Research problems in the use of a shallow Artificial Intelligence model of personality and emotion*

Clark Elliott
Institute for Applied Artificial Intelligence
DePaul University
243 South Wabash Avenue
Chicago, IL 60604
and
School of Education and Social Policy
Northwestern University
Evanston, IL
email: elliott@ils.nwu.edu

Abstract

This paper presents an overview of some open research problems in the representation of emotion on computers. The issues discussed arise in the context of a broad, albeit shallow, emotion reasoning platform based originally on the ideas of Ortony, Clore, and Collins(Ortony, Clore, & Collins 1988). In addressing these problems we hope to (1) correct and expand our content theory of emotion, and pseudo personality, which underlies all aspects of the research; (2) answer feasibility questions regarding a usable representation of the emotion domain in the computer, and (3) build agents capable of emotional interaction with users. A brief description of a semantics-based AI program, the *Affective Reasoner*, and its recent multi-media extensions is given. Issues pertaining to affective user modeling, an expert system on emotion eliciting situations, the building of a sympathetic computer, models of relationship, personality in games, and the motivation behind the study of emotion on computers are discussed. References to the current literature and recent workshops are made.

Introduction

The emotion reasoning platform discussed in this paper has developed over the course of several years and currently includes, besides the underlying emotion engine (described briefly below), a speech recognition package which is able to discriminate at least some broad categories of emotion content; a music indexing and playback mechanism allowing virtually instant access to hundreds of hours of midi format music used to aid in the expression of emotion; a schematic representation of approximately 70 emotion faces; and a text-to-speech module for expressing dynamically constructed text, including a minimal amount of emotion inflection (through runtime control of speed, pitch and volume).

*Preparation of this article was supported in part by Andersen Consulting through Northwestern University's Institute for the Learning Sciences. Citation appears in bib.

In the spirit of one of this year's conference themes, we view this project as being in the "platform and concept hacking" stage, wherein we explore what plausible expectations we may make with respect to a shallow model of emotion and personality representation on the computer. Findings are strictly preliminary, and yet enough work has been done to raise what we feel to be some interesting questions.

What we refer to as "emotions" in this paper arise naturally in many human social situations as a byproduct of goal-driven and principled (or unprincipled) behavior, simple preferences, and relationships with other agents. This applies to many situations that one would not ordinary refer to as *emotional*: a social *agent* becoming *annoyed* with someone who is wasting her time (a mild case of that person violating the agent's principle of social efficiency thus blocking of one of the agent's goals through the reduction of a valued resource, time), enjoying a piece of music because it is appealing (liking it, through a simple, unjustifiable, preference), and so forth. We limit our consideration of emotion states, and intensities, to states and intensities similar to what Frijda et al. (Frijda *et al.* 1992) were describing when they referred to the *overall felt intensity* of an emotion as comprising "whatever would go into the generation of a response to a global question such as this: 'How intense was your emotional reaction to situation S?"' Physical manifestations, neural processes, and much about duration are not included in the model.

Lastly, we suggest that the representation of human emotion and personality in a social context, using AI techniques, is long overdue as a major area of study (c.f. (Norman 1980)). It is our belief that for each of the issues raised below, enough background work has been done that partial solutions are within the grasp of the AI community.

Background

In our current research, embodied in a large AI program called the Affective Reasoner, we simulate simple worlds populated with with agents capable of responding "emotionally" as a function of their concerns. Agents are given unique pseudo-personalities modeled as both a set of *appraisal frames* representing their individual goals, principles, preferences, and moods, and as a set of *channels* for the expression of emotions. Combinations of appraisal frames are used to create agents' interpretations of situations that unfold in the simulation. These interpretations, in turn, can be characterized by the simulator in terms of the eliciting conditions for emotions. As a result, in some cases agents "have emotions," which then may be expressed in ways that are observable by other agents, and as new simulation events which might perturb future situations (Elliott 1992). Additionally, agents use a case-based heuristic classification system (based on (Bareiss 1989)) to reason about the emotions other agents are presumed to be having, and to form representations of those other agents' personalities that will help them to predict and explain future emotion episodes involving the observed agent (Elliott & Ortony 1992; Elliott 1992).

Ortony, et al. (Ortony, Clore, & Collins 1988) discuss twenty-two emotion types based on valenced reactions to situations being construed as goal-relevant events, acts of accountable agents, or attractive or unattractive objects (including agents interpreted as objects). This theory has been extended to include the two additional emotion types of *love* and *hate* (Elliott 1992). See figure 1.

