Building Emotional Agents

W. Scott Reilly Joseph Bates

May 1992
CMU-CS-92-143

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Abstract

The Oz project is developing technology for high quality interactive fiction and virtual realities. An important aspect of this research is the development of intelligent, emotional agents to occupy these micro-worlds. We present here some preliminary work on Em, a implemented model of emotion for use in these agents. Em is based on the theoretical, cognition-based emotion model of Ortony, Collins, and Clore and has shown promise of being a flexible model of emotion.

This work was supported in part by Fujitsu Laboratories, Ltd. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of Fujitsu Laboratories or any other parties.

Keywords: artificial intelligence, agent architecture, autonomous agents, emotion, social knowledge, interactive media, virtual reality, Oz

1 Introduction

The Oz project at the Carnegie Mellon School of Computer Science is developing technology for high quality interactive fiction and virtual realities. Our goal is to provide users with the experience of living in dramatically interesting micro-worlds that include moderately competent, emotional agents. An Oz world is composed of (1) a simulated physical environment, (2) the agents which live in the environment, (3) a user interface and its associated theory of presentation and style, and (4) a theory of drama, which plans and gently controls the overall flow of events in the world. The goals and overall method of Oz is described by Bates [1].

Our goal is to allow users of the system to suspend their disbelief when interacting with an Oz world. In order to achieve this, the agents that populate these worlds must give the illusion of being real in a number of ways (at least with respect to the micro-world in which they exist). These agents must display competent action in the world, reactivity to a changing environment, goal-directed behavior, appropriate emotions, individual personality traits, social skills, language skills, and some degree of intelligent inferencing about the world around them. In order to create the illusion of reality, we feel it is more important to produce a broad range of abilities that are integrated into a coherent system than it is to perfect any individual aspect of the agents. For this reason we call our agents "broad agents" [2].

We are developing an architecture, named Tok, in which to build this broad range of capabilities, and are developing the Em subsystem to handle most of the emotional and social aspects of the agent's behavior. Em's model of emotion is based largely on Ortony, Collins, and Clore's cognitive model of emotion [12], and has demonstrated some ability to produce reasonable emotional behavior for Oz agents. Em can also model simple social relationships between agents, how these relationships change over time, and how these relationships interact with emotion and behavior.

To integrate Em with the rest of the architecture, Tok uses a system of behavioral features that are influenced by the agent's emotional state and which in turn influence many aspects of the agent's behavior. Combining Em and the behavioral features system shows promise of letting us model certain aspects of personality for Oz agents.

After sketching the structure of Tok, we will present Em and the model that it is based on in some detail. Then we will discuss how Em is integrated into the Tok architecture. After that we will look at how Em models social relationships. We will conclude by looking briefly at evidence that this overall framework will be both fairly culturally independent and flexible, which suggests that Em can support reasonable imitations of the emotional systems of most human or artistically conceived cultures.

