Appears in Proceedings of the 10th Conference on Computer Generated Forces and Behavioral Representation, Orlando, FL pp.457-466

Modeling Emotions in the Mission Rehearsal Exercise

Jonathan Gratch
University of Southern California
Institute for Creative Technologies
13274 Fiji Way Marina del Rey, CA 90292
gratch@ict.usc.edu

Stacy Marsella
University of Southern California
Information Sciences Institute
4676 Admiralty, Marina del Rey, CA 90292
marsella@isi.edu

Keywords: human behavior moderators

ABSTRACT: This paper discusses our attempts to model realistic human behavior in the context of the Mission Rehearsal Exercise system (MRE), a high-end virtual training environment designed to support dismounted infantry training between a human participant and elements of his command. The system combines immersive graphics, sound, and interactive characters controlled by artificial intelligence programs. Our goal in this paper is to show how some of the daunting subtlety in human behavior can be modeled by intelligent agents and in particular to focus on the role of modeling typical human emotional responses to environmental stimuli.

1. Introduction

There is a growing acknowledgement that effective leaders are "emotionally intelligent" (Goldman, 1997). Emotions suffuse our social relationships. People pursue or avoid activities base on how they make us feel. We carefully attend to emotions expressed by those close to us to infer something of their internal state or assess how they are responding to our actions. We communicate emotion, consciously and unconsciously, that alter the behavior of those around us. Good leaders understand how to make use of this information - to build group cohesion, to motivate, and to reach a better understanding of those around them.

Emotions also play a critical role in creating engaging and believable entities to populate training simulations. Existing behavioral models the lack believability and realism, which is problematic even when modeling the behavior of airplanes and tanks (NRC, 1998), but is especially jarring is face-to-face interactions between agents and human participants. Modeling such humanistic aspects of military operations is a difficult challenge, but if successful, would significantly enhance the leadership training potential and realism of training simulations.

This paper begins to address this challenge. We describe an integration of two research efforts focused on creating realistic, engaging, and believable characters to populate training environments. Gratch's Emile system focuses on the problem of emotional appraisal: how emotions arise from an evaluation of how environmental events relate to an agent's plans and goals (Gratch, 2000). Marsella's IPD approach focuses more on the impact of expressions of emotional state through suitable choice of gestures and body language (Marsella et al. 2000). We describe how this combined approach contributes to the effectiveness of the Mission Rehearsal Exercise, a virtual reality training environment designed to exercise decision making in stressful circumstances at the level of dismounted infantry (Figure 1). In this environment, intelligent agents control characters in the virtual environment with which the participants must interact in the course of their training, and our emotional models attempt to augment the believability, realism and suspense of these interactions.

2. Mission Rehearsal Exercise

A young Army lieutenant drives into a Balkan village expecting to meet up with the rest of his platoon, only to find that there has been an accident. A young boy lies hurt in the street, while his mother rocks back and forth, moaning, rubbing her arms in anguish, murmuring encouragement to her son in Serbo-Croatian. The lieutenant's platoon sergeant and medic are on the scene.
The lieutenant inquires, “Sergeant, what happened here?” The sergeant, who had been bending over the mother and boy, stands up and faces the lieutenant. “They just shot out from the side street, sir. The driver couldn't see them coming.”

“How many people are hurt?”

“The boy and one of our drivers.”

“Are the injuries serious?”

Looking up, the medic answers, “The driver's got a cracked rib, but the kid's….” Glancing at the mother, the medic checks himself. “Sir, we've gotta get a medevac in here ASAP.”

Figure 1: Virtual Bosnian village

The lieutenant faces a dilemma. His platoon already has an urgent mission in another part of town, where an angry crowd surrounds a weapons inspection team. If he continues into town, the boy may die. On the other hand, if he doesn't help the weapons inspection team their safety will be in jeopardy. Should he split his forces or keep them together? If he splits them, how should they be organized? If not, which crisis takes priority? The pressure of the decision grows as a crowd of local civilians begins to form around the accident site. A TV cameraman shows up and begins to film the scene.

This is the sort of dilemma that daily confronts young Army decision-makers in a growing number of peacekeeping and disaster relief missions around the world. The challenge for the Army is to prepare its leaders to make sound decisions in similar situations. Not only must leaders be experts at the Army's tactics, techniques and procedures, but they must also be familiar with the local culture, how to handle intense situations with civilians and crowds and the media, and how to make decisions in a wide range of non-standard (in military terms) situations.