Additionally, using the work of Ortony, et al. (Ortony, Clore, & Collins 1988) as a guide, we analyzed a set of descriptions of emotion eliciting situations and created a modified set of *emotion intensity variables* to explain the causes of varying emotion intensity, within a coarse-grained simulation paradigm (Elliott & Siegle 1993). We reduced the resulting set of variables to a computable formalism, and represented sample situations in the Affective Reasoner. We then isolated three areas of the simulation where variables in either the short-term *state* of an agent, the long-term *disposition* of an agent, or the *emotion-eliciting situation* itself, helped to determine the intensity of the agent's subsequent affective state. For each area there is an associated group of variables. The first group, *simulation-event variables*, comprises variables whose values change independently of situation interpretation mechanisms. The second group, *stable disposition variables*, consists of variables that are involved in an agent's interpretation of situations, tend to be constant, and help to determine an agent's personality and role in the simulation. We felt that, for the purposes of implementation, the distinction between these two groups was underspecified in the work of Ortony, et al. (Ortony, Clore, & Collins 1988). The last group, *mood-relevant variables*, contains those variables that contribute to an agent's mood state. In all there are approximately twenty such variables, although not all apply to each emotion category.

For example, the variable *blameworthiness-praiseworthiness* might roughly be described as the degree to which an observing agent interprets an observed agent as having upheld or violated one of the observing agent's principles, in some situation. It is derived from a set of simulation values, which might include values for the amount of effort expected in a given situation, the accountability of an agent as determined by role, and so forth. It has no default value, being determined entirely by one agent's construal of the simulated act of an another agent.

Our work has focused primarily on the detailed working out of a computational method for representing the antecedents and expression of human emotion in diverse human social situations. This has included the analysis, within the constraints of the underlying emotion theory, of many hundreds of informally described social situations which have given rise to emotions. To date the computational mechanism has included, among other components, (a) the construction of hundreds of *appraisal frames*, which include slots for reasoning about intenstiy and mood, in domains as diverse as, for example, financial accounting, stories, playing poker, and sales, (b) pseudo personality types made up of these appraisal frames, (c) the use of these pseudo personalities for *construing* situations with respect to the concerns of simulated agents, thus giving rise to "emotion generation" in simulation runs, (d) the generation of simple emotion instances based on the twenty-four emotion categories, (e) the generation of actions through approximately 450 channels (about twenty for each emotion category) consistent with the simulated emotions they are intended to express (each of which may, in turn, contain multiple manifestation *instances*), (f) abductive reasoning about the emotions expressed by other agents in the system, (g) the internal representation of the presumed pseudo personalities of observed agents by observing agents, (h) simple "explanations" of the emotions with respect to their emotional antecedents within the simulation, (i) simple models of relationship between the agents, allowing for "emotions" based on the concerns of others, and (j) the inclusion of the user as one of the agents in the system about which reasoning may take place.

Most recently we have been working on opening up communication channels with the user of the system through the addition of modules for speech recognition, inflected speech generation, indexed music-on-demand (using a midi interface and a 400 voice Proteus MS-PLUS synthesizer), and facial expression for the agents.

The broad long-range goals we would like to see pursued include number of applications we envision as made possible by the representational capabilities of a system such as this. Among these are the building of a computer that has some capability of categorizing, and responding to, a user's affective state; the building of systems that allow users to interactively explore the emotional content of a broad range of simulated social situations (e.g., for tutoring, for socially unskilled psychotherapy patients, and for military stress applications); testing the use of "emotionally aware" automated software agents as a way of enhancing the user's engagement in educational, and other software; using emotionally aware agents to communicate priorities naturally to the user, such as with an automated assistant for meetings, or the communication of technical concerns to less technical users through the focus of attention (e.g., operating systems, financial analysis), the use of computer-based emotion expression as an authoring tool (e.g., as online feedback for students), the construction of games that include a complex, yet cohesive, emotion and personality component; the use of the underlying emotion theory to analyze, and manipulate, the automated telling of stories; and platforms for testing the link between music and emotion expression.

Research Questions

Affective user modeling.