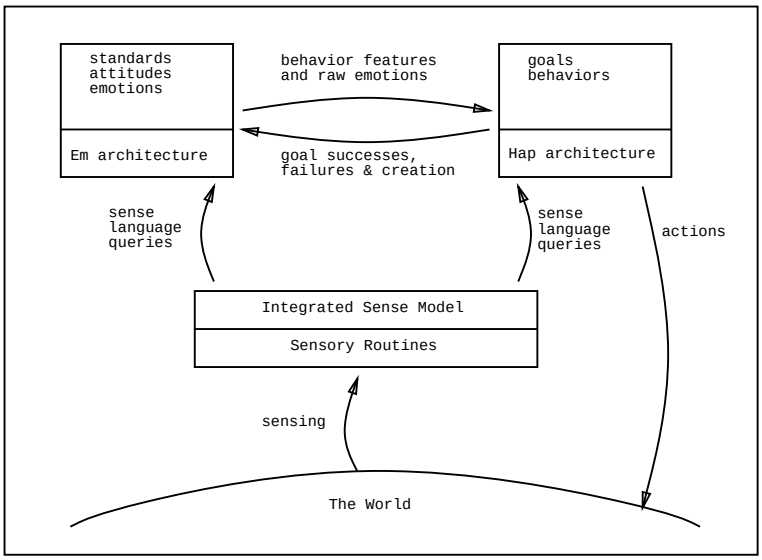

Figure 1: Tok Architecture

2 Overview of Tok

Tok contains a number of integrated components [4, 8, 9]. To explain these components, and particularly Em, we will look in abstract terms at how the entire architecture works (See Figure 1) [1]. The Tok agent is embedded in a physical world simulation and repeats a cycle of sense-think-act for the duration of the simulation. On a given cycle, the agent first receives sense data. This data is recorded in raw form by the Sensory Routines and then may be integrated into a partial internal world model by the Integrated Sense Model.

Once this data is integrated, Em generates emotions based on the agent's internal state and the currently sensed environment. Any time Em changes the emotional state, it informs the other subsystems of Tok. This includes notifying the part of Em that manages the agent's interpersonal relationships. Once Em finishes its initial update, the Hap system takes control. Hap is the behavioral engine for the Tok architecture. Hap chooses an action to perform based on the agent's physical environment, its active goals, the importance of those goals, what it had been working on previously, the emotional state, and personality constraints. For more details on Hap, see [9].

While processing, Hap may create new goals or discover that old goals have suc-

[1] This is a description of the Lisp-based version of the Tok architecture. We are also at work on a C-based version of the architecture that is used to drive creatures we call Woggles in an animated Oz world. There are minor differences in the two systems that will not be discussed here.

ceeded or failed. When this occurs, Hap allows Em to update its state. This coordination is important as it allows Em's state to change because of Hap's execution, which can in turn affect Hap's behavior choice, giving us a dynamic system with immediate emotional feedback.

Further, when Hap is executing it may need to analyze text it receives from another agent. Hap will send the text to Gump (not shown) for parsing into a simple case-frame language that Hap is designed to process. Then, if Hap decides it wants to respond, it sends a similar case-frame representation to Glinda (not shown) which generates appropriate output based on both the case-frame and pragmatic considerations such as the agent's emotional state [8].

The integration of these modules is an interesting problem in itself [4], but not one that we will discuss here except for a brief description of the Em-Hap integration in the behavioral features section below.

3 Models of Emotion

Our emotion system is modeled on a scaled down version of the cognition-based emotion model of Ortony, Collins, and Clore (OCC) [12], though we were also influenced by the ideas of Dyer [6]. Here we summarize the scaled down version of the OCC model we are using and then describe the Em model and the emotional system of a specific Tok agent that we have implemented, Lyotard the cat. We will also explain where we have diverged from the OCC model and why we found such changes useful or necessary.

3.1 The OCC Emotion Model

Instead of attempting to describe every possible emotion, the OCC model works at the level of emotional clusters, called emotion types, where the emotions within each cluster share similar causes. For example, the distress type describes all emotions caused by displeasing events. Distress includes individual emotions, such as sad, distraught, and lovesick that differ in less significant dimensions such as intensity and the reason the event was found to be unpleasant. We will discuss emotion types here as if they were actually emotions, so keep in mind that we are really discussing clusters of similar emotions.

The OCC model proposes that emotions are the results of three types of subjective appraisals: (1) the appraisal of the pleasingness of events with respect to the agent's goals, (2) the appraisal of the approval of the actions of the agent or another agent with respect to a set of standards for behavior, and (3) the appraisal of the liking of objects with respect to the attitudes of the agent. The model also proposes another set of emotions that are caused by combinations of other emotions. Altogether, these

give rise to a number of emotions, of which we will examine: joy, distress, hope, fear, pride, shame, admiration, reproach, anger, gratitude, gratification, and remorse.

Events are things that happen, including the actions of agents. They are judged to be either pleasing or displeasing according to the agent's *goals*. Goals represent anything that the agent wants, so they may be actively pursued (e.g. "I want to eat") or they may be more passive (e.g. "I want the Mets to win the World Series"). So, the event of eating dinner when there is a goal to eat will be judged as being pleasing. Events give rise to joy and distress emotions with an intensity based on a number of factors, including how (dis)pleasing the event was found to be. The prospect of future events gives rise to fear and hope. Like joy and distress the intensity of these emotions is determined by the pleasingness or displeasingness of the potential future event.

