Using the Affective Reasoner to Support Social Simulations*

Clark Elliott
The Institute for the Learning Sciences
Northwestern University
1890 Maple Avenue
Evanston, Illinois, 60201
and
Institute for Applied Artificial Intelligence
DePaul University, Chicago
email: elliott@ils.nwu.edu

Abstract

This paper is in two parts. In the first part, the outline of an emotion reasoning architecture, embodied in a simulation program called the Affective Reasoner, is presented, and a rudimentary personality representation for simulated agents is introduced. In the second part, an exercise is reviewed in which the Affective Reasoner is given the task of representing agents with different personality types in such a way as to allow the user to engage in a simulated interaction with them. Representational issues pertaining to the unique appraisal and behavioral styles of the different personality types are addressed. Conclusions are drawn about the usefulness of the Affective Reasoner in such a paradigm.

1 Introduction

A central assumption of this paper is that simulations of social interactions between agents should incorporate models of individual affect and personality. Most human interaction revolves around people's individual needs and goals. These lead to idiosyncratic, internally motivated behavior, and to emotional responses to situations that arise. If these aspects of mental life are not captured in simulations of interpersonal interactions, then simulated agents will be at best bland, lifeless, and unrealistic. This problem has been largely ignored in AI.[1]

This paper describes a general emotion reasoning architecture, embodied in a program called the Affective Reasoner, which has been used as a basis for simulating such interpersonal interactions. Also discussed are the results of using this architecture to solve the problem of representing four distinct client types in an interactive simulation designed to teach selling.

2 The Affective Reasoner

The Affective Reasoner (hereafter AR) is a simulation platform that embraces a wide range of emotion-reasoning issues. In the current research, various worlds are simulated, and populated with agents capable of participating in emotional episodes based on their concerns. Agents are given unique dispositions modeled as a hierarchical set of appraisal frames. These frames represent their individual goals, principles, preferences, and moods. Combinations of the appraisal frames are used to interpret situations that unfold in the simulation. The interpretations, in turn, are characterized in terms of the way they may or may not meet the eliciting conditions for emotions. In some cases emotions result, which then may be expressed, through a set of selectively activated behavioral channels, in ways that are observable by other agents, and as new simulation events which might perturb future situations. Additionally, agents use a case-based heuristic classification system to reason about the emotions of other agents, and to build representations of those other agents' personalities that will help them to predict and explain future emotion episodes involving those observed agents.

2.1 The emotion eliciting condition theory

Embodied in the AR is a set of rules for mapping from emotion eliciting conditions into *emotion types*, based on the work of Ortony, Clore, and Collins [Ortony *et al.*, 1988]. Their theory has been adapted for the simulation and extended from twenty-two to twenty-four categories of valenced reactions to situations, as outlined in figure 1. Each of the twenty-four emotion types has a set of eliciting conditions. When the eliciting conditions are met, and various thresholds have been crossed, corresponding emotions result. A key element of the theory is that the way an emotion-eliciting situation maps into eliciting conditions depends on how an individual

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

[1]But see [Bates *et al.*, 1992] and [Frijda and Swagerman, 1987] for interesting approaches to related problems.

agent interprets that situation. For example, suppose that a captain heroically/foolishly goes down with his ship, while trying to salvage it. On the one hand, his fellow seamen might see this as praiseworthy, since he has upheld the principle of dedication to maritime service. On the other hand his wife might perceive it as blameworthy, since she sees him as having violated the principle of putting the needs of his family foremost. In both cases the act is the same; it is only the *construal* of the situation which is different.

Emotion-eliciting conditions leading to emotions fall into four major categories: those rooted in the effect of events on the *goals* of an agent, those rooted in the *standards and principles* invoked by an act of some agent, those rooted in *tastes and preferences* with respect to objects (including other agents treated as objects), and lastly, selected combinations of these three categories. Another way to view these categories is that they are rooted in an agent's assessment of the *desirability* of some event, the *praiseworthiness or blameworthiness* of some act, the *attractiveness* of some object, or selected combinations of these assessments.