One hard, and divisive, problem facing the AI community is that of building user models. Rather than more traditional models which focus on the mental processes of the user in problem solving situations (Van Lehn 1988), we propose an alternative wherein only certain components of the *affective* state of the user is modeled. This is a much smaller problem, but one which should provide useful leverage. It might be considered akin to the feedback a responsive speaker might make use of when "playing" to her audience. We do *not* propose this as a full model of a user's emotional states which would then also require that all of the hard cognitive modeling problems be solved as well.

To implement simple affective user modeling, several components are required: (1) A structure which allows us to capture an agent's (in this case, the user's) outlook on situations that arise. This structure must include some concept of the role, personality, and current state of the user (within the context of shallow emotion reasoning), which together comprise the basis for the user's idiosyncratic way of construing the world. (2) A lexicon through which the user expresses his or her emotions to the computer. (3) A comprehensive set of emotion categories which allow for the mapping of emotion eliciting situations to the emotion expression lexicon, and vice versa. In our current work, as discussed above, we have implemented a broad, albeit shallow, representation of the first component, and a comprehensive, descriptive representation of the third.

The weakest link in such a system is in the lexicon. How does a computer, which has no understanding of faces[1], and which presumably has no mechanism for generating plausible explanations which might allow it to determine which emotions are likely to have arisen, know what emotion a user is expressing? In addressing this question we consider several leverage points which show promise in allowing us to work around this problem, at least to some degree. First, and most importantly, it might well prove to be true that users are *motivated* to express their emotions to the computer, provided that there is at least the *illusion* that the computer understands how they are feeling. Should this be so, then some latitude is afforded us in requiring that the user, who is adaptable in communication, conform to the protocol of the computer, which is not. Second,

[1] Although reportedly the work of Eckman might begin to address this issue.

Figure 1: Emotion types (Table based on [O'Rorke and Ortony, 1992] and [Elliott, 1992])

Group	Specification	**Name** and Emotion Type
Well-Being	appraisal of a situation as an *event*	**joy**: pleased about an *event* **distress**: displeased about an *event*
Fortunes of Others	presumed value of a a situation as an *event* affecting another	**happy-for**: pleased about an *event* desirable for another **gloating**: pleased about an *event* undesirable for another **resentment**: displeased about an *event* desirable for another **sorry-for**: displeased about an *event* undesirable for another
Prospect based	appraisal of a situation as a prospective *event*	**hope**: pleased about a prospective desirable *event* **fear**: displeased about a prospective undesirable *event*
Confirmation	appraisal of a situation as confirming or disconfirming an expectation	**satisfaction**: pleased about a confirmed desirable *event* **relief**: pleased about a disconfirmed undesirable *event* **fears-confirmed**: displeased about a confirmed undesirable *event* **disappointment**: displeased about a disconfirmed desirable *event*
Attribution	appraisal of a situation as an accountable *act* of some agent	**pride**: approving of one's own *act* **admiration**: approving of another's *act* **shame**: disapproving of one's own *act* **reproach**: disapproving of another's *act*
Attraction	appraisal of a situation as containing an attractive or unattractive *object*	**liking**: finding an *object* appealing **disliking**: finding an *object* unappealing
Well-being / Attribution	compound emotions	**gratitude**: admiration + joy **anger**: reproach + distress **gratification**: pride + joy **remorse**: shame + distress
Attraction / Attribution	compound emotion extensions	**love**: admiration + liking **hate**: reproach + disliking

the comprehensive emotion model allows us to develop a large lexicon categorized by both emotion category and intensity. Third, speech recognition packages are advanced enough to capture some of what is of interest to us with respect to the lexicon.

For example, using the work of Ortony et al. as a guide (Ortony, Clore, & Foss 1987), we built expressions containing emotion words, intensity modifiers, and pronoun references to different roles (e.g., *I am a bit sad because he..., I am rather sick at heart about her..., I was pretty embarrassed after my...*) (Elliott & Carlino 1994). We built phrases containing 198 emotion words (e.g. ...,bothered, brokenhearted, calm, carefree, chagrined, charmed, cheered, cheerful, cheerless,...). In preliminary runs we were able to detect 188 of the emotion words correctly on the first try, in context, with 10 false positives. Misses tended to be cases such as confusing "anguish" with "anguished," and "displeased" with "at ease." There were 10 other instances of difficulty with other parts of the phrases, such as confusing "my" with "I." Most of these would have been caught by a system with even rudimentary knowledge of English grammar.