Agents' actions can be approved of or disapproved of according to a set of *standards*. Standards represent both moral beliefs of right and wrong (e.g. "Thou shall not hit people on the head") as well as personal beliefs about level of performance (e.g. "I should be able to bowl 150"). If the agent doing the judging approves or disapproves of one of their own actions, then they experience pride or shame. Similar actions by another agent give rise to admiration or reproach. The intensities of these emotions are based primarily on the level of (dis)approval of the action.

Objects (including agents) can be liked or disliked according to an agent's *attitudes*. Attitudes represent personal tastes and preferences (e.g. "I dislike modern art"). When an agent attends to any object (e.g. thinks about it, imagines it, notices it) this can give rise to an emotional response. Attending to a liked object gives rise to love and attending to a disliked object gives rise to hate. A distinction to keep clear is between attitudes and the emotions that they generate, especially as they are often given identical names (e.g. we feel *love* [the emotion] toward people that we *love* [the attitude]). The intensities of these emotions are based primarily on the level of (dis)like for the object.

Finally, emotions may be combinations of other emotions. For example, if Alice hits Bob on the head, then Bob will likely disapprove of the action (reproach) as well as find it displeasing (distress). This combination will lead Bob to feel anger toward Alice in addition to the other emotions. In addition, the intensity of the anger is based on the intensities of the two component emotions, so it will tend to overwhelm the other two emotions in effect. Similarly, gratitude is a combination of joy and admiration, gratification is a combination of joy and pride, and remorse is a combination of distress and shame.

3.2 The Em Emotion Model

The Em system is based on the OCC model, but differs in a number of ways that seemed useful or necessary to implement the model and use it to develop a Tok

EMOTION	CAUSE
Joy	Goal success (*)
Distress	Goal failure (*)
Hope	Prospect of goal success (*)
Fear	Prospect of goal failure (*)
Pride	Action of self approved according to standards
Shame	Action of self disapproved according to standards
Admiration	Action of other approved according to standards
Reproach	Action of other disapproved according to standards
Love	Attention to liked object
Hate	Attention to disliked object
Gratification	Action of self causes joy and pride
Gratitude	Action of other causes joy and admiration
Remorse	Action of self causes distress and shame
Anger	Action of other causes distress and reproach

(*) Denotes difference from OCC model

Table 1: Current Emotions in Em and their Causes

agent, namely Lyotard the cat. (For a more detailed account of Lyotard's structure and behavior, see [3].) Table 1 shows the emotions that Em models, what Em takes to be the causes for those emotions, and which of the causes differ from those described in the OCC model. In this section we will describe how Em actually implements these emotions and how some of our implementation decisions have led us away from the OCC model.

The way Hap represents goals makes it trivial for Em to determine when a goal has succeeded or failed, but determining the cause of the success/failure is not necessarily even possible. Also, it is possible for an agent to have a goal succeed or fail due to an event that is never sensed or inferred by the agent. In these cases, we still want the agent to react in an emotional manner. Because of these reasons, Em, unlike the OCC model, does not judge events directly. Instead, Em judges the outcomes of the agent's goals. This mechanism is actually closer in nature to conflict theories of emotion as described by Mandler in [10].

There are two sets of goals in Tok: Hap keeps a goal tree of the current active goals and Em keeps a list of the permanent passive goals. Both types of goals succeed and fail in various ways, and both types of goals have an importance level associated with them. Whenever a goal succeeds or fails, Em creates a joy or distress emotion with intensity equal to the importance of the goal. All of the passive goals are checked during Em's initial processing in each sense-think-act cycle. Although the internal state may change subsequent to these checks, the external world does not alter during the sense-think-act cycle. This means that some internal goal failures may not immediately be noticed, but in general this works fairly well. The active

goals are Hap's responsibility, so when Hap discovers a goal success/failure it alerts Em and gives Em access to that goal for determining the goal's importance and the cause of its success or failure.