This theory was used as an organizing principle for the appraisal mechanisms of agents. Added to the implementation of the theory are components for actually mapping from situations into the eliciting conditions, an expressive component for generating actions, and a component for reasoning about the emotions of other agents. In addition, extensions to permit the representation of mood and emotional intensity are under development [Elliott and Siegle, 1993]. In the sections that follow, overviews of the three main components of the AR are given. Complexities such as those that arise in the implementation of *expectations*, multiple and conflicting emotions, relationships between agents, and so forth, are not discussed here, but are given a full treatment in [Elliott, 1992].[2]

2.2 Generating "emotions"

Emotion-eliciting situations may arise in the course of the simulation. These situations may or may not be of concern to one or more agents. If they are, then varied interpretations of the situation may be made, depending upon the makeup of the relevant appraisal frames in each agent's *interpretive personality component*. These interpretations are reduced to Emotion Eliciting Condition Relations (EECRs), which in turn are used to generate instantiated emotion templates. Figure 2.2 illustrates the different sources for these emotion templates. Below are annotations for the steps illustrated in the figure:

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

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

[2] For example, mood representation alone affects both the appraisal mechanisms and the expressive components of simulated agents; it affects both the intensity and duration of subsequent affective states; and, it may require as many variables to adequately represent.[Elliott and Siegle, 1993; Frijda *et al.*, 1992; Gilboa *et al.*, 1993; Clore, 1992].