Additionally, in other preliminary runs of our speech recognition package the computer was able to recognize the seven emotion categories, anger, hatred, sadness, love, joy, fear, and neutral, which we did our best to communicate to it, when speaking the sentence, "Hello Sam, I want to talk to you." In this small exercise we broke the sentence up into three parts, identifying each part as a "word" to the speech recognition system. We then trained each phrase for the seven different inflections. With practice we were able to get close to 100% recognition of the intended emotional state.[2] To achieve this we had to be slightly theatrical, but not overly so, and there was a flavor of speaking with someone who was hard of hearing, but again, not overly so.[3]

Once the lexicon is established, and minimal natural language parsing is in place through the use of an ATN or other simple system, tokens can either be interpreted directly as situations in themselves, or as a set of features indexing into a case-base to retrieve similar cases indicating a particular emotion category. To illustrate, on the one hand, from user input of "I am satisfied with the results" we might, for example, yield the *situation*: "the user is *satisfied* now in response to the comparison of her answer with the one just provided by the computer"). On the other hand, given user input of "I am happy now" (spoken with a hateful inflection), we might yield the set of *features*: *user expresses happiness*, and *user's inflection expresses hate*, which in turn retrieves the cases of *hatred* masked by a contrasting verbal expression.

Assuming that such a system can be built, it raises the possibility of testing its numerous applications in diverse domains. We touch on these in the following sections, in concert with other issues.

A sympathetic computer

Sidestepping the issue of whether a a user's *engagement* with a computer system will be increased by dealing with an emotionally believable agent, we might first ask what is required to achieve believability. How much does it matter if the computer is transparent in its capabilities and motivations? What level of sophistication is required? How intelligent does the system have to be to achieve a minimal level of believability?

Consider the following: A computer system detects, or is told by the user (as discussed above), that the user is fearful (anxious, worried, scared – each of these would be categorized as different intensities of the emotion category *fear*). The system is able to respond by asking what might be considered reasonable questions: "What is the name of the *event* over which you are fearful?," "How far is this event in the future?" (According to the underlying theory, as embodied in the Affective Reasoner, *fear* is a prospect-based emotion resulting from the likely blocking of a future *event*.) "What is the likelihood of this coming about?" (*Likelihood* is an intensity variable). It is then able to comment on relationships between the various *intensity variables* and then to recall cases from a library, where either the automated agent itself ("I was afraid of getting turned off!" (Frijda & Swagerman 1987)), or some other agent ("I was previously told the following story by another user...") was *fearful* with similar intensities. Lastly, after verifying that the retrieved case involves an emotion similar to that of the user, the computer agent responds by saying, "I am only a stupid computer program. Nonetheless, in my own simple way, the following is true: I consider you my friend (see *models of relationship* below). I *pity* you for your *fearful* state. I *hope* to become aware in the future that you are *relieved* about the outcome of situation *<input-situation>*, and that your *preservation goals* with respect to this are not blocked after all."

Such a system is within the range of our current technology, and representational capabilities (and see the reference to Scherer's work below). At what level of sophistication are users willing to accept that, in its own (remarkably) simple way, the computer does nonetheless feel pity for them? Given the tendency of users to anthropomorphize even the simplest video game characters, this would seem to be an important question to answer.

[2] And this, even though the goal of the developers of the package was to *eliminate* inflection and simply recognize the words.

[3] Interestingly, when not *imagining* a situation consistent with the emotion we were attempting to express by voice, we were consistently less accurate. In addition, it appeared that the computer was biased somewhat in its errors toward the category the user was imagining, for certain pairs of categories, which might suggest an interesting study.

Expert System on Emotion

Scherer describes an expert system on emotion motivated by the need to "use computer modelling and experimentation as a powerful tool to further theoretical development and collect pertinent data on the emotion-antecedent appraisal process" (Scherer 1993). His system captures user input feature vectors representing an emotional situation and shows the relative distance from various predicted emotion concepts (categories in our terminology). He uses this as a way of validating the underlying representation of the appraisal-to-emotion process in his system.

Can we repeat or extend Scherer's results using the differing emotion categories, and appraisal structures? In particular, in the Scherer work the user is asked to identify the intensity of the emotional experience being described. Using our current theory, it would be suitable to draw on the antecedents of emotion intensity, embodied in the twenty or so emotion intensity variables given in (Elliott & Siegle 1993). Such work would also dovetail with the "sympathetic computer" mentioned above.