In the post-cold war era, peacekeeping and other operations similar to the one outlined above are increasingly common. A key aspect of such operations is that close interaction occurs between the military and the local population. Thus, it is necessary that soldiers understand the local culture and how people are likely to react.

Unfortunately, training options are limited. The military does stage exercises using physical mockups of villages with actors playing the part of villagers to give soldiers some experience with such operations. However, these are expensive to produce, and the fact that the actors must be trained and the sets built makes it difficult to adapt to new situations or crises. Computer-based simulators have been used by the military for years for training but these have focused almost exclusively on vehicles: tanks, humvees, helicopters and the like. Very little exists for the soldier on foot to train him in decision making in difficult circumstances.

The goal of the Mission Rehearsal Exercise (MRE) is to create a virtual reality training environment in which scenarios like the one described above can be played out. Participants are immersed in the sights and sounds of the setting and interact with virtual humans acting as characters in the scenario. At times, these characters may also act as coaches, dispensing advice to the trainee to help him achieve pedagogical goals. The underlying assumption is that people learn through experiencing a situation and making decisions in the context of a stressful and sometimes confusing environment.

Training goals require an emphasis on immersive reality and "broad agents" that integrate motor skills, problem solving, emotion, gestures, facial expressions, and language. The MRE system pushes the state-of-the-art in simulation technology through the integration of high-fidelity real-time graphics, intelligent agents, immersive audio and interactive story. An initial prototype of the system now exists and its improvement is a subject of ongoing research (reference pending). Intelligent agents control characters (virtual humans) in the virtual environment, playing the roles of locals, friendly and hostile forces, and other mission team members.

3. PLAN-BASED Emotions

To support the complexity of training environments such as the MRE, intelligent agents must incorporate a variety of capabilities. Here we focus on the issue of emotional modeling. Many psychological theories of emotion emphasize the relationship between emotions and cognition. How one responds to some external events seems closely tied to their implications for ones plans and goals (Ortony et al, 1988; Lazarus, 1991). Even purely mental "events" can evoke strong emotions: most of us have experienced a flash of insight in our research that leaves us with intense feelings of joy, only to be crestfallen seconds later by the realization of some crucial flaw. Emotions clearly have a strong influsome crucial flaw. Emotions clearly have a strong influence over our decision-making abilities as well (Damasio 1994; Sloman, 1987).

Gratch (2000) has argued that artificial intelligence planning techniques provide a powerful and general mechanism for modeling this interplay between cognition and emotion, at least with regard to "task-oriented" emotions (those emotions that arise from the performance of some concrete task). Adopting a plan-based approach has some key advantages. By maintaining an explicit representation of an agent's plans one can easily reason about future possible outcomes – a key requirement for handling emotions like hope and fear that involve future expectations. Explicit representations allow one to detect interactions between plans, for example, as when the plans of one agent are incompatible with those of another – a key requirement for handling emotions like anger or reproach which typically involve multiple actors. Planning algorithms provide general mechanisms for making these assessments that we can leverage this generality in creating a model of emotional reasoning.

A plan-based approach also allows a richer model of how cognition influences one's emotional state. We can model some of the dynamic ebb and flow of human emotion by relating emotional appraisals to the current state of plans in an agent's memory. As plans grow and change through the planning process, so too the emotional state will change as a reflection of this process – in a sense providing a window into an agent's mental processes.

Finally, by providing an explicit and rich reasoning infrastructure, plan-based approaches facilitate models of how emotions impact decision-making. Emotional state can act as search control, focusing cognitive resources on specific goals or threats. It can also alter the overall character of problem solving. For example, negative emotions seem to lead to narrow focused problem solving while positive emotions lead to broader problem solving that attempts to achieve multiple goals simultaneously (Sloman, 1987).

4. ÉMILE

Émile is an agent architecture that incorporates a planning system, a plan-based model of emotions, and a simple model of personality and social behavior. It is designed to model agents that can engage in social interactions in relatively structured domains. In this sense it fills the niche between more improvisational agents (Rousseau and Hayes-Roth, 1998) that support greater interactivity but have relatively simple behaviors, and purely scripted agents that can display quite complex behavior but no interactivity.