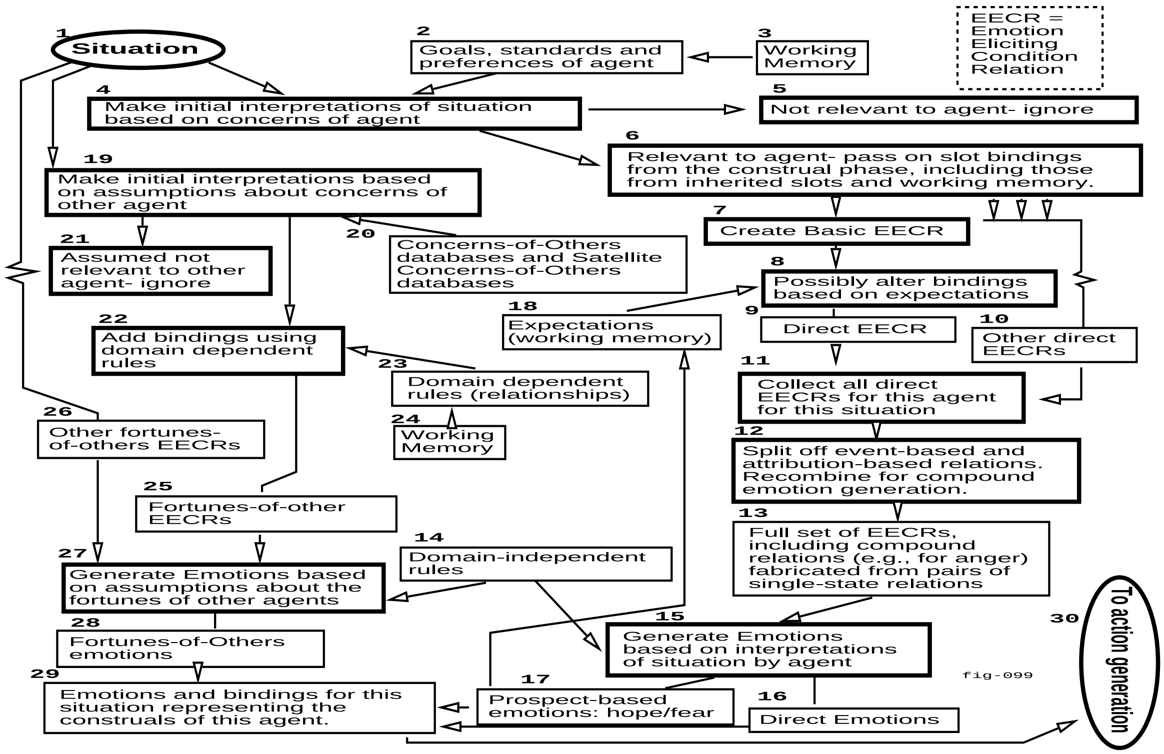

Figure 2: Structure of the mapping from a situation to emotions

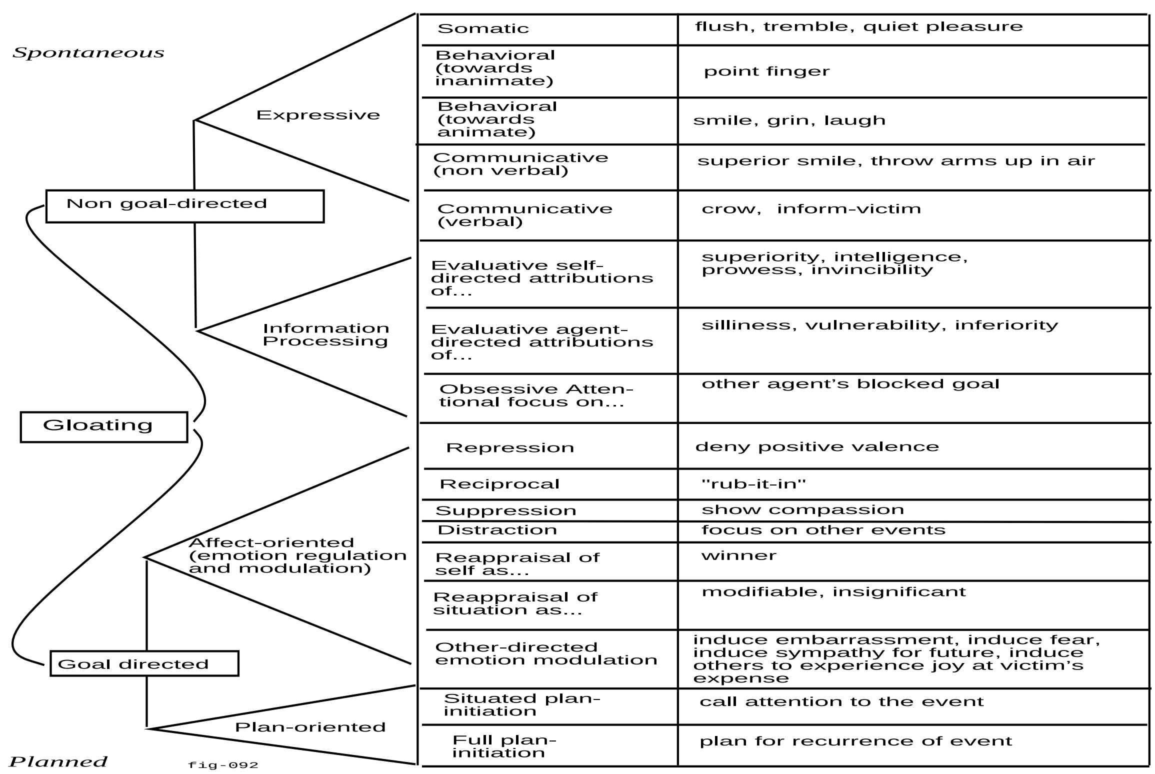

Figure 3: Action Response Categories for *gloating* (Adapted from Gilboa and Ortony, 1991).

1. A situation is created when a simulation event is popped off the queue and the state of the simulated world is altered in a way that might be of concern to one or more agents.

2, 3, 4, 5. The *construal* frames representing the goals, principles, and preferences of some agent are matched against the eliciting situation frame. Working memory is attached to slots in the construal frames and can alter the success of the match, as well as the resultant bindings that are generated. Working memory is also used to implement that portion of *mood* which affects the appraisals of agents [Elliott and Siegle, 1993]. When a match succeeds for some construal frame (and its inherited properties) then the situation is considered to be relevant to the agent's concerns. Since an agent may simultaneously have more than one different construal of the same situation, multiple interpretations of that situation may result. When all matches fail, the situation is not considered relevant to the agent's concerns and is ignored.