Because Em cannot directly assess the cause of a goal's success/failure, we have created a mechanism to help in this task. Em associates with each goal a function that examines the agent's sensory memory for an event that may have led to the goal's success or failure. For example, when Lyotard's being-pet goal succeeds, Em searches for one of several possible actions (e.g. pet, rub) of which Lyotard was the object and credits that event with causing the joy. We also do not model the agent's perceived likelihood of goal success or failure, so emotions based on changes of goal likelihood are not yet modeled.

Tok does not yet have a system for imagining potential future events. We plan to add an internal world simulator to Tok, which will help solve this problem as well as help with the goal likelihood problem. For now generating hope and fear is somewhat complicated. In Lyotard we associate with each goal two tests that examine the environment for various cues that would indicate potential for goal success or failure. These tests are also able to return some cause for the fear or hope that Em uses to define the emotion more accurately. In Lyotard, the get-human-to-open-door goal has a hope test that looks around for nearby humans that might open the door for Lyotard. If a human is noticed, then the test notes the specific human agent responsible for the hope, and this information is then used to determine the intensity of the hope (a liked human creates more hope than a neutral or disliked human). Em checks all active and passive goals for hope/fear during its initial processing. Also, as Hap creates new goals it alerts Em and Em checks the new goals for potential hope and fear.

Em judges its (dis)approval of actions by evaluating every act that it senses according to its standards. Lyotard is a house cat and our first complete Tok agent, so we keep things simple by only modeling two standards, help-my-goals-to-succeed and do-not-cause-my-goals-to-fail. For this reason, we are able to speed this process up by only checking the standards when goals succeed or fail, but this is not possible for standards in general. Lyotard's standards include a test for each goal to which it can possibly assign credit/blame. For example, the help-my-goals-succeed standard, when invoked because of the success of the eat goal, will use a test to look around for a nearby human that may have just filled the catbowl with food. Em will generate admiration if a nearby human is found.

Objects are (dis)liked according to the agent's attitudes. Attitudes contain a function for determining if a certain object matches the given attitude and a level of like/dislike for such an object. Em can run any set of sensed, physical objects specified by the agent builder through these attitude functions. For example, Lyotard checks all objects that he senses against his attitudes. We say that this is the list of objects

that is *attended to*. This is a rather crude attentional mechanism for determining what objects an agent is currently thinking about or noticing. This mechanism is unrealistic for two reasons: agents should not be able to attend to all sensed objects and agents should have emotional reactions to objects that they are not currently sensing (e.g. remembered or imagined objects).

Once the list of objects attended to is determined, Em matches them against the attitudes. Whenever a match is found an appropriate love or hate emotion is generated. The only object that Lyotard responds to emotionally is a human agent. Initially, Lyotard dislikes the human. Lyotard also contains rules for updating this attitude based on the human doing things to cause Lyotard to be either angry (leading to greater dislike) or grateful (leading to greater like). Em updates the attitudes during its initial processing stage of the sense-think-act cycle as it has access to all of the sensed objects at that point. This will have to change once we allow emotional reactions to non-sensed objects.

The combination emotions are created simply by searching through all of the current emotions looking for appropriate matches. For example, if distress and shame emotions are found that were both caused by unintentionally hurting someone's feelings, then a remorse emotion will also be generated with an intensity intensity equal to the sum of the distress and shame emotions. For Lyotard this search is not really necessary. Because Lyotard's standards are tied to his goals, every time a standard-based emotion is generated, a combination emotion will also be generated. For more complicated agents, the search mechanism is necessary.

In addition to generating emotions, Em has a number of other emotion-related responsibilities. For example, Em keeps track of thresholds for every emotion type. This means that agents have built-in level of emotional tolerance so that not everything in the world will spark an emotional reaction, but a series of little events will still be able to generate an emotional response. Em also manages emotion decay. Each time Tok performs a sense-think-act cycle, Em lowers the intensity of all of its individual emotions. The distinction between attitudes and the emotions that they generate is important here – attitudes may change with time but they tend to be much more stable than the emotions that they generate. Both the thresholds and decay rates are specifiable for each emotion type, so Em is customizable in these dimensions.

