the assembly area", secure the area", "secure it" (after talking about the area), "do it" (after a suggestion to secure the area).
- a simulation operator/controller can change the conditions in the simulator, e.g., the severity of injuries of some characters,

1	Lt	Sergeant, secure a landing zone.
2	Sgt	Sir, first we should secure the assembly area.
3	Lt	Secure the area.
4	Sgt	Yes sir.
5	Sgt	Squad leaders listen up!
6	Sgt	I want 360 degree security here.
7	Sgt	1st squad take 12-4.
8	Sgt	2nd squad take 4-8.
9	Sgt	3rd squad take 8-12.
10	Sgt	4th squad secure the accident site
11	Squads	(move into position to secure the area)
12	Sgt	Johnson, send a fire team up to the square to secure an LZ.
13	Johnson	Yes sergeant.
14	Johnson	Sergeant Duran!
15	Johnson	Send your team up to the square and secure an LZ.

Figure 2: Dialogue Fragment 1

the degree of danger to friendly forces, the number, position, and friendliness of local civilians.

To see how the negotiation model works in practice, let us look at some examples of how the implemented Sergeant character negotiates within the peacekeeping scenario. In Figure 2, the lieutenant (Lt) orders the sergeant (Sgt) to secure a landing zone (1). This creates a committed stance of the Lt toward this act, thus authorizing the act, since the Lt is the authority. The sergeant intends this act, but the plan-state is premature, and so the sergeant counterproposes with the highest-level next step, securing the assembly area (2). The counterproposal creates both a disparaged stance of the Sgt toward securing the landing zone, but also an endorsed stance toward securing the area. The Lt then orders the Sgt to secure the area (3) (again committing to and authorizing this act). This time the plan-state is good, and so the sergeant accepts (4). Since the action is authorized, as the responsible party, the Sgt can carry out the act. This is a team act, so the Sgt must bring all the team members on board. This is done first by getting their attention (5), and then expressing the goal (6) (giving the Sgt a stance toward the squad leaders since his previous stance from the acceptance was just toward the Lt). Next, individual orders are given (7-10), which function the same way, one level down, as the orders between Lt and Sgt. In this case, rather than explicitly accepting, the squad leaders just order their troops to move into action (11), giving non-verbal evidence of having heard, understood, and accepted the order. Once the sergeant is satisfied that the area is secure, he can move on to a next task in his plan. In this case, securing the landing zone is appropriate, and since the Lt had already authorized it and had not removed his stance, the Sgt can carry out this action. Again, this is an abstract team task, and so the Sgt orders Johnson, the third squad leader, to secure the LZ (12). Similar reasoning leads Johnson to accept (13) and order one of his subordinates to move into action and secure the LZ (14-15).

Figure 3 shows another kind of example. Here, the Lt orders something that the Sgt considers to be a bad idea (1). Again the Sgt counterproposes, but here explicitly rejects the proposed action (2), explains that it will lead to an undesirable effect (3), and then proposes what he considers a better action (4). If the Lt orders this new action, things would proceed similarly to the flow in Figure 3. In this case, however, the Lt sticks to his original order. Now the Sgt sees the plan state as discussed and reluctantly accepts (6), as the only way to fulfill his obligation in the case in which the Lt will not reconsider. From this point, he adopts an intention to perform

1	Lt	Sergeant, send two squads forward.
2	Sgt	Sir, that's a bad idea.
3	Sgt	We shouldn't split our forces.
4	Sgt	Instead we should send one squad to recon forward along our route to Celic.
5	Lt	Send two squads forward!
6	Sgt	Against my recommendation sir.
7	Sgt	Lopez, send first squad and fourth squads to Celic!

Figure 3: Dialogue Fragment 2

the action, even though he knows it is not optimal. Given the intention, the ability to do the action next, and the authorization, he then carries out the action, again by ordering subordinates to carry out their parts of a team task (7).

So far, the model has mainly been evaluated with respect to coverage in the domain. We have begun pilot studies with target users (Army cadets with knowledge of the domain), who are able to negotiate some tasks successfully, but we have not yet reached the point of careful evaluation of many users in a real training situation.

6. CONCLUSIONS AND FUTURE WORK

In this paper, we have sketched the teamwork model of our virtual human agents, including the task model, the dialogue model, and aspects of the emotional evaluation model, that are integrated to allow complex team behavior, including negotiation and delegation. The model is currently implemented and used by agents in the MRE project involved in a peacekeeping training scenario. All of the reasoning described in this paper is domain-independent. Domain-specific features include the specific tasks and constraints, the pre-defined social relationships between the characters, and vocabulary items used in both speech recognition and speech synthesis. Future work includes further evaluation, as well as expanding the length of negotiation sequences, flexibility of explanation, and personality and emotion-driven responses, as well as also applying to new domains.

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.

7. 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, J. Nivre, and E. Ahlsen. On the semantics and pragmatics of linguistic feedback. *Journal of Semantics*, 9, 1992.

[3] R. Bindiganavale, W. Schuler, J. M. Allbeck, N. I. Badler, A. K. Joshi, and M. Palmer. Dynamically altering agent behaviors using natural language instructions. In *Proceedings of the Fourth International Conference on Autonomous Agents*, pages 293–300, New York, 2000. ACM Press.