6. If the situation is relevant to the concerns of the agent, bindings will have been created during the match process. In addition to those bindings created when slots in the situation frame are unified with pattern-matching variables in slots of the construal frame (including inherited slots), additional bindings may come from attached procedures and from working memory.

7. A basic Emotion Eliciting Condition Relation (EECR) is created for *each* construal of the situation.

8, 9. The *status* attribute of the EECR may be changed to *confirmed* or *disconfirmed* if the situation is relevant to a stored expected outcome.

10, 11. Since there may be multiple construals there may be multiple EECRs. These are collected together before further processing.

12, 13. The event-based construals and attribution-based construals are split off into two (non-intersecting) sets. They are recombined, as applicable, to form the compound-emotions EECRs, subsuming the component EECRs.[3] When recombination is *not* applicable the original EECRs are passed along instead. (For example, there is no compound emotion that corresponds to the eliciting conditions arising when a blameworthy act helps the appraising agent to achieve a goal.) See figure 1.

14. The domain-independent rules contain the backbone of the emotion eliciting condition theory. All emotions are generated using these rules. When a completed EECR matches the left-hand side of one of these rules, an emotion instance is generated.

15, 16, 17, 18. Those emotions not about the fortunes of others are the *direct* emotions. A subclass of these is the prospect-based emotions of *hope* and *fear*. These latter emotions stem from unconfirmed events and so generate *expectation* frames. These frames are stored, and may be used to interpret future situations.

19, 20, 21. For each other agent involved in the situation, the observing agent maintains representations of their concerns in *Concerns-of-Others* databases. These representations are used to interpret the situation in ways presumed to be of concern to other agents, using steps similar to 2, 3 and 4. In addition, Satellite Concerns-of-Others databases, and presumed relationships of others, allow agents to interpret situations with respect to second and third party agents [Elliott and Ortony, 1992].

22, 23, 24. If the situation is determined (partly by assumption) to be of concern to the other agent, then the domain-dependent rules are used to determine the meaning for the target agent. Is the agent pleased about the outcome for the other agent? Relationships between the agents are stored in working memory and may be altered as the simulation progresses.

25, 26, 27, 28. Each of the resulting fortunes-of-other EECRs results in an emotion using the domain-independent rules.

29, 30. The *direct* and *fortunes-of-others* emotions are all collected into a group and passed to the action-generation module. Included in each emotion instance is a set of bindings from both the original match and the intermediate processing.

2.3 Expressing emotions

Once emotional states have been generated for agents in response to eliciting situations, agents manifest these states according to their unique temperament traits. For example, when one agent is angry she might tend to shout, whereas another might smile and deny that anything is wrong. Actions having nothing to do with emotional states are not covered by this mechanism. Actions that do stem from emotional states, however, may be classified into one of the four high-level categories, *expressive, information-processing, affect-oriented (emotion regulation and modulation)*, and *plan-oriented*.[4] Figure 3 gives a breakdown of the action response categories and their theoretical groupings for the emotion *gloating*.

The high-level categories are arranged, as far as possible, from the more spontaneous to the more planned (e.g, *expressive* actions tend to be less planned than those for *emotion regulation and modulation*). Within these high-level categories the particular action response categories are similarly arranged (e.g., *somatic* expressive responses tend to be more spontaneous than *communicative-verbal* expressive responses). Lastly, within the action sets themselves the tokens (and mini-plans) are arranged from the least intense to the most intense (e.g., *smiling* as a *non goal-directed, expressive behavioral-toward-animate* expression of *gloating* is less intense than is *laughing*). In all cases, these are at best partial orderings.[5]

Once an action has been selected, it is instantiated using the set of variable bindings generated from the original match between the construal frame(s) and the eliciting situation frame. For example, the variable *?other-agent* might be bound during the original match process, which could then

[3] Allan Collins points out that this subsumption is a matter of debate.[Collins, 1992] This implementation was arbitrarily chosen.