Models of relationship

In (Elliott & Ortony 1992) the authors discuss the ability of emotionally cognizant agents to model the concerns of one another, and the user. One aspect of these models is that they allow for the modeling of simple *relationships* between the agents, including the user. To wit, we might define a simple model of friendship as requiring that an agent be *happy for* another agent when good fortune strikes that other agent, and feel *pity* when bad fortune strikes. To do this the first agent must have a model of how a situation is presumed to be *construed by that other agent*. For example, to feel sorry for a friend when her basketball team has lost, it is important to know which team the friend is rooting for. This can only be done if each agent maintains some internal model of each other agent's presumed concerns.

In our current work we have built simple models of friendship and animosity, and to some degree, *identification* with another agent (i.e., where an agent takes on another's goals as its own.) When this is extended to include the user, interesting possibilities arise, such as those discussed in the next section. How sophisticated a set of relationships can we build based on the emotional interaction of the (human and) automated agents?

Games

Consider the game of poker, and how an emotion model can enhance the believability, and possibly interest, of a computerized version of the game. Using a preliminary model, we have discussed (and built a simple versions of) the following system (Marquis & Elliott 1994):

The computer simulates one, or more, agents who play five card draw poker against the user. Each agent knows the basic rules of poker, including betting. Agents have their own *goals*, *principles*, and *preferences* so that they respond differently to different situations that arise. Bluffing, required if one is to truly simulate poker, leads to fear for some agents, hope for others. Some agents show their emotions (through facial expression, music selection, and verbal communication of physical manifestations), others suppress them (as in real life). An agent who has been bluffing and wins, might *gloat*, or if it loses it might feel *remorse*. An increase in the amount of the "pot" increases the *importance* of winning and thus the intensity of many of the emotions. Likewise *surprise* (generally hard to represent) can be derived from losing even though one has a good hand, or vice versa, thus intensifying the respective emotion. Agents might feel reproach towards, for example, a player (such as the user) who is too cautious, or takes to long to decide what to do (based on a violation of the principled customary way of playing). Other emotion-eliciting situations include having the computer attempt to cheat without being caught by the user, and so forth.

One claim that might be made is that to approach a realistic computer model of games such as poker, some model of personality and emotion, minimally a broad and shallow model as described here, is essential. Whether or not this makes the computer a better playing companion is an open question, but again, one that seems worth answering.

Are schematic facial models sufficient?

It is clear that sophisticated dynamic three-dimensional models of faces, such as that shown by Pelechaud at IJCAI93, have the power to delight audiences and convey expressive information(Pelachaud, Viaud, & Yahia 1993). Nonetheless, it may also be argued that much emotional content can be delivered in schematic format as well. Cartoon faces (such as those in *Calvin and Hobbs*, for example, convey much about the current appraisals of the cartoon character, and schematic faces have been used in clinical settings to help children identify their emotions. Are such faces, which are much less computationally expensive to manipulate, able to convey emotions in a consistent way?

In our own work we use a set of approximately seventy schematic faces, covering up to three intensities in each of the twenty-four emotion categories (Elliott, Yang, & Nerheim-Wolfe 1993). Sixty of these have been included in a morphing module so that faces gradually break into a smile, decay from *rage* back to a default state, and so forth. The module runs in real time, allowing run-time control over face size, rate of morph, and rudimentary mouth movement (for when the agent is speaking). The system thus allows for over 3000 different morphs, a range not possible with 3D representation. The morphs run on a 66 Mhz IBM PC (with sound and speech cards) concurrently with midi playback, text-to-speech, speech recognition, and the background emotion simulation.

Assuming that either representation can be effective, the question still arises about the effectiveness of the emotion representation on which the dynamic face depends. Our approach is that, at present, low-level personality and emotion representations are too complex to simulate complex social interaction, and that content theories of personality and emotion embedded in individual domains are too simplistic. Hence the middle ground, using an *architectural* approach (e.g., at the level of schematic face representation) consistent across all aspects of the system.

Other areas of applicability

There are a number of other areas where a broad, but shallow computer representation of emotion and personality might be useful. These include testing the ability of the computer to represent the underlying principles in many areas of human endeavor. To this end we are formalizing representations in many different domains, including, in addition to the above mentioned domains, business applications (including selling, and financial statement analysis – where personal interpretation and preference of the analyst is actually somewhat common), and storytelling (using emotion themes as organizing principles similar to Lehnert's plot units (Lehnert 1981), (cf. (Reeves 1991))).