3.3 Relationship Between Em and OCC

The Em system does not currently implement the entire OCC model; Em doesn't even model the entire subset of the OCC model presented here. One limitation is that Em does not model a number of emotion types, such as happy-for, gloating, resentment, pity, satisfaction, fears-confirmed, relief, and disappointment. This limitation is due

to time constraints and limitations in Tok, such as the lack of models of the internal state of other agents.

The emotions that Em does model are also limited in at least two ways. First, the OCC model contains a list of features that influence the intensity of each of their emotion types. We use some of their suggestions, but do not exhaust their lists. For example, the intensity of joy emotions is supposed to be a function of the unexpectedness of the joy-causing event. We have this limitation for a number of reasons, such as time and the current state of the rest of the Tok architecture, but many of these features are also quite difficult to quantify properly and the OCC model makes no suggestions about how to approach this.

Second, we do not yet differentiate between emotions of the same type. As was explained above, emotion types are conglomerates of similar emotions that differ in a number of ways, and these differences may have relevance for behavior. For example, lovesick is a distress emotion that may give rise to specific behavior that other types of distress do not. The OCC model only hints at some of the possible emotions of each type that might exist. We will have to decide (possibly on an agent by agent basis) which ones are useful to differentiate. This will include writing inference rules for recognizing specific emotions and writing appropriate behaviors that express these emotions.

Though Em does not implement all of the OCC model, in four ways it is (or will be) an extension on it: (1) Em is an implemented version of their theoretical model and includes a number of differences (as described above) that were necessary to create a working system; (2) Em handles social knowledge and how this interacts with emotions (more on this in section 5); (3) Our current plan (though this is not yet implemented) is to extend the concept of attitudes well beyond that of the OCC model – this will be especially important for representing attitudes such as trust and respect that agents can have for other agents; (4) Em is part of a larger agent architecture, so we have had to create a flexible way to specify how emotions influence behavior.

4 Integration: Behavioral Features

Once Em develops an emotional state, that state must influence the agent's behavior. Behavioral features are components of the mental state that modulate the activity of Hap. Em contains a set of agent-dependent rules for setting these features based on the current emotional state of the agent. For example, the behavioral feature "aggressive" might arise when an agent of a certain personality type experiences either of the emotions fear or anger. The aggressive feature as a result of fear might exist in an agent that has "bravado" as a personality trait. It is then up to Hap to generate behavior that is consistent with the current features. For more details on

how this is done, see [3] and [4].

This mechanism gives us the power to represent a large number of personality types. For example, the fear emotion may lead to any of a number of features, like fearful, aggressive, withdrawn, or excited, depending on the personality we are trying to model. These rules may give us the flexibility to represent cultural differences of agents. For example, though we call the emotion that occurs when a goal fails distress, this does not actually carry any direct significance for behavior. We may want most of our agents to act lethargic or withdrawn when distressed, but we can just as easily model a culture where goal failures are seen as learning experiences that cause an agent to act happily or even proudly. We are looking at Carbonell's work on personality traits [5] and Hovy's work on rhetorical features [7] for additional ideas in this area.

5 Social Knowledge

In addition to generating emotions and features for modifying behavior based on those emotions, Em is also responsible for another aspect of broad agents: social knowledge about relationships. One element of social knowledge that Em is not responsible for is representing societal restraints on action. For example, when an agent gets tired we would expect that agent to sleep. However, if the agent happens to be at a party at the time, then just lying down and dozing off is considered unacceptable behavior. This aspect of social knowledge is built into the Hap behaviors of the agent.

The facet of social knowledge that Em models is the interpersonal relationships between agents, how these relationships change over time, and how they interact with emotion. Interpersonal relationships can be complicated to model, so in keeping with our "broad" approach to agents, we have limited our initial efforts to a few simple relationships.