[4] The high-level categories are not presently used for action generation. One possible use of the theory underlying these categories, however, might be as a basis for a functional representation of emotions as initiators of purposeful, *motivated* actions. In this regard, the *information processing* actions, for example, might be seen as representing aspects of an agent's attempts to "understand what is going on."

[5] This is not an attempt to list action *words* that might be associated with a particular emotion. While this might be interesting with respect to understanding or generating text, it is not relevant to the goal of characterizing psychological states.

be used to instantiate the action response template <*laugh at ?other-agent*>.

From figure 3 we can see that *superior smile* and *throwing arms up in the air* are two ways of non-verbally communicating *gloating*, and that *inducing embarrassment* and *inducing others to experience joy at the victim's expense* are two ways of *modulating the emotions of others*. The expressive component is closely related to the domain in which the AR is being used, so only those expressive actions which do not further perturb the system (i.e., are not expressed as relevant simulation events) are directly portable from one application to another. The largest implementation, TaxiWorld [Elliott, 1992], used about 1000 different action paths. Multiple actions may simultaneously express a single emotion instance.[6]

2.4 Observing and modeling other agents

The AR has mechanisms whereby simulated agents may reason about the emotions, and the personality makeup, of other agents. To do this, features derived from an observed eliciting situation, and the observed agent's response to the situation, are formatted as a *case* for a slightly modified version of Protos [Bareiss, 1989], a heuristic classification system. Once an observing agent has classified the case as belonging to an instance of one of the 24 emotion types, the agent stores this knowledge as part of its internal representation of the other agent. For example, one agent might observe another frown and speak rudely when kept waiting. From prior experience the observing agent knows that its counterpart often expresses anger by speaking rudely and frowning. Using this knowledge to abduce anger it then looks for interpretations of the eliciting situation that would lead to anger. Once a suitable explanation is found (such as the violation of a principle that one *should not* be kept waiting, and the thwarting of a time-dependent goal) it updates its internal representation of the other agent so that those concerns are included.

Several levels of defaults are required to implement this mechanism. Searches might have to be made through the representation of the current presumed concerns of the other agent, through defaults for the *role* of the agent, through defaults for agents that are "similar" in nature, and through the system catalog of "what might explain" the response.[7] Also necessary are *satellite concerns-of-others* knowledge bases used in conjunction with a representation of the *presumed relationships-of-others* since observed agents may have emotional responses based on the fortunes of third-party agents as well [Elliott and Ortony, 1992].

[6]Other issues that have had to be addressed include determining the dimensions along which to tune a particular agent's temperament traits (e.g., an agent who grows quiet when under the influence of negative emotions, or an agent who is task-oriented, and vocal, with respect to standards-based emotions), the interplay and consistency of multiple manifestations of an emotion (e.g., some agent might deny anger when she is angry, or might shout about it, but will not do both at the same time), dynamic changes in the agent's temperament (e.g., tending toward low energy before meal times), and so forth.

[7]It is worth mentioning that versions of the AR have been run which focus almost entirely on these aspects. In one, the system itself, as an automated agent, reasons about the emotions of the user within a constrained format. Although this idea has not been developed, it suggests at least the possibility of a simple form of affective user modeling.

3 Social simulation using the Affective Reasoner

The AR has been used as a basis for constructing a social simulations. In this environment, users interact with one or more simulated agents via the computer keyboard and mouse. Users get feedback through simulated facial expressions for the agents, and through formalized text describing what is taking place in the simulated world. In this scheme, users take part in the simulation by making decisions for one of the agents. The other agents respond to the situations that arise in the course of the simulation, including those that come about as a result of the user's input. Simulated agents represent individuals with emotional depth and their own idiosyncratic ways of expressing their feelings. A user's task is to identify the concerns of the other agents so that he or she can interact with them in ways that successfully bring about some best-case outcome.

During the course of the simulation, agents are continually appraising situations that arise in their world. Some of these are of concern to them, and lead to emotional states. Once such a state has been established, expressions of it may be manifested within the simulation. For example, the displayed face of the agent might change, working memory representing the tolerance level of the agent for further similar situations might change, the level of attentional focus for the agent might change, or the agent might be described, in text, as laughing or turning red.