Émile is described in greater detail in (Gratch, 2000). To summarize, Émile builds on Clark Elliott's (1992) construal theory which assesses the relationship between events and an agent's disposition (described by its goals, social standards, and preferences) through a set of knowledge structures called *construal frames*. These frames characterize the relationship in terms of a set of abstract features which are then mapped to an emotional state. Émile adds a level of indirection that generalizes this process. Rather than appraising events directly, Émile appraises the state of plans in memory. The abstract relationship between events and an agent's disposition is derived in a domain-independent fashion by a general-purpose planning algorithm. Thus, Émile replaces a large number of domain-specific construal frames needed by construal theory with a small number of domain-independent rules. Domain-specific information, for the most part, is restricted to the operator descriptions (the domain theory) from which plans are built, and which an intelligent agent needs anyway to inform planning and action selection.

Émile also draws heavily on the explicit plan representation to derive the intensity of emotional response. Émile incorporates the view of Oatley and Johnson-Laird (1987) and Neal Reilly (1996) that emotions are related to changes in the perceived probability of goal attainment. Intensity is broken down into the probability of the event in question (e.g. the probability of goal achievement or the probability of a threat) and the importance (utility) of the event to the agent, both of which are derived from the current plan structure. As intensity is based on the current plans, the assessment is a reflection of their current state and changes with further planning.

At any moment in time, the appraisal mechanism will have produced a number of appraisals from the current plan structure. Émile uses Velásquez's (1997) Cathaxis model to integrate them into an overall emotional state and model the decay of emotional intensities. As long as individual appraisals persist, Émile decays their intensity by a constant rate. The integration model adds these decaying intensities into different buckets based on their emotional label. Thus, if one has several separate appraisals of "fear", these are added together into an overall fear intensity. Collectively, these buckets correspond to the current activation vector. According to the Cathaxis model, the activation of a given emotion. is excited by some emotional states and inhibited by others. These influences are expressed in a cross-correlation matrix.

5. Emotional Expression

Observers can reliably infer a person's emotions and attitudes from their nonverbal behaviors (Ekman et al. 1969) and therefore potentially respond in a variety of ways. Thus, when creating virtual humans that maintain and convey an internal emotional state, we must ensure that the agent's performance suggests a corresponding emotional state to the observer, or run the risk of creating confusion or disbelief.

People exhibit a wide repertoire of nonverbal behaviors consistent with their emotional state, some intentionally communicative, other not. For example, threatening glances or shaking a fist at someone play an intended role in communicating information to another person. On the other hand, behaviors such as rubbing one's thigh, averting gaze or a facial expression of fear may have no explicitly intended role in communication. Nevertheless, they do suggest considerable information about them, their emotional arousal, their attitudes and what they are attending to.

For our purposes, we need a model of agent behavior that appropriately suggests an emotional undercurrent. Such a model must address particular concerns. Of particular concern for the agent characters we design is that they provide convincing portrayals of humans facing difficult, dangerous problems. To that end, they must have emotionally revealing nonverbal behaviors and expressions consistent with deeply evocative/disturbing situations. These behaviors must also change in concert with the emotional state of the characters; obviously people express themselves differently when sad, happy or angry. Further, they must have behaviors unique to the individual since not everyone exhibits the same behaviors, in the same way.

Another key concern here is that the agent's mix of nonverbal behavior at any time appear emotionally consistent. Consider severe depression. There are many ways to convey severe depression; it may be effective for an agent to appear withdrawn, inattentive, or perhaps hugging themselves. However, if a supposedly depressed agent used various open, communicative gestures such as beats (McNeill) while expressing something to another agent, then the performance may not "read" correctly. The behavior may not appear consistent with depression. This is especially so if the agent had previously been exhibiting behaviors more consistent with depression. In fact, the mix of gestures used by an agent must be coherent and avoid unintended interpretations. For example, people don't tend to nonchalantly use deictic gesture while simultaneously averting their gaze due to mild feelings of anger or guilt. Such behavior may look unnatural, inconsistent, or may convey a different shade of meaning depending on context. Which is not to say that the overall mix of behaviors should always be monolithic. People do say one thing while expressing another. At the least, the mix of nonverbal behaviors often shade the meaning of what is said or communicated nonverbally. Returning to the previous example, if an agent does combine deictic gesture with gaze aversion, it may shade the interpretation dramatically, towards an expression of extreme emotion and a desire to control that emotion. For example, the agent is so disgusted with the "listener", they can't bear to look at them.