Interpersonal relationships are intertwined with the attitudes that Em uses in determining emotions toward objects. In this case the objects just happen to be other agents. In fact, Lyotard's relationship with the human is simple enough that we are able to use the attitude structures to completely represent it. This is certainly not necessary or desirable in all agents. Lyotard can only have relationships that differ along the like-dislike dimension, though these relationships can have varying strengths (e.g. likes very much, neutral, mild dislike). We plan to expand the attitude system to let us represent relationships that involve elements like trust, fear, and admiration. We will also have to adjust the emotion generation rules of Em to reflect emotions that may arise from attitudes besides like and dislike. In any case, some aspects of relationships, like intimacy/closeness and relative social standing, need structure beyond that offered by attitudes.

Because relationships and attitudes are so closely related, the social relationship

directly influences the emotional system. For example, when Lyotard is around a friend (someone he likes) he will feel love; when around an enemy (someone he dislikes), Lyotard feels hate; when around an agent he has no attitude toward or who he feels neutral toward, Lyotard will not have an emotional experience.

Emotions likewise affect social relationships. Emotions are a common cause of changes in relationships. An agent making Lyotard angry will be liked less by Lyotard, whereas an agent making Lyotard grateful will be liked more. Lyotard initially dislikes the user, but actions like feeding him will eventually win him over.

6 Cultural Independence and Flexibility

Em makes assumptions about the emotional systems to be modeled, such as: (1) the agent has explicit goals with varying levels of importance; (2) the agent has standards; (3) the agent has personal tastes and attitudes; (4) these goals, standards, and attitudes in conjunction with appropriate events and objects in the world lead to emotional experiences; and (5) the agent recognizes interpersonal relationships. Each assumption places restrictions on the architecture that we build and the types of emotional systems that we can model. For example, we would have trouble modeling a system where emotions arose randomly.

Fortunately, these assumptions appear to be consistent with Lutz's view of the cultural universals of human emotional experience [11]. Lutz is a cultural anthropologist who has studied non-western cultures to better understand their emotional experiences. She believes that there are some universals that arise in all human social settings that give rise to emotional experience. Her list of cultural universals, similar to Em's basic assumptions, contain references to goals, standards, and social relationships. Her proposed universals are: (1) the conflicting goals of multiple actors or others violations of cultural standards; (2) ego's own violation of cultural standards, or anticipation of such; (3) danger (physical or psychosocial) to ego or significant others; (4) the loss of significant relationships or the threat of such loss; and (5) the receipt of resources including the tangible (e.g. food) and the intangible (e.g. praise, affection). The fact that Em's assumptions seem consistent with this theory of the culturally universal causes for emotion gives us hope that this model does not overly restrict us in the kinds of emotional frameworks that we can imitate.

Just because our system does not make any assumptions that inherently rule out modeling most human emotional systems does not mean that our system actually has the power to represent these systems. In fact, we do believe that our system is flexible enough to model a large number of different emotional agents. Each agent's goals, standards, and attitudes are completely specifiable. By specifying the agent's goals we can specify that agent's needs, wants, and desires – including those that are culturally based. By specifying the importance of those goals we model the agent's

priorities. By specifying the agent's standards we can specify the agent's moral code, which is obviously culturally dependent. By specifying the agent's attitudes, we can model culturally biased tastes and preferences. Modeling attitudes is a little more subtle because we need to be able to specify what attitudes an agent can have, but we plan to allow this flexibility within the Em framework. Finally, the various emotional thresholds and decay properties of emotions can be specified to produce various levels of emotional restraint.

The way emotion affects behavior is also flexible. The emotions themselves all have names, but these names have nothing to do with how these emotions are processed. So, the distress emotion will occur when an event causes an agent's goal to fail, but the distress emotion does not have to lead to what we would normally consider distressed behavior. The indirection provided by the behavioral features system allows us to customize how emotions influence behavior. It is just as easy to model distress causing enthusiasm or aggression as depression or withdrawal. We will continue to study whether the flexibility of this mechanism is able to achieve all of the emotional expressiveness that our agent builders want.