Rudimentary personality types are configured before the simulation is run, either by default, or through user-selection of attributes. In this paper, four statically pre-determined personalities will be discussed. In general, however, the appraisal frames and expression paths represented in the knowledge base may be combined as desired, creating a wide variety of personalities. The role of an agent in the simulation determines which situations that agent will be involved in, and thus which situations might be subject to interpretation. For example, supposing that we are simulating a sales situation, a *secretary* agent might not be concerned with situations arising during the *close*, but might be concerned with situations arising during the *initial greeting* phase of the salesperson's visit.

As in real social situations, the state of an agent changes over time. For example, a client might gradually become bored, especially attentive, warm, or hostile. Users must respond in quasi-real time to the situations that arise. If a user does nothing, or responds too slowly, then agents might take the initiative, because their concerns are not being addressed.

4 The problem

4.1 Problem summary

For the work reported in this paper, the attempt was made to model a set of four client types identified as part of a training program designed to teach novice account executives how to sell telephone-book advertising (cf. [Kass *et al.*, 1992]). These four types, here referred to as *dominant, political, steady*, and *wary*, were used to characterize the typical sorts of clients one might expect to encounter in the field. The definitions were both practical and operational: they were described in a manner congenial to use in training systems, rather than as set of formal psychological dimensions, and they were partly defined by the the way in which such client types might be expected to respond to particular actions of a salesperson.

Descriptions were given in two ways. First, the client types were partially described by a set of rules prescribing what the salesperson should and should not do in sales situations with different kinds of clients. One rule, for example, says that if one "takes charge" with a dominant client, then one should expect to succeed only in making the client angry. Approximately 21 such rules were given for each client type. Second, the personality types were described in terms of their observable behavioral styles (e.g., laconic or talkative) Approximately 14 such style characteristics were given for each client type.

The given task was not to build a full sales-training system, but rather to explore the representational power of the theory-based architecture in a practical domain. For this reason no attempt was made to model those processes that had nothing to do with either personality or the generation or expression of affective states. Additionally, the exercise was further constrained by using only a very simple conceptualization of negotiation.

4.2 Representing situations and agents

For this exercise, a sales episode was broken down into seven steps: (1) The greeting, (2) the fact-gathering sessions, (3) the making of the proposal, (4) support for the client's decision, (5) the client's agreement or disagreement, (6) the negotiation process, and (7) the parting. Each of these steps was represented as one or more simulation events. Step (5) had four different instantiations depending on whether or not the client agreed with the proposal, and on whether the salesperson wished to modify the client's stance. In some instances it caused a loop back to step (3).

The four client types represented were: (1) *Dom*, a dominant personality who prefers logical, rapid decisions, (2) *Pol*, a political personality who likes to socialize and is very status conscious, (3) *Pas*, a passive personality who does not express objections, but who might act on them, and (4) *War*, a wary personality who often tends to be anxious about anything unfamiliar.

To illustrate these different client types, let us consider how each of the them might be expected to appraise the salesperson's style at the opening of the interview. **Dom** holds a principle that the salesperson should be brief, specific and to the point. Rambling will waste her time, violating that principle and interfering with her goals, thus, ultimately, incurring anger. **Pol** has a preference for social niceties. He has a goal of defining roles in the social structure. To rush immediately into the business at hand thwarts this goal and will not please him. **Pas** is not comfortable with rushing into business either, but, unlike Pol, he is primarily seeking to establish that the salesman has a genuine interest in his needs and preferences. **War**, like Dom, prefers getting right to business. Unlike Dom, however, she wants all the steps laid out methodically. In War's case, the introduction is especially important in setting the tone for the session. War tends to imagine the negative consequences of suggested courses of action, and it is important that the opening of the interview set her at ease.

The simulated agents also tend to express their emotions differently. Dom, for example, had activated action paths for *evaluative self-directed attributions tending towards superiority, competence, intelligence*, an *obsessive attentional focus on goals*, and so forth. By contrast, War had paths emphasizing *evaluative self-directed attributions tending toward inferiority and powerlessness*, tendencies to *reappraise a situation in a negative light*, and so on.