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

[5] H. H. Clark. *Using Language*. Cambridge University Press, Cambridge, England, 1996.

[6] H. H. Clark and E. F. Schaefer. Contributing to discourse. *Cognitive Science*, 13:259–294, 1989.

[7] 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.

[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] C. A. Connolly, J. Johnson, and C. Lexa. AVATAR: An intelligent air traffic control simulator and trainer. In *Proceedings of the Fourth International Conference on Intelligent Tutoring Systems (ITS '98)*, number 1452 in Lecture Notes in Computer Science, pages 534–543, Berlin, 1998. Springer.

[10] P. Dillenbourg, D. Traum, and D. Schneider. Grounding in multi-modal task-oriented collaboration. In *Proceedings of the European Conference on AI in Education*, 1996.

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

[12] C. Eliot and B. P. Woolf. An adaptive student centered curriculum for an intelligent training system. *User Modeling and User-Adapted Instruction*, 5:67–86, 1995.

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

[14] J. Gratch. Émile: Marshalling passions in training and education. In *Proceedings of the Fourth International Conference on Autonomous Agents*, pages 325–332, New York, 2000. ACM Press.

[15] B. J. Grosz and S. Kraus. Collaborative plans for complex group action. *Artificial Intelligence*, 86(2):269–357, 1996.

[16] N. Jennings. Controlling cooperative problem solving in industrial multi-agent systems using joint intentions. *Artificial Intelligence*, 75, 1995.

[17] W. L. Johnson, J. W. Rickel, and J. C. Lester. Animated pedagogical agents: Face-to-face interaction in interactive learning environments. *International Journal of Artificial Intelligence in Education*, 11:47–78, 2000.

[18] 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.

[19] H. J. Levesque, P. R. Cohen, and J. H. T. Nunes. On acting together. In *Proceedings of the Eighth National Conference on Artificial Intelligence (AAAI-90)*, pages 94–99, Los Altos, CA, 1990. Morgan Kaufmann.

[20] S. Marsella and J. Gratch. A step toward irrationality: Using emotion to change belief. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems*, New York, 2002. ACM Press.

[21] C. Matheson, M. Poesio, and D. Traum. 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*, 2000.

[22] D. McAllester and D. Rosenblitt. Systematic nonlinear planning. In *Proceedings of the Ninth National Conference on Artificial Intelligence (AAAI-91)*, pages 634–639, Menlo Park, CA, 1991. AAAI Press.

[23] J. D. Moore. *A Reactive Approach to Explanation in Expert and Advice Giving Systems*. PhD thesis, University of California, Los Angeles, 1989.

[24] C. H. Nakatani and D. R. Traum. Coding discourse structure in dialogue (version 1.0). Technical Report UMIACS-TR-99-03, University of Maryland, 1999.

[25] D. Novick. *Control of Mixed-Initiative Discourse Through Meta-Locutionary Acts: A Computational Model*. PhD thesis, University of Oregon, 1988. also available as U. Oregon Computer and Information Science Tech Report CIS-TR-88-18.

[26] J. Rickel. An intelligent tutoring framework for task-oriented domains. In *Proceedings of the International Conference on Intelligent Tutoring Systems*, pages 109–115, Montréal, Canada, June 1988. Université de Montréal.

[27] 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.

[28] J. Rickel and W. L. Johnson. Extending virtual humans to support team training in virtual reality. In G. Lakemayer and B. Nebel, editors, *Exploring Artificial Intelligence in the New Millenium*, pages 217–238. Morgan Kaufmann, San Francisco, 2002.

[29] 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.

[30] B. Roberts, N. J. Pioch, and W. Ferguson. Verbal coaching during a real-time task. In *Proceedings of the Fourth International Conference on Intelligent Tutoring Systems (ITS '98)*, pages 344–353, Berlin, 1998. Springer.

[31] W. Swartout, R. Hill, J. Gratch, W. Johnson, C. Kyriakakis, K. Labore, R. Lindheim, S. Marsella, D. Miraglia, B. Moore, J. Morie, J. Rickel, M. Thiebaux, L. Tuch, R. Whitney, and J. Douglas. Toward the holodeck: Integrating graphics, sound, character and story. In *Proceedings of 5th International Conference on Autonomous Agents*, 2001.

[32] M. Tambe. Towards flexible teamwork. *Journal of Artificial Intelligence Research*, 7:83–124, 1997.

[33] 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.

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

[35] D. R. Traum and E. A. Hinkelman. Conversation acts in task-oriented spoken dialogue. *Computational Intelligence*, 8(3):575–599, 1992.

[36] 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.

[37] M. A. Walker and S. Whittaker. Mixed initiative in dialogue: An investigation into discourse segmentation. In *Proceedings ACL-90*, pages 70–78, 1990.

[38] B. Woolf, D. Blegan, J. H. Jansen, and A. Verloop. Teaching a complex industrial process. In *Proceedings of the Fifth National Conference on Artificial Intelligence (AAAI-86)*, pages 722–728, Los Altos, CA, 1986. Morgan Kaufmann.