Emotionally Responsive Poker Playing Agents *

Stuart Marquis
Department of Computer Science
DePaul University
243 South Wabash Avenue
Chicago, IL 60604
email: smarquis@hawk.depaul.edu

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

Introduction

This paper describes a set of poker-playing agents that embody components of our work on broad, shallow, emotion and personality representation on the computer. The agents, with whom the user plays, have rudimentary "emotional," as well as purely cognitive, responses to the events that arise in the course of the game. Emotion eliciting events include those based on the game itself, such as being dealt a strong hand, and also spoken input from the user, an opponent. Cognitive responses include, for example, the choice of when to bluff, and when to play a hand "straight". Such cognitive decisions, in turn, can in themselves also lead to emotional responses. The poker application code is run on top of portions of a large LISP program known as the Affective Reasoner. The Affective Reasoner (hereafter AR), as well as preliminary efforts at using a multi-media interface for this program to communicate with the user, have been discussed elsewhere (Elliott 1992; 1994; Elliott & Ortony 1992).

The automated agents in the system manifest "real" and "faked" emotions through facial expressions, inflected speech synthesis, and music. An agent's face continually morphs between destination bitmaps, the same set of 72 faces used in the Affective Reasoner. Spoken text is first constructed in LISP (primarily from relatively fixed strings, currently) and speech is generated by a sound chip using different simple inflections based on changes in speed, pitch, and volume. Music is in MIDI format, yielding almost instant access capability for hundreds of hours of recorded music. Music playback is designed to run through a 400-voice Proteus MS-PLUS synthesizer. Quality is such that, depending on circumstances (i.e., in non-critical listening situations), listeners have often mistaken the music as sound recordings of acoustic orchestral, and other, instruments.

The game has been played with up to five players, including the user, but current development but has focused more on games between the user and a single agent.

The Affective Reasoner derives partially from a theory of the cognitive antecedents of emotion originally specified by Ortony, Clore and Collins (Ortony, Clore, & Collins 1988). Agents appraise their simulated world according to their individual concerns. Sometimes these appraisals, based on the agents' goals, principles, and preferences, give rise to "emotion" states which fall into twenty-four distinct categories. These states then, may in turn be manifested through approximately 450 "expression" channels. Channels range from the most spontaneous (such as somatic responses) to the most planned (such as planning for a recurrence of the event), for each of the emotion categories. Additionally, agents can do rudimentary reasoning about each other, and about the user, using a case-based heuristic emotion classification system based on Bareiss's Protos (Bareiss 1989), and internal models of the presumed concerns of the other agents (and the user).

In exploring applications of the AR, and the more recent multimedia extensions to it, we chose poker playing as one area of study because there are some simple, well-defined emotion episodes that can be seen as arising throughout the course of a game of poker, and a real-time environment in which to observe relatively fast-paced interaction with the user. There is also a degree of constrained, and somewhat formalized, interaction between the emotional and cognitive concerns of poker playing agents: to wit, good and bad situations may give rise to positive and negative emotions, but it is in the interests of the poker player to be able to filter their expression. For example, poker has *bluffing*, which introduces additional interactions between the honest "emotion" expression of earlier AR simulations, and the deliberate cognitively controlled expression of false emotions.

Exploiting the Multimedia Platform

This simulation runs on a 66 Mhz IBM clone Multimedia PC, with 16M of real memory, using Windows, and Franz Inc.'s Allegro Common LISP for the PC. The LISP engine communicates with the text-to-

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

speech, speech recognition, and face-morphing engines through Windows DDE interfaces, and to the MIDI playback system through direct DLL calls.

The application runs with a graphical interface, with areas of the screen used to display the players' chips, the pot, bitmaps of the cards, etc. Commands are issued through buttons and check-boxes or voice control. Our voice recognition software is designed only to produce keyboard events from recognized speech. Since the Windows environment traditionally provides keyboard shortcuts for most mouse input, this has not been a limitation. The interface is then run hands-free, and the recognized grammar is simple enough (and redundant enough) that anyone familiar with the game of poker can play with almost no assistance.

Using an event-driven graphical interface allows some degree of real-time processing in our simulation. Game events take time to occur, and our agents take time to make decisions, based on our ad hoc estimates of how difficult their decisions are. Also, our agents measure how long the user takes to bet or exchange cards. If the user takes an especially long time to make a decision, our agents will revise their estimate of the user's hand and possibly have "emotional" reactions to a user taking too long.

