Do You Want To Talk About It?
A First Step Towards Emotion Integrated Dialogue

T. J. Muller, A. Hartholt, S. Marsella, J. Gratch, and D. Traum

University of Twente and Univesity of Southern California

Abstract. In this paper, we describe an implemented system for emotion-referring dialogue. An agent can engage in emotion-referring dialogue if it first has a model of its own emotions, and secondly has a way of talking about them. We create this facility in the MRE Project's virtual humans, building upon the existing emotion and dialogue facilities of these agents.

1 Introduction

It is widely recognized that in order to make agents more human-like, the use of emotions is essential. Although some work has been done in giving agents emotions and having those emotions influence the agent's behavior and dialogue, few systems offer the agents a chance to actually talk about their emotions, to say they are sad for example. We call this *emotion-referring dialogue*.

In this paper, we present a first step towards emotion-referring dialogue. As a vehicle, we use the Mission Rehearsal Exercise (MRE) immersive learning environment. MRE features virtual humans, autonomous agents the user can enter into conversation with [Rickel et al., 2002]. The virtual humans possess a deep process model of emotion, EMA, that not only derives a current emotional state, but represents an "explanation" of how the emotion arose from features of the social and physical environment and an agent's beliefs, desires and intentions [Gratch and Marsella, 2004] [Marsella and Gratch, 2003]. Till now, an agent's emotions could be conveyed non-verbally by gestures, body language and facial expressions and also influence both the interpretation and generation of dialog. However, the agent could not explicitly communicate its emotional state or respond to questions about it. We show how emotions are made explicit by implementing a question-answering capability, enabling the agent to verbally express his feelings when confronted with questions like "How do you feel?" and "Why are you mad, it's such a lovely day?".

2 Emotion-referring Dialogue

A first step in designing emotion-referencing dialogue is deciding which kinds of information could be conveyed. Following appraisal theory [Lazarus, 1991], emotions in EMA arise from an agent's interpretation of its relationship to the physical and social environment. "Cognition" informs this interpretation by inferring (via domain-independent reasoning mechanisms) how events impact the

agent's goals, beliefs, etc. For example, some consequence of an action could, with some probability, inhibit or facilitate a goal, leading to joy or distress. These events and their implications are *appraised* in terms of a number of specific dimensions, including *desirability*, *likelihood*, *blame*, *sense of control*, etc. These appraised dimensions then determine an instance of emotion of a specific *type* (anger, distress, etc.) and *intensity*. This multi-level representation maintains not only the top-level emotional response, but its justification in terms of appraisal dimensions and lower-level domain propositions and relationships.

The original dialogue capability in MRE virtual humans was designed to support team task interaction, including orders and negotiation over actions, and questions and answers about actions and states. For example, (1) is the representation of a state proposition (a boy is not healthy). (2) is the representation of a question "who is hurt", which can be answered using knowledge such as (1).

(1) (C262 ^attribute health-status ^object-id boy ^polarity negative ^time present ^type state ^value healthy)

(2) a. (S8290 ^prop P7433 ^q-slot object-id ^type question)
b. (P7433 ^attribute health-status ^polarity negative ^time present ^type state ^value healthy)

As a step towards talking about emotions, we derive emotion state representations like (3) from EMA. Using these state representations required some extensions to the `current-state` model. Previously, these state objects were present throughout the run-time of the agent, with different belief values (true, false, or unknown). While this is feasible for a fairly small closed task, it is difficult for representing emotions, since the intensities of different emotions take on continuous values, with many changes as time passes and events occur. Thus we have created new *on-the-fly* states which only exist while the question is under discussion. These represent the emotion type and intensity, but also the justification structure underlying the emotion instance. This allows the current dialogue question answering mechanisms [Traum, 2003] to be used as-is, without creating an unwieldy number of extra states.

(3) (C262 ^attribute max-feeling ^object-id <self> ^type state ^value <emotion>)

As we are still experimenting with the appropriate character of emotion-referring dialogue, and as the existing natural understanding and generation routines do not support such dialogue, we have currently constructed special purpose techniques to handle these capabilities. Rather than using the existing semantic parser, which would have required either adding to a finite state grammar, or providing training data, or both, a keyword spotting parser was built specifically to look for emotion-related questions. For every question a set of both present and missing keywords is defined, which uniquely identifies that question. The input is scanned for these keywords, and if a sufficient match is found, a semantic representation similar to (3) is created. This is then used to match against possible states that could be used to answer the question. For

now, most of the semantics will result in the generation of just one reference. When ambiguity about the user's intention arises, rules will select the most appropriate state: all possibilities that do not lead to an emotion with a certain threshold intensity will be discarded, and if multiple high intensity states exist, one will be picked randomly. The system will then update the dialogue state, after which the natural language generation phase is entered, when the dialogue manager decides to answer the question.

In the generation phase, again we implemented a special-purpose routine rather than create the knowledge need to do full linguistic realization. We designed a template system that could take the values in the emotion state to chose specific filler words for slots rather than modifying the existing realization system to be able to produce emotion descriptions. The emotion state serves as the main source of information, providing, where available, the emotion type and intensity, and additional information about the associated state. Special look-up tables translate these information bits into strings of natural language.