Implicit in these various concerns is that the agent has what amounts to a resource allocation problems. The agent has limited physical assets, e.g., two hands, one body, etc. At any point in time, the agent must allocate these assets according to a variety of demands, such as performing a task, communicating, or emotionally soothing themselves. For instance, the agent's dialog may be suggestive of a specific gesture for the agent's arms and hands while the emotional state is suggestive of another. The agent must mediate between these alternative demands in a fashion consistent with their goals and their emotional state

6. PHYSICAL FOCUS

To address these concerns, the emotional behavior component of this agent architecture relies on the Physical Focus model of the IPD system (Marsella et al. 2000). The IPD work was in turn heavily influenced by work on non-communicative but emotionally revealing nonverbal behavior (Freedman 1972) as well as Lazarus's (1991) delineation of emotion-directed versus problem-directed strategies for coping with stress.

The Physical Focus model bases an agent's physical behavior in terms of what the character attends to, how they relate to themselves and the world around them, specifically whether they are focusing on themselves and thereby withdrawing from the world or whether they are focusing on the world, engaging it. The intent of the model is to refine down all the variegated ways in which emotional state impacts the agent's nonverbal behavior into distinct modes of relating to the world that provide a consistent resolution of the resource allocation problem.

The choice of nonverbal behaviors is determined by the agent's Physical Focus mode, which characterizes the mix of behaviors exhibited by an agent. At any point in time, the agent will be in a specific mode based on emotional state that predisposes it to use particular nonverbal behavior in a particular fashion. Each behavior available to an agent is categorized according to which subset of these modes it is consistent with. Any specific nonverbal behavior, such as a particular nod of the head, may exist in more than one mode and conversely a type of behavior, such as head nods in general, may be realized differently in different modes. Transitions between modes are based on emotional state.

By grouping behaviors into modes, the physical focus mode attempts to mediate competing communicative and non-communicative demands on an agent's physical resources,

```
defPlan handle-accident
  tasks: {accident, lt-arrives, evaluate, implore, evacuate, move-out, reassure, treat}
  causal constraints:
    {accident    {disables child-healthy}              end-handle-accident}
    {evaluate    {disables facilities-ok}              end-handle-accident}
    {move-out    {disables troops-helping}             evacuate}
    {treat       {enables  child-healthy       0.4}    end-handle-accident}
    {implore     {enables  help-requested}             reassure}
    {lt-arrives  {enables  authority-present 0.7}      evacuate}
    {lt-arrives  {enables  authority-present 0.7}      implore}
    {lt-arrives  {enables  authority-present 0.7}      treat}
    {evacuate    {enables  facilities-ok      0.65}    treat}
    {reassure    {enables  troops-helping      0.5}    evacuate}
 ordering constraints:
    {accident before move-out}
    {accident before treat}
    {lt-arrives before reassure}
 role assignments:
    {mother  implore}
    {lt      accident move-out reassure {evacuate lt}}
    {sgt     lt-arrives}
    {medic   treat evaluate}
    {ldr}

defGoal child-healthy
    "(<percept> ^boy-health good)"  "the child to be healthy"
     :probability 0.2 :location victim :concerns {{mother 80.0} {lt 40.0}}
```

Figure 2: A portion of the mother's domain knowledge

As originally developed, the Physical Focus model was designed to take input from both an agent's cognitive and emotional model. Based on that input, and the subset of behaviors consistent with the current mode, a behavior is selected. In addition, the current mode also was used as feedback to the cognitive model. In the current effort, the Physical Focus model was employed somewhat differently, due to underlying differences in the agent model used here as well as difference in the realization of the agent bodies. We will return to this discussion in the section on integration issues.

7. STEVE

The three interactive agents in the MRE scenario are modeled using the Steve system of Rickel and Johnson (1998), and have been integrated with greatly improved body and enables the condition that authority is present with 70% probability, which is a precondition of treating the victim). Finally, the schema specifies which agents are responsible for executing which tasks (the medic is responsible for evaluating and treating the child). The figure also illustrates how one defines conditions used as preconditions or effects of plan steps. "Child-healthy" is a proposition that is true if the perceptual state indicates that the "boy-health" attribute has a value of "good". The system *a priori* expects with twenty percent likelihood that this goal, should it be unsatisfied, can be attained. The location attribute tells Steve where to look or gesture when referring to this condition. Finally one can specify a set of agents who are concerned with the truth value of this condition and the utility they place on it being satisfied (the mother cares a lot about the boy being healthy). This information is used to infer the intrinsic and extrinsic utility of goals.