4.2.1 User choices

The user, taking the role of the salesperson, affects the course of the simulation. At each stage during the simulated interview the user has menu choices for what to do next. Once the user selects an option, a frame representing the features of the user's choice is posted to the simulation queue and executed. These and other simulation events give rise to situations to which the simulated agents respond in accordance with their concerns.

New situations, and the agents' responses to them, arise as a result of the user's choices (e.g., attempting to "close" too soon might lead to anger on the part of Dom, but anxiety on the part of War), in response to actions on the part of the simulated agents (e.g., when Dom is angry she might terminate the interview), or by default (e.g., if the user does nothing Pas grows uncomfortable, and Dom gets angry).

In addition to the user choices that arise at the different stages in the course of the interview, there are also additional user choices representing the overall *style* of the salesperson. For example, the user might set *pacing* (moderate, slow-tutorial, fast-hyped), *tone* (friendly, cool), *attitude* (deferential, take-charge), *style* (brief-summary, anecdotal, extremely-detailed), and so forth, which also can affect the course of the interview. The style of the salesperson is dynamically configurable as the simulation progresses.

4.3 Conclusions

On the overall representational task there were mixed results. The approach was successful in giving agents the ability to appraise situations so that their concerns led to emotions consistent with their intended personalities. Eighty-six rules specifying "what to do" and "what not to do" were successfully represented within the simulation. The rules could all be directly related, for each of the client types, to appraisals of situations that arose in the course of the simulation. This was the most effective component in the AR, with respect to the constrained representational task.

To account for the content of the 86 sales domain rules, each was first analyzed with respect to the way in which it uniquely characterized the associated personality type. For each rule, typical emotions such as anger, joy, fear, and so forth, could be inferred from the clients' stated or implied reactions to the salesperson's initiatives. Working back from these emotions, it was possible to determine high-level frames for goals, principles, and preferences which would lead to the associated emotion types. From there, lower-level appraisal frames suitable for interpreting the actual emotion-eliciting situations within the simulation could be identified. These were then used as the basis for building the interpretive personality for that particular client type. For example, from the rule that says *never talk down to Pol* (the political client type) we might infer that Pol could grow angry when treated this way. Within the constraints of the theory, this would be true if his personal goals involved maintaining a certain status within interpersonal relationships, and his principles said that salespeople should be aware of this.

By contrast, there was less success in representing the given behavioral styles of the client types within the AR. The manifestations of these styles fell into three groups. The first group contained those characteristics that suggested certain actions. For example, *harsh* might suggest that, when angry, a client might be blunt, or shout. The second group contained those behavioral styles directly related to appraisals, such as *judgmental*, *perfectionist*, and *excitable*. The third group contained those behavioral styles that were ubiquitous in the client types' respective dispositional manners. For example, clients were described as *dependable* and *persistent*.

The first group was somewhat amenable to representation, but an interpretation of the data had to made. For example, the concept of a *manipulative* behavioral style is generally comprehensible by people without further specification. However, it seems likely that the AR's representation of this trait as a tendency toward *other-directed emotion modulation* falls short of the mark. One might argue that manipulative people behave as they do quite independently of their emotions.

The second group of behavioral styles was readily represented as part of the interpretive personalities of the agents. For example, for an agent to be *judgmental* simply meant that either the thresholds for invoking standards leading to negative emotions were lowered, or that more appraisal frames leading to those emotions were in place.

The manifestations of behavioral style in the third group were not amenable to representation in this architecture. For example, being *slow to make decisions* probably does not routinely have an affective origin. The best the AR could do was to represent those emotional reactions an agent might have, for example, when under pressure to make a decision quickly.

As illustrated above, because the AR's expressive component is designed strictly as a set of categorical responses to simulated emotions, it is not well suited for representing dispositional behaviors rooted in more cognitive motivations. A better solution would have to additionally integrate some representation of individuals' long term planning and survival strategies.