Multimedia Support for Expressing Emotions

The current simulation uses line-drawings of a set of "emotion" faces, recently converted to a Windows DDE server interface by Yee-Yee Yang (Elliott, Yang, & Nerheim-Wolfe 1993). Description of such graphical presentations is sometimes difficult in text. Briefly, the impression is one of a quite responsive agent. For example, based on real-time events, such as the agent being dealt a particularly fine hand, or the user speaking to the agent, the facial expression may immediately begin to morph, gradually breaking from a frown into a smile: the eye-brows slowly wrinkle, the mouth contorts, the eyes widen. There are 72 bitmap faces available in the system, and up to three intensity levels for each of the 24 emotion categories. The algorithm runs in real time, morphing from any bitmap to any other, effecting roughly 4,000 possible morphs thus yielding a relatively rich platform for expression. At any point during the morph, or at a static facial destination, if the agent starts to speak, the mouth moves up and down. Also, we have run-time control over the speed of the morph, so that, for example, it is possible to make strong emotions appear almost instantaneously, and less intense emotions take longer. Each agent has a default face defined, and in the absence of emotion eliciting situations their face will slowly decay back to its default state over the course of several minutes. That the faces are not "realistic" is not of concern to us: it is seems clear from our preliminary work that they convey much of the intended message to the user in the same way that other line drawings, such as those in "Calvin and Hobbes," do.

The agents communicate with the user through speech. The text-to-speech package itself is rather simple, but does allow signals to control its pitch, speed, and volume while it speaks a phrase. We have used this speech modulation capability in conjunction with the AR to produce "emotionally appropriate" speech. For example, when the agent is *distressed*, any phrases it has to communicate will be sent to the speech engine tagged *<low volume><slow speed><falling pitch>*. Our mapping of emotions to volume, speed, and pitch is based on Collier's summary of previous research(Collier 1985), extended to the AR's 24 emotion categories. Using speech inflection allows emotionally neutral phrases (e.g., "it's your bet") to be communicated in an "emotional" way. Thus the emotion, to some degree, can be expressed in the medium, independently of the content. Each of our emotion categories also has a number of expressive phrases associated with it that are used from time to time (e.g., "hooray", "oh, rats", "congratulations").

The simulation also has a large library of midi recordings. These files, mostly snippits from classical music, have been ad-hoc tagged with emotional content. The agents have the capability of playing music that roughly corresponds to their "emotion" state. This hotly debated area is one of great expressive capability: music has long been understood (by musicians and filmmakers, as well as psychology researchers) to have strong mood-induction properties. While we have only just begun to explore the possible mappings between emotion expression and music, the claim might well be made that a platform such as this can provide a fruitful platform for studying this interaction.

It is our belief that these components, together, show promise of making an effective interface. The combination of several channels of expression has been shown to significantly increase the accuracy of communicating emotional state (Scherer 1988; Collier 1985). This improvement is compounded by using dynamic channels – morphing faces instead of static bitmaps, and fluctuating vocal inflection instead of monotone. Suppose the computer agent is dealt his hand. He then frowns, a dirge starts playing in the background, and he mumbles "oh rats" in a voice almost too low to discern. Together this conveys his "emotion" state fairly well – that is unless he is deceiving us! (See *Bluffing* below.) It might then be assumed by the user that his dejection is as a result of recent events, (e.g., the hand just dealt to him), and this might encourage the user to bolder betting against the computer.

Emotion Recognition through Voice Inflection

We have built on our voice-command interface, described above, to add rudimentary emotion recognition from speech. In other (non-poker) work we have built vocabularies of simple phrases, said with different emotion inflections. In our preliminary efforts we found that the package was good enough to recognize each phrase (i.e., the same phrase repeated with a different inflection) as being in one of seven broad emotion categories (neutral, joy, fear, anger, hate, sad, love). This is not so fine-grained as the 24 separate emotion and mood categories known by the affective reasoner, but it allows for a rough appraisal by the computer of the emotional state of the user. The success rate is surprisingly good on correctly distinguishing the inflections.

The method for achieving emotion recognition is somewhat baroque, and worth mentioning. Direct detection of voice inflection is not intended to be supported by our voice recognition software, and in fact is actually considered a deficiency. (That is, for a *command* interface, the primary purpose of the system, inflection should be ignored.) To work around this, we added extra "words" to the base vocabulary such as, for example, "joy-bet" and "sad-bet." We then trained these words to be the spoken words, respectively, "Bet" said in a joyful voice and "Bet" said in a sad voice. In this way the voice-recognition software keeps track of "joy-bet", "bet", and "sad-bet" as three independent sound-patterns, but after parsing, and stripping off, the emotion inflection, we define the same keyboard-translation to all three, so they all result in a bet action in the interface.