8. EXPRESSIVE CHARACTERS

The Physical Focus routines interface with human avatars modeld in Boston Dynamics, Inc.'s PeopleShop run-time environment. PeopleShop provides body models that can be either pre-scripted or controlled in real-time through an API. Character animation is based on motion capture: an actor wearing special sensors is recorded performing certain actions and this data is carved into segments and played back on demand. Boston Dynamics worked under subcontract to provided a number of custom features and behaviors including procedural control of gaze and the integration of their software with face models provided by another corporation, Haptek, that provides procedural control over facial expressions.

Motion capture is good for creating natural body movements but it is rather awkward to use in conjunction with our reasoning and emotional models. It is inflexible and you have to anticipate in advance all of the actions and gestures that you will require for the scenario. This inflexibility is especially problematic for our emotional models. A character's motions and gestures should change noticeably as a function of the current emotional state. Ideally, we could procedurally adjust the behavior in real-time. In fact, some research have begun to explored how to alter motion-capture in just such a fashion (Chi et al., 2000). Until such technology is available, our solution has been to carefully organize the motion capture segments to get the desired flexibility and range of emotional expression.

Figure 3 illustrates the set of motion capture segments. They have been organized into a finite-state machine that is loosely organized as hub-and-spoke. Hubs are a set of stationary body poses that correspond to the three Physical Focus modes: body, transitional, and communicative. Spokes are various behavior segments that transition from a hub, through a sequence of movements, then back to the hub. Behaviors are further divided into task related behaviors (such as imploring the lieutenant to stay) and non-task related behaviors (such as rocking back and forth). Behaviors generated call-backs to the agent, informing it when the behavior is complete and what state the body is in.

When selecting a behavior, the agent compares the current body state, emotional state, physical focus, and whether a behavior is currently executing. Some behaviors, such as task related behaviors and reactions to perceptual events (look at an explosion) have precedence and interrupt other ongoing behaviors. If neither of these behaviors are pending, the system chooses some behavior that is consistent. Responses to external events are further modulated by Physical Focus (the mother doesn't respond to low intensity perceptual events when in body focus). In some cases multiple behaviors apply (a resource conflict). Soar provides a general arbitration scheme to resolve such conflicts.

9. INTEGRATION ISSUES

Steve is designed to model team behavior, however, in this scenario the mother and the soldiers, while sharing some similar goals, would hardly be described as being on the same team. In particular, they have expectations about the desired course of events. We chose to model this by providing different domain knowledge to the mother and the soldier agents. The models are similar and refer to many of the same tasks and perceptual events, but this allows the mother to have a different understanding of the flow of events. For example, the mother understand the soldiers plans in much less detail and in one case mis-interprets the intent of on of the soldiers actions (when the lieutenant sends some squads forward to reinforce the other platoon – "move-out" – the mother infers that the troops are no longer helping her child – {disables troops-helping}).

Some software modifications were necessary to integrate Steve with Émile. Steve's representation language had to be extended to represent the probabilities and utilities needed for Émile to calculate the intensity of certain emotional responses. Steve also had to be extended to infer that certain tasks could disable conditions needed by other tasks (after the medic evaluates the child it is clear that the facilities are inadequate to treat the child without evacuating him to another location). We also slightly changed how Steve processes information, essentially slowing its reaction time to allow for greater visualization of the changes in mental state. Finally, we incorporated some knowledge from Émile's planning system to allow Steve to detect un-planned for perceptual events and express an appropriate startle reflex.