Another problem was in making the personality types accessible to the casual user. The intuitive grouping of appraisal frames into sets that form identifiable components of complete interpretive personality types is a hard problem on which little progress was made. Standard personality distinctions from personality theory in psychology (e.g., [Digman, 1990]) were felt to be unhelpful because the constructs they employ are not likely to be generally accessible to most users. In addition, attempting to use them to identify the types of rudimentary personalities represented in the AR seems pointless because there is so much about a realistic characterization of personality for which no attempt at modeling was made. It would be difficult, for example, to adequately characterize the relationship between sociability, or extraversion, and the interpretation of a salesman's greeting.

From this modest exercise the following conclusions may be drawn: (1) the Affective Reasoner is well suited for representing the interpretive disposition of simulated agents in social simulations, (2) it is capable of representing some aspects of the intended behavioral styles of such agents, (3) it is awkward at best, in its current form, for representing the numerous commonly understood behavioral style characteristics which have little to do with emotional states, and (4) it is capable of creating many internally consistent rudimentary personality types, but that subcomponents of these types are difficult to characterize in a way meaningful to users.

Acknowledgments

I would like to thank Greg Siegle, and Robin Burke for their editorial comments on this paper. I would especially like to thank Andrew Ortony for not only his editorial comments, but also for his support and advice throughout the development of this research project.

References

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

[Bates *et al.*, 1992] Joseph Bates, A. Bryan Loyall, and W. Scott Reilly. Integrating reactivity, goals, and emotion in a broad agent. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*, Bloomington, IN, August 1992. Cognitive Science Society.

[Clore, 1992] Gerald Clore. Cognitive phenomenology: Feelings and the construction of judgment. In L. L. Martin and A. Tesser, editors, *The Construction of Social Judgments.* Lawrence Erlbaum Associates, 1992. To appear.

[Collins, 1992] Allan Collins, April 1992. Personal communication.

[Digman, 1990] John Digman. Personality structure: Emergence of the five-factor model. *Annual Review of Psychology*, 41:417–440, 1990.

[Elliott and Ortony, 1992] Clark Elliott and Andrew Ortony. Point of view: Reasoning about the concerns of others. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*, pages 809–814, Bloomington, IN, August 1992. Cognitive Science Society.

[Elliott and Siegle, 1993] Clark Elliott and Greg Siegle. 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*, pages 58–67. American Association for Artificial Intelligence, 1993. Stanford University, March 23-25, Palo Alto, CA.

[Elliott, 1992] Clark Elliott. *The Affective Reasoner: A Process Model of Emotions in a Multi-agent System.* PhD thesis, Northwestern University, May 1992. The Institute for the Learning Sciences, Technical Report No. 32.

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

[Frijda *et al.*, 1992] Nico H. Frijda, Andrew Ortony, Joep Sonnemans, and Gerald L. Clore. The complexity of intensity: Issues concerning the structure of emotion intensity. In M. Clark, editor, *Emotion: Review of Personality and Social Psychology*, volume 13. Sage, Newbury Park, CA, 1992.

[Gilboa and Ortony, 1991] Eva Gilboa and Andrew Ortony. The structure of emotion response tendencies. Work in progress, 1991.

[Gilboa *et al.*, 1993] Eva Gilboa, Andrew Ortony, and William Revelle. Personality and the structure of emotional responses. In S. Van Goozen, N. E. Van de Poll, and J. A. Sargent, editors, *Emotions: Essays on Emotion Theory.* Lawrence Erlbaum Associates, 1993.

[Kass *et al.*, 1992] Alex Kass, Robin Burke, Eli Blevis, and Mary Williamson. The GuSS project: Integrating instruction and practice through guided social simulation. Technical Report 34, The Institute for the Learning Sciences, Northwestern University, May 1992.

[O'Rorke and Ortony, 1994] Paul O'Rorke and Andrew Ortony. Explaining emotions. *Cognitive Science*, 18:283–323, 1994.

[Ortony *et al.*, 1988] Andrew Ortony, Gerald L. Clore, and Allan Collins. *The Cognitive Structure of Emotions.* Cambridge University Press, 1988.