7 Conclusions and Future Work

The Tok architecture and the Em system have shown promise in modeling broad, emotional agents. In particular, Em is beginning to provide the power to model fairly deep phenomena, including personal and cultural standards, attitudes, interpersonal relationships, emotions, and some level of personality. Em is also interesting because it has been successfully integrated into an action architecture that both influences and is influenced by the workings of Em. Finally, Em shows promise of being flexible enough to model most of the agents that Oz world builders might want. For a more detailed account of the behavior we have achieved with this system, see [3] and [4].

We still have many obstacles between our current work, as it exists in Lyotard the cat and our animated Woggles, and our final goal of a broad agent architecture. Improvements to our emotional and social modeling require extensions in both Tok and Em. In Tok we plan to add models of the internal states of other agents which will allow generation of emotions such as happy-for and pity, and to add an internal physical world simulation which should improve the forward-looking emotions of hope and fear as well as help with estimating likelihood of goal success for generating joy and distress emotions. In Em we also have some work to do. We need to flesh out the attitude system, using for instance ideas from Schank et al. [13], to include attitudes such as trust and respect. We also plan to build a more flexible and expressive system for representing interpersonal relationships, using ideas from Hovy [7] and Dyer [6].

8 Acknowledgments

Work on Tok is being done in conjunction with Bryan Loyall and Mark Kantrowitz of the CMU Oz group. We gratefully acknowledge the support of Fujitsu Laboratories, Ltd. for our research.

References

[1] Joseph Bates. Virtual reality, art, and entertainment. *PRESENCE: Teleoperators and Virtual Environments*, 1(1):133–138, 1992.

[2] Joseph Bates, A. Bryan Loyall, and W. Scott Reilly. Broad agents. In *Proceedings of AAAI Spring Symposium on Integrated Intelligent Architectures*, Stanford, CA, March 1991. Available in *SIGART Bulletin*, Volume 2, Number 4, August 1991, pp. 38-40.

[3] Joseph Bates, A. Bryan Loyall, and W. Scott Reilly. An architecture for action, emotion, and social behavior. Technical Report CMU-CS-92-144, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, May 1992. To appear in Proceedings of the Fourth European Workshop on Modeling Autonomous Agents in a Multi-Agent World, S.Martino al Cimino, Italy.

[4] Joseph Bates, A. Bryan Loyall, and W. Scott Reilly. Integrating reactivity, goals, and emotion in a broad agent. Technical Report CMU-CS-92-142, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, May 1992. To appear in Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society, Bloomington, Indiana.

[5] Jaime Carbonell. Computer models of human personality traits. Technical Report CMU-CS-79-154, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, November 1979.

[6] Michael Dyer. *In-Depth Understanding.* The MIT Press, Cambridge, MA, 1983.

[7] Eduard Hovy. *Generating Natural Language under Pragmatic Constraints.* Lawrence Erlbaum Associates, Hillsdale, NJ, 1988.

[8] Mark Kantrowitz and Joseph Bates. Integrated natural language generation systems. In R. Dale, E. Hovy, D. Rosner, and O. Stock, editors, *Aspects of Automated Natural Language Generation*, volume 587 of *Lecture Notes in Artificial Intelligence*, pages 13–28. Springer-Verlag, 1992. (This is the Proceedings of the Sixth International Workshop on Natural Language Generation, Trento, Italy, April 1992.).

[9] A. Bryan Loyall and Joseph Bates. Hap: A reactive, adaptive architecture for agents. Technical Report CMU-CS-91-147, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, June 1991.

[10] George Mandler. *Mind and Body.* W. W. Norton & Company, New York, NY, 1984.

[11] Keith Oatley. A Review of *Unnatural Emotions: Everyday sentiments on a Micronesian atoll and their challenge to Western theory* by Catherine Lutz (1988). Chicago University Press. *Cognition and Emotion*, 5:65–79, January 1991.

[12] A. Ortony, G. Clore, and A. Collins. *The Cognitive Structure of Emotions.* Cambridge University Press, 1988.

[13] Roger C. Schank, Robert Wilensky, Jaime G. Carbonell Jr., Janet L. Kolodner, and James A. Hendler. Representing attitudes: Some primitive states. Technical Report 128, Department of Computer Science, Yale University, New Haven, CT, January 1978.