Some changes were also needed to integrate IPD's Physical Focus into the current system. The original body models in

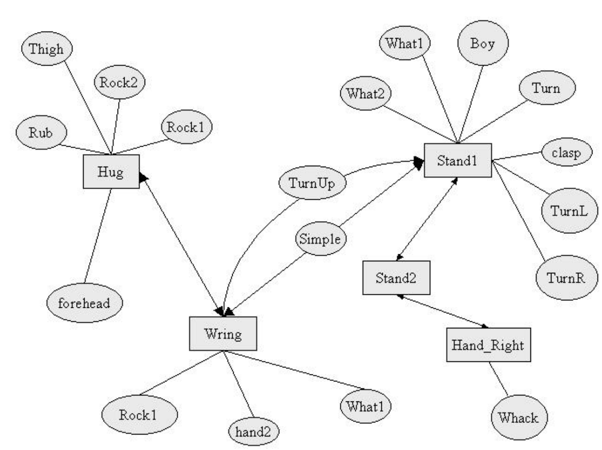

IPD were two dimensional, composed of many roughly orthogonal parts (hands, arms, etc) that could be separately animated. In MRE, the animation is three dimensional, far more realistic looking, though much more constrained as it is based on motion capture. This led to several simplifications. Most notably, because of the reduced flexibility of motion capture, and consequently the reduced need to manage the agent's behavior, we chose to reduce the number of modes to three. These modes then served to drive our specification of what behavior to capture. In the IPD system, Physical Focus was always engaged, mediating behavior and attention. However in the implementation, Physical Focus made behavior choices only when the agent was not explicitly engaged in some task. However to augment the impact of Physical Focus, we modified the underlying Steve system so that when performing a task, the selection of the specific behavior that achieved the task could be determined by the physical focus mode. As an example, the mother will implore the Lieutenant to help her child different when in communicative mode as opposed to transitional mode. In additional, we added the ability to react or not to react to unexpected events in the environment based on physical focus. For instance, when in body focus mode the mother is less attentive to minor events that occur in the environment.

Physical Focus also requires an appraisal of anxiety, which Émile did not support. According to most psychological theories, anxiety is treated as a non-specific threat to a goal in contrast to fear, which is treated as a specific threat. Émile previously only considered specific threats in its models (i.e., one task has an effect that disables a precondition of some other task). In the current implementation, we use the probability model to infer non-specific threats. If a task achieves predicate P with some probability less that 1.0, there is a non-specific threat to the achievement of P. It is non-specific in the sense that the goal may not be achieved, with probability 1-Pr(P), yet there is no explicit reason why not (as opposed to a goal which has a low probability of achievement because an anticipated task disables it with high probability). This covers anxiety arising from non-specific threats to goal achievement, but does not account for other sources of anxiety, for example non-specific threats to already achieved goals. A more complete model of anxiety is the subject of future work.

10. ILLUSTRATION

We now walk through some of the key points of the scenario as they relate to the mother to illustrate how the emotional model influences her behavior. In the opening scene, the mother is waiting for the lieutenant to arrive, which she views as a precondition for her child to be treated. She as somewhat angry at the lieutenant as she perceives him are responsible for the accident (as the lieutenant is assigned the role of executing the "accident" task). Initially she believes the facilities-ok is satisfied, meaning she has the simple plan in memory that the lieutenant should arrive and her child will be treated, neither task being under her control. Since her child is hurt she has high levels of distress. Since the lieutenant arriving and the treatment tasks have low probability effects (non-specific threats), she is also extremely anxious, though also somewhat hopeful. The high anxiety leads her to have an inner-directed Physical Focus. Her body gestures are directed inward and she will not attend to most stimuli.

When the lieutenant arrives in his jeep the mother perceives that "authority-present" is now satisfied in the current state. As this subgoal is now attained, the non-specific threat associated with its attainment disappears, the probability that the child will be treated increases somewhat, and the mother's anxiety and distress diminish somewhat. This is enough to transition her into transitional focus, her gestures become more outward directed and she attends to more perceptual stimuli and her child.

The lieutenant asks for a report of the child's health. The mother attends to this exchange and essentially eavesdrops on the medic statement that the facilities are inadequate. Steve's reasoning mechanism infers the current plan is invalid and that the child must now be evacuated. This change in plans leads to a change in evaluation of her goals and thus a change in emotional state. She lowers her estimate that the child will be successfully treated and the evacuation introduces several new sources of anxiety. She transitions back to body focus, which is articulated physically through visible and audible weeping.