3 Results

The implemented emotion dialogue system allows the user to ask a variety of questions, e.g. which emotion an agent feels most, if he feels a certain emotion, how he feels about a certain state, what causes him to feel a certain way, etcetera. These types of questions can be asked in various ways. For instance, the question "Are you worried?" could also be rephrased as "Do you feel concerned?" or "You seem upset." – all map to the same semantics. They are answered depending on the current state of the system. If, for instance, the user wants to know who is making the agent feel a certain emotion, the agent could answer that he in fact is not feeling that emotion at all, point to a certain person, blame a certain person, or state that no one in particular is responsible. Personality also plays a role, particularly in the amount of information an agent conveys. For example, if the agent feels worried (for instance, because the agent is helping an accident victim that may die), the question "Are you worried?" could be answered with "Yes, sir.", but another possible answer is "Yes, sir, because the boy is probably going to die.", depending on certain system parameters that modulate the generation of dialogue.

(4) illustrates a dialogue between the implemented system and a user (acting as the system's boss).

	User	**Agent**
	What happened here?	There was an accident, sir.
(4)	Who's hurt?	The boy and our driver.
	How do you feel about the boy?	I'm feeling worried about him.
	Why?	The boy has critical injuries, sir.

4 Conclusion and Future Research

The key lesson from this work is that having a deep model of emotional reasoning can straightforwardly enhance the naturalness and expressivity of automated di-

alogue systems. With a natural language framework in place, we showed that an already existing emotion structure can be used for an emotion-referring dialogue with relative ease. The mentioned question-answer structure enables the user to ask an agent a variety of questions about it's internal state, getting feedback in a natural way. Although we are still far from a complete natural implementation, this first step gives us valuable insights on how to proceed. New utterances can be added quite easily by adding new emotion states and natural language rules, resulting in the possible use of utterances like "Why are you mad at the mother?" The use of keyword scanning and templating delivered relatively quick results when using a small, known domain and enabled us to make easy changes as there's no need for training natural language modules.

The current implementation of the emotion dialogue has several limitations. Some could be eliminated by integrating it fully into the MRE system and making more use of the richness of the dialogue and emotion models. Others require further developments to the emotion and dialogue modules, including appraisals of obligations, justifications, (see [Mao and Gratch, 2004]), excuses, and questions about coping strategies.

The following dialogue, suggested by a anonymous reviewer, provides a good starting point to illustrate these limitations:

	User	**Agent**
	How do you feel about me?	I'm feeling angry with you
	Why?	you didn't come to the meeting.
(5)	But it was not my fault! I had to finish some urgent work. Why do you get angry for such a stupid thing?	Because I'm a serious and careful person: I hate waiting.

The first two exchanges could be supported with minimal modification. The first involves collecting all emotion instances that make reference to the user (e.g., as an agent deserving blame) and reporting the emotion category of the most intense instance. A slight generalization of the "on-the-fly" state representation discussed above would suffice. The second exchange is handled without modification, assuming the task model included an "attend meeting' action in the domain model that the user failed to execute. Implicit in this dialogue is that the user had an obligation to attend the meeting (probably acquired through dialogue). Although the current dialogue system maintains such obligations, they are not as yet appraised directly by EMA. See [Mao and Gratch, 2004] for some developments along these lines.

The last exchange involves three utterances by the user: a rejection of blame, a justification of this rejection, and a disparagement of the system's prior attribution. The first two sentences are representable by recent extensions of the emotion module proposed by Mao [Mao and Gratch, 2004], though none of them are correctly parsed by the understanding module, nor do the underlying cognitive or dialogue modules provide a means of reconciling contradictory justifications. Our work on negotiation might be extensible along these lines.

The agent's final response could be handled without too much violence by encoding a "not wait" goal that is threatened by the user's failure to attend the meeting. The justification in terms of an appeal to personality traits, however, exceeds the scope of our current emotion module.

A remaining issue concerns the naturalness of the resulting dialogue. For a variety of reasons, people often do not communicate true emotional state unaltered. People may be in denial about their actual state, or may chose to express something other than their actual feelings to shift-blame, manipulate social interactions or seek social support. EMA explicitly models such emotion-induced *coping strategies*, and thus provides modeled a natural mechanism to incorporate such an "intentional stance" toward expressing emotion. For example, an agent in MRE may cope with anger by applying a coping strategy of denial. When asked about this emotion, a natural response might be, "I'm not angry!"

Acknowledgements

We thank the MRE team for providing an environment in which our research could take place. Anton Nijholt contributed valuable insights. This work was partly funded by the Department of the Army under contract DAAD 19-99-D-0046. Any opinions, findings and conclusions or recommendations expressed are those of the authors and do not necessarily reflect the views of the Department of the Army.

References

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

[Lazarus, 1991] Lazarus, R. (1991). *Emotion and Adaptation*. Oxford University Press.

[Mao and Gratch, 2004] Mao, W. and Gratch, J. (2004). Social judgment in multiagent interactions. In *to appear in proceedings of AAMAS 2004*.

[Marsella and Gratch, 2003] Marsella, S. and Gratch, J. (2003). Modeling coping behavior in virtual humans: Don't worry, be happy. In *In proceedings of AAMAS 2003*.

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

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