Other appropriate uses include military applications where training to deal with stressful situations *must* include a model of personality and emotion and social interaction, as well as role; understanding of political personalities (do world leaders act according to logic, or according to a personal codes based on different principles, and goals?)(Bannerjee 1991); family politics; and social simulations (Kass *et al.* 1992).

Why should this work be pursued?

Emotion representation on the computer can be seen to fall into at least four categories of pursuit: (1) the testing of design issues raised by theories of emotion (Colby 1981; Toda 1982; Frijda & Swagerman 1987; Pfeifer & Nicholas 1985; Sloman 1987), (2) the placing of neuro-motor control in an environmental or social context (Gray 1993; Rolls 1993), (3) the use of a "folk" representation of emotion to control the behavior of automated agents in social situations; and to predict, or attempt to understand, the behavior of other agents in such situations (Elliott 1993; Reilly 1993; Bates, A. Bryan Loyall, & Reilly 1992) and, (4) the use of emotions for process control (Birnbaum & Collins 1984).

At the recent *Workshop on Architectures Underlying Motivation and Emotion* [4] it became clear that the issue of which approach is more promising is far from settled. Even commonly held beliefs about emotions, such as their use as some sort of reactive-planning mechanism were questioned (e.g., Jeffrey Gray posed the question, if this were true, why would all the manifestations of fear arise many seconds *after* slamming on one's brakes to avoid an auto accident?). Perhaps what we consider to be emotions arise only as a byproduct of more essential mechanisms? Nonetheless, it seems that emotions are ubiquitous in human society and an integral part of the social fabric thereof. Until shown that it is wrong, we will continue to make the following argument: even if we were to completely understand, and be able to recreate, the neural-procedural architecture of the the part of the brain where emotions reside, we still would need to have an understanding of the *software* that was to run on the machine. (To wit: consider that *deservingness* is a somewhat universal concept, and its effects on our emotions are also common fare (e.g., our pity for people suffering the effects of war is increased when those people are *innocent children.*). How is this represented in the biological hardware?) Additionally, unless we can "download" human reasoning into a machine, it is necessary to specify the rules underlying the tasks and responses we wish to make our agents capable of, within the various domains. In the short term at least, this will require a scientifically based, analytical understanding of how personality and emotion affect the interaction and motivation of human agents in social situations. Note that we are not so pure in our arguments: such bottom up approaches will yield tremendous insight into design, and will most certainly constrain the top down approach. Likewise, our top-down approach will yield insights into determining the more salient emotion issues in developing a smooth social interaction with computers. We see the use of intelligent computer systems, especially those that interact with human agents, as the most promising path of study.

[4] Birmingham, England, August, 1993

Acknowledgement

The author wishes to gratefully acknowledge the assistance of Stuart Marquis, Greg Siegle, Yee-Yee Yang, and Eric Carlino in the development of this work.

References

Bannerjee, S. 1991. Reproduction of opppositions in historical structures. In *Working Notes for the AAAI Fall Symposium on Knowledge and Action at the Social and Organizational Levels.* AAAI.

Bareiss, R. 1989. *Exemplar-Based Knowledge Acquisition, A Unified Approach to Concept Representation, Classification, and Learning.* Academic Press, Inc.

Bates, J.; A. Bryan Loyall; and Reilly, W. S. 1992. Integrating reactivity, goals, and emotion in a broad agent. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society.* Bloomington, IN: Cognitive Science Society.

Birnbaum, L., and Collins, G. 1984. Opportunistic planning and freudian slips. In *Proceedings of the Sixth Annual Conference of the Cognitive Science Society.* Boulder, CO: Cognitive Science Society.

Colby, K. M. 1981. Modeling a paranoid mind. *The Behavioral and Brain Sciences* 4(4):515–560.

Elliott, C., and Carlino, E. 1994. *Detecting user emotion in a speech-driven interface*. Work in progress.

Elliott, C., and Ortony, A. 1992. Point of view: Reasoning about the concerns of others. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*, 809–814. Bloomington, IN: Cognitive Science Society.