Later in the scenario, the lieutenant orders one or two squads forward ("move-out") to reinforce the platoon downtown. The mother interprets this as disabling her subgoal that the troops are helping her child. The strength of this interpretation is influenced by the number of squads that move forward (implemented by domain-specific rules that infer conclusions from the agent's perceptual input). The emotional model treats this as a blameworthy event, causing the mother to become angrier at the troops. This anger is sufficient to transition her into communicative mode. The mother also updates her plans, deciding that the troops will return to helping her child if she implores them to stay (via the "implore" task). Her body language in performing this action is colored by her body focus and anger level, either remaining seated and gesturing mildly or raising to a standing position and gesturing strongly (Figure 4)

11. Discussion

This project is still in its early stages (the initial prototype was completed at the end of September 2000). From a research perspective the biggest limitation is the lack of evaluation. Is MRE a viable learning environment? Does the addition of emotional models increase the realism of the scenario? Do people find the character's reactions plausible? How do emotional models impact the learning experience? Our plan is to begin formal evaluations in the coming year in conjunction with other research groups in the psychology and communications departments at the University of Southern California. Our anecdotal feedback has been encouraging. We have demonstrated the system to a number of military personal and those who served in Bosnia or Kosova seemed strongly affected by the experience. One U.S. Army Colonel began relating a related incident after seeing the demo, became quite emotional, and concluded by saying, "this system makes people feel, and we need that." In another anecdote, someone playing the role of the lieutenant became agitated when the mother character began yelling at him and when she wouldn't respond to his reassurances (she cannot be mollified when her anger exceeds some threshold.

While this is encouraging, a number of problems must be addressed before we can exercise the MRE systems potential as a learning environment and evaluate its effectiveness. The prototype is not very interactive. Although the system uses speech recognition, the recognition grammar is quite limited. Furthermore, while there is some variability in the order events can occur, the scenario is essentially a linear narrative with one branch point (based on how many squads the lieutenant sends to reinforce the other platoon). As such, the scenario does not exercise the flexibility of our emotional models, and provides little evidence that the emotional responses would appear appropriate over a wider range of interactions. Before performing any rigorous evaluation we need allow the student to exercise more flexibility by adding domain knowledge to cover other possible decisions. Steve's reasoning capabilities will also have to be augmented as Steve has been designed to teach a single correct procedure (e.g. how to repair an engine) rather than a range of possible alternatives. This lack of alternatives also makes it difficult to model the impact of emotional state on decision making, which is most naturally encoded as some preference over alternative courses of action.

Another limitation is our current reliance on motion-capture data for the motions and gestures of the animated characters. Motion capture generates fluid and realistic motion but it is not well suited for real-time interactions. Our solution – a hup and spoke model with short motion-capture segments – allowed us to express some of the dynamics of the mother's emotional state, but there is no substitute for procedural control. As a solution we propose to integrate our work with Badler's EMOTE system (Chi et al., 2000). Emote can procedurally "morph" motion capture date along a number of dimensions, making a gesture seem to have more or less energy and gestures to be directed more inward or outward, much as is advocated by Marsella's Physical Focus.

Finally, there are a number of limitations in how the system infers emotional state that need adjustment or re-thinking in light of this application. One key issue is the notion of responsibility. For example, whom should the mother blame for the accident? The troops? Herself? Our sense is she should have a shared sense of responsibility and that this sense should change dynamically, influenced by her emotional state and subsequent actions of the troops. Currently, we simply use Steve's responsibility constraints to assign blame. Our treatment of anger is also too simplistic. Anger seems influenced by the extent to which we decide someone intended the offending action and the extent to which they show remorse or attempt to redress the offence. We suspect the explicit use of plans can assist in forming such assessments, but we still sorting out how.

These limitations not withstanding, the integration of plan-based appraisal of emotional state with the model Physical Focus provides a great deal of architectural support for emotional modeling. Furthermore, anecdotal evidence suggests that people not only find the agent's emotions to be plausible, but, to our surprise, people occasionally responded emotionally to our agents.

Acknowledgements

This research was funded by the Army Research Institute under contract TAPC-ARI-BR and by the U.S. Army under contract to the Institute for Creative Technologies. The content of this article does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred

References

Beaudoin, L. 1995. *Goal Processing in Autonomous Agents*. Ph.D Thesis (CSRP-95-2), Univ. of Birmingham.

Berkowitz, L. 2000. *Causes and Consequences of Feelings*. Cambridge University Press, 2000.

Bower, G.H. 1992 How might emotions affect learning? In Christianson (ed), *Handbook of emotion and memory*. Erlbaum

Cassell, J. & Stone, M. 1999. Living Hand to Mouth: Psychological Theories about Speech and Gesture in Interactive Dialogue Systems. AAAI Fall Symposium on Narrative Intelligence.