In this manner then, our voice recognition software has a primitive ability to distinguish three separate emotional states based on the user's speech inflection. Detection of *joy* or *sadness* generates an additional, hidden event (besides the *Bet*, or *Pass* and so forth), a *user-feels-joy* or *user-feels-distress* simulation event. This simulation event is fed to the poker-hand logic, to update the computer agents' estimates of the player's hand, and also entered into the appraisal-frame processing of the emotion reasoning code. All the defined computer-agent personalities update their hand estimates in the expected manner (*user-feels-joy* → increase estimate, *user-feels-distress* → lower estimate), but different agents have different emotional appraisals of the event. "Rocky" is defined as empathetic to the human – when the user is happy, he is happy for the user. Conversely, when the user is sad, Rocky feels pity for the user. "Boris", on the other hand, is not so friendly. He is defined just the opposite, so that the human's sadness brings him gloating, and vice-versa.

There are some issues that have arisen out of developing even this rudimentary "emotion" detection. For one, the tone of voice that the human uses can be a result of many things besides emotions: distractions in the room, posture, eating or drinking, even deliberately misleading intonation by the user. However, this is not necessarily *wrong*, because bad signals are a problem that real poker players have to deal with as well, albeit on a more subtle basis. This is especially true when the user is either bluffing his or her tone, or intentionally holding to a monotone (the acoustical "poker face"). Another issue is the limited inflection detection capabilities of the system, although we see this as a problem that improvements in technology may well address.

Cognitive Control Over Emotional Expression – Bluffing

Our agents sometimes bluff. To simulate bluffing we have given the agents (in this case, minimal) logic for deciding when to bluff and how to bluff. We have also simulated the varying degrees of skill and effectiveness at bluffing implemented as the agents' ability to intercept and alter the "natural" manifestations of their "emotions."

In our rudimentary model, the basic idea behind bluffing is to choose a false worth for your hand, then make bets consistently with that false worth instead of the true worth. Throughout the rounds of betting, every time an agent is called upon to make a bet, he determines whether this would be a good time to begin bluffing. This determination is based on the current worth of the hand, the total chips bet so far, how early in the rounds of betting it is, a simple memory of how successful betting has been against this opponent in the past, and a personality-defined propensity for bluffing.

If an agent "decides" to begin bluffing, it first generates a false worth for its hand. A bluffer may bluff high, hoping to make other players fold, or bluff low, hoping to induce higher bets from other players. The computer agents will bluff in one of these two ways depending on the true worth of their hand. After the false worth has been decided, all future betting decisions will be based on that worth. Within a hand an agent never stops bluffing to recommence honest betting. Our bluffing strategies were culled from a standard poker text (Spanier 1977).

In addition, now that we have effectively added a new kind of event to the simulation, *bluffing*, we also add appraisal frames for *bluffing* simulation events which then may in themselves lead to "emotions." For example, the agent may appraise a successful bluff as a goal achieved, leading to a joy "emotion", or may appraise an opponent's bluffing as going against a standard of honesty leading to "reproach."

Mechanisms for Control of Expression

It would be useless for an agent to bet large amounts on a poor hand if her face clearly showed great fear. Users would immediately see through this. A successful bluffer misrepresents his true emotional state by using the regular emotion expression channels to communicate a fictitious emotional state.

The poker agents bluff by "cognitively" intercepting emotional behaviors before they are expressed. The AR always causes the agents to appraise the situation as they normally would (e.g., *distressed* over a poor hand), and generates any "emotion" states that might arise. These "emotions" are then mapped to actions (i.e., one or more of face morphing, music, and speech). If the agent is not bluffing, those emotional actions are expressed, and the user can observe the results. If the agent is bluffing, then the agent tries to block these expressions of emotion, and replace them with expressions of false emotions.

This is only one of many explanations of deceit in emotional expression. We assume that the agents really have the emotions, as defined by their appraisal of events, but that they do not express those emotions. Other emotion researchers have suggested that emotions are always accurately expressed, and that the only way to eliminate *leakage* of true emotions is through self-deception. The agent deliberately experiences an untrue emotion (Ekman 1982). Another approach is for an agent to dampen their emotional reactions while bluffing, so that although they experience honest emotions, any expressions of these emotions would be of such low intensity as to go undetected. Buck distinguishes between *stressful* and *unemotional* deception, and cites studies that have shown that subjects in unemotional situations do not display the telltale non-verbal cues of deception (Buck 1984).

Each agent is defined with a skill at bluffing, which determines his lesser or greater ability to successfully intercept these expressions of "emotions." If the expression is successfully intercepted, the agent expresses the inverse of its true "emotion." (One of the characteristics of the underlying emotion theory is that emotions are *valenced* reactions to situations, and that for each positively valenced emotion there is a negatively valenced one, and vice versa). *Distress* becomes *joy*, *fear* becomes *hope*, etc., and these inverse emotions are then mapped to actions, just as before, which are then expressed. This implies either that our agents have a perfect understanding of the emotional theory, or that they use some mechanism to generate normal responses to planned emotional states. This is probably not accurate, but is a sufficient model for this preliminary work. A more advanced method would be for our agents to each have a memory of how emotions get expressed, and to act according to this memory, whether or not it accurately reflected true emotional expression. A bluffed expression would be the agent's best guess at how the false emotion should be expressed.