Elliott, C., and Siegle, G. 1993. Variables influencing the intensity of simulated affective states. In *AAAI technical report for the Spring Symposium on Reasoning about Mental States: Formal Theories and Applications*, 58–67. American Association for Artificial Intelligence. Stanford University, March 23-25, Palo Alto, CA.

Elliott, C.; Yang, Y.-Y.; and Nerheim-Wolfe, R. 1993. Using faces to express simulated emotions. unpublished manuscript.

Elliott, C. 1992. *The Affective Reasoner: A Process Model of Emotions in a Multi-agent System.* Ph.D. Dissertation, Northwestern University. The Institute for the Learning Sciences, Technical Report No. 32.

Elliott, C. 1993. Using the affective reasoner to support social simulations. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, 194–200. Chambery, France: Morgan Kaufmann.

Elliott, C. 1994. Research problems in the use of a shallow artificial intelligence model of personality and emotion. In *Proceedings of the Twelfth National Conference on Artificial Intelligence*, 9–15. Seattle, WA: AAAI. THIS PAPER.

Frijda, N., and Swagerman, J. 1987. Can computers feel? theory and design of an emotional system. *Cognition & Emotion* 1(3):235–257.

Frijda, N. H.; Ortony, A.; Sonnemans, J.; and Clore, G. L. 1992. The complexity of intensity: Issues concerning the structure of emotion intensity. In Clark, M., ed., *Emotion: Review of Personality and Social Psychology*, volume 13. Newbury Park, CA: Sage.

Gray, J. 1993. A general model of the limbic system and basal ganglia; applications to anxiety and schizophrenia. In *Workshop on Architectures underlying motivation and emotion.* The University of Birmingham School of Computer Science and Centre for Research in Cognitive Science. August 11-12, Birmingham, England.

Kass, A.; Burke, R.; Blevis, E.; and Williamson, M. 1992. The GuSS project: Integrating instruction and practice through guided social simulation. Technical Report 34, The Institute for the Learning Sciences, Northwestern University.

Lehnert, W. 1981. Plot units and narrative summarization. *Cognitive Science* 5:293–331.

Marquis, S., and Elliott, C. 1994. Emotionally responsive poker playing agents. In *Notes for the Twelfth National Conference on Artificial Intelligence (AAAI-94) Workshop on Artificial Intelligence, Artificial Life, and Entertainment*, 11–15. American Association for Artificial Intelligence.

Norman, D. A. 1980. Twelve issues for cognitive science. *Cognitive Science* 4(1):1–32. Keynote speech for First Cognitive Science Conference.

O'Rorke, P., and Ortony, A. 1994. Explaining emotions. *Cognitive Science* 18:283–323.

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

Ortony, A.; Clore, G.; and Foss, M. 1987. The referential structure of the affective lexicon. *Cognitive Science* 11:341–364.

Pelachaud, C.; Viaud, M.-L.; and Yahia, H. 1993. Rule-structured facial animation system. In *Proceedings of the Thirteenth Annual Joint Conference on Artificial Intelligence.* Chambery, France: Morgan Kaufmann.

Pfeifer, R., and Nicholas, D. W. 1985. Toward computational models of emotion. In Steels, L., and Campbell, J. A., eds., *Progress in Artificial Intelligence.* Ellis Horwood, Chichester, UK. 184–192.

Reeves, J. F. 1991. Computational morality: A process model of belief conflict and resolution for story understanding. Technical Report UCLA-AI-91-05, UCLA Artificial Intelligence Laboratory.

Reilly, W. S. 1993. Emotion as part of a broad agent architecture. In *Workshop on Architectures underlying motivation and emotion.* The University of Birmingham School of Computer Science and Centre for Research in Cognitive Science.

Rolls, E. T. 1993. A theory of emotion, and its application to understanding the neural basis of emotion. In *Workshop on Architectures underlying motivation and emotion.* The University of Birmingham School of Computer Science and Centre for Research in Cognitive Science.

Scherer, K. 1993. Studying the emotion-antecedent appraisal process: An expert system approach. *Cognition & Emotion* 7(3):325–356.

Sloman, A. 1987. Motives, mechanisms and emotions. *Cognition & Emotion* 1(3):217–234.

Toda, M. 1982. *Man, Robot and Society.* Boston: Martinus Nijhoff Publishing.

Van Lehn, K. 1988. Student modeling. In Polson, M. C., and Richardson, J. J., eds., *Foundations of Intelligent Tutoring Systems.* Lawrence Erlbaum Associates.