D. Chi, M. Costa, L. Zhao, and N. Badler: "The EMOTE model for Effort and Shape," ACM SIGGRAPH '00, New Orleans, LA, July, 2000, pp. 173-182,

Damasio, A. 1994. *Descartes' Error*. Avon Books, NY, NY.

de Souza, R. 1987. *The rationality of emotion*. MIT Press.

Elliott C. D. 1992. *The Affective Reasoner: A Process Model of Emotions in a Multi-agent System*. Ph.D Thesis (TR#32), Northwestern University.

Elliott C. & Siegle, G., 1993. *AAAI Spring Symposium on Reasoning about Mental States*.

Ekman, P. and Friesen, W.V. The Repertoire of NonVerbal Behavior: Categories, Origins, Usage and Coding. *Semiotica* 1:49-97, 1969.

Fikes, R.E. & Nilsson, N. J. 1971. STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving. *Artificial Intelligence*, 2(3-4)

Freedman, N. The analysis of movement behavior during clinical interview. In *Studies in Dyadic Communcation*, 1972, 153-175.

Frijda, N.H. & Swagerman, 1987 J. Can computers Feel? Theory and Design of an Emotional System. Cognition and Emotion, 1 (3) 235-257.

Gratch, J., 1998. Reasoning about multiple plans in dynamic multi-agent environments. AAAI Fall Symposium on Distributed Continual Planning, Orlando, FL.

Hill, R., Chen, J., Gratch, G., Rosenbloom, P., & Tambe, M. 1997. "Intelligent Agents for the Synthetic Battlefield," in AAAI-97/IAAI-97, pp. 1006-1012.

Hudlica, E., & Billingsley, J., 1999. Representing behavior moderators in military human performance models. *8th Con./ on Computer Generated Forces and Behavioral Representation*.

Johnson, L., Rickel, Jo. & Lester, J., 2000. Animated pedagogical agents: face-to-face interaction in interactive learning environments. *International* Journal of *Artificial Intelligence in Education. Forthcomming*

Lazarus, R.S. 1991. *Emotion and Adaptation*. Oxford Press.

Lester, J., Towns, S.G.,FitzGerald, P.J. 1999 Achieving Affective Impact: Visual Emotive Communication in Lifelike Pedagogical Agents. *International Journal of AI in Education, 10 (3-4)* pp. 278-291.

Marsella, S. Johnson, W.L. & LaBore, C. 2000. Interactive Pedagogical Drama. Proceedings of the Fourth International Conference on Autonomous Agents. Pp 301-308.

McNeil, D. *Hand and Mind*. University of Chicago Press, Chicago IL, 1992.

National Research Council (NRC), 1998. *Modeling Human and Organizational Behavior*. National Academy Press, Washington D.C.

Neal Reilly, W.S., 1996. *Believable Social and Emotional Agents*. Ph.D Thesis CMU-CS-96-138. Carnegie Melon Univ.

Newell, A. 1990. *Unified Theories of Cognition*. Harvard Press.

Oatley, K. & Johnson-Laird, P.N. 1987. Towards a Cognitive Theory of Emotions. *Cognition and Emotion*, 1 (1).

Ortony A., Clore, G. L., & Collins, A. 1988. *The Cognitive Structure of Emotions*. Cambridge University Press.

Picard, R. 1997. *Affective Computing*. MIT Press.

Rickel, J. & Johnson, L. 1998. Animated agents for procedural training in virtual reality: perception, cognition, and motor control. *Applied Artificial Intelligence, v 13..*

Rousseau, J., 1762. *Émile*. Reissued, Everyman Press, 1993.

Rousseau, D. and Hayes-Roth, B. A Social-Psychological Model for Synthetic Actors *Proceedings of the Second International Conference on Autonomous Agents*, Minneapolis, MN, May 1998.

Sloman, A. 1987. Motives, mechanisms and emotions. *Cognition and Emotion*, 1, pp 217-234.

Tversky, A. & Kahneman, D. 1983. Extensional versus intuitive reasoning: the conjunction fallacy in probability judgement. *Psychological Review, v. 90 n.4.*

Velásquez, J. 1998. When robots weep: emotional memories and decision-making. *AAAI98*.

Young, M. 1999. Cooperative plan identification: constructing concise and effective plan descriptions. *AAAI99*.