Recall that our agents currently have three channels in which to express their emotions: face, music, and voice. The interaction between these three channels can be handled in a number of ways. Some agents intercept all three expression channels at once. In this case, based on their predefined skill, each time a "real emotion" is to be expressed in a bluffing situation, the agent will try to bluff, and if successful, will cause all three channels to consistently express the false emotion.

For other agents, each of the three action paths are independently intercepted. (People too seem to have partial, not all-or-nothing, control over their emotional behavior.) This means that an agent may correctly control his facial expression, but fail to control his musical expression, allowing the wrong (that is, honest) background music to play. If the user is paying attention, this inconsistency is a clear indication of bluffing (honest behavior is never inconsistent), but because it is mixed, it is still ambiguous which channels are honest expressions, and which are the bluffed.

For some agents, channels have also been tied together, without making them completely unified. In this case, if an agent can control her face, she has a better chance of getting her voice under control. If she looses control of what she is saying, she looses her facial composure also.

It is also possible that channels should affect each other inversely. If an agent manages to successfully fake her facial expression, then perhaps her true emotion has a better chance of being expressed along different paths. In other words, through displacement, if the agent blocks one channel of expression, others simply become *more* likely to express the emotion. We have also ran agents defined in this way.

Another approach is to treat the blocking of the honest emotion and the expression of the fictitious emotion as separate functions. This leads to more complex interactions. If the blocking and the feigning are both unsuccessful, the agent expresses an honest emotion. If both are successful, the agent expresses the false emotion. If the agent blocks his true emotions but doesn't manage to express his false emotions, he will express nothing – the interface will have a flat expression on the face, and nothing coming out of the speaker. If the agent fails to block the true emotion, but expresses the false emotion anyway, multiple emotions will seek expression along the same paths. People might display this confusion by oscillating between two facial expressions, or trying to say two different things at once. Our agents can simulate both of these behaviors.

Concluding Remarks

The experience of playing against such agents is quite different ¿from playing against more static agents: Facial movement is arresting, real-time cues (such as the agents growing reproachful if the user takes too long to bet) demand that the user address the concerns of the automated agents. The responsiveness of the music to the changing situation enhances the intensity of the playing experience. That agents speak, and listen, gives interaction with them a natural feel. It is quite easy to become engaged in attempting to decipher an agent's ploys: is it happy because it has a good hand, or is it bluffing? If we guess wrong and it is actually bluffing, is the agent going to gloat over us?

Our investigations have brought out some issues in human-computer interaction. There are many advantages to using non-traditional channels of communication. Using multiple channels of communication increases the "bandwidth" of communication. Raw data and an appraisal of that data can be transmitted at the same time. The multiple channels can reinforce each other, to deliver a stronger or clearer message. In many cases, information is ambiguous unless transformed into a more understandable format. Emotional expression serves, in some sense, the same role as charts or graphs, to present the same information in a more accessible format. Multiple channels can qualify a message, toning down dramatic news, or emphasize something that might otherwise not stand out.

There are some risks associated with this. Non-textual communication cues have been shown to increase the perceiver's confidence in the data (Collier 1985). In human-human interaction, non-verbal channels are more reliable, simply because they are more difficult to intentionally mis-communicate. There are certain physiological signs of stress that are nearly impossible to mask during deliberate deception (Frijda 1986). For our computer agents, however, all emotion expression is deliberate. The computer can misrepresent its emotional state just as easily as it can express its true state. Even for "honest" agents, emotions are based on their own appraisal of events, and do not necessarily correspond what a human would expect. The addition of emotional channels of communications may lead the human perceiver to put unwarranted trust in the message.

References

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

[2] Buck, R. 1984. *The Communication of Emotion.* New York: The Guilford Press.

[3] Collier, G. 1985. *Emotional Expression.* Hillsdale, N.J.: Lawrence Erlbaum Associates.

[4] Ekman, P. 1982. *Emotion in the Human Face.* Cambridge: Cambridge University Press.

[5] 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.

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

[7] 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.

[8] 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.

[9] Frijda, N. 1986. *The Emotions.* Cambridge: Cambridge University Press.

[10] 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. THIS PAPER.

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

[12] Scherer, K., ed. 1988. *Facets of Emotion: Recent Research.* Hillsdale, N.J.: Lawrence Erlbaum Associates.

[13] Spanier, D. 1977. *Total Poker.* New York: Simon and Schuster.