Emotionally Expressive Agents

Magy Seif El-Nasr*, Thomas R. Ioerger, John Yen
Department of Computer Science
Texas A&M University

Donald H. House, Frederic I. Parke
Visualization Laboratory
Texas A&M University

Abstract

The ability to express emotions is important for creating believable interactive characters. To simulate emotional expressions in an interactive environment, an intelligent agent needs both an adaptive model for generating believable responses, and a visualization model for mapping emotions into facial expressions. Recent advances in intelligent agents and in facial modeling have produced effective algorithms for these tasks independently. In this paper, we describe a method for integrating these algorithms to create an interactive simulation of an agent that produces appropriate facial expressions in a dynamic environment. Our approach to combining a model of emotions with a facial model represents a first step towards developing the technology of a truly believable interactive agent, which has a wide range of applications from designing intelligent training systems to video games and animation tools.

1 Introduction

The ability to simulate believable characters or agents in an interactive system is a challenging problem that has applications in many areas [20, 7]. These types of agents present a visual interface, typically a human face, which is familiar and non-intimidating, making users feel comfortable when interacting with a computer. These agents can play the role of a personal assistant, and advisor, a tutor, or a synthetic character in a simulation or entertainment software system. Faces provide good interface elements in computers because they take advantage of people's capability for recognizing expressions, and hence can be used to effectively communicate certain types of information, particularly emotions [4]. Believable agents could enhance interactive applications such as information kiosks by making them more accessible to the general public. There are also many applications in training and education that could benefit from the simulation of believable characters, such as training teachers how to deal with various types of students by using artificially-generated classroom scenarios. Visualization of faces and expressions is often employed in animations and computer games. From inverse kinematics to artificially-generated textures to simulations of fabric dynamics, there has been a powerful drive to automate the animation of various subjects with increasing speed and realism [13], and incorporating an emotional model in characters is another key step toward creating effective behavior-based models.

An important aspect of systems that incorporate believable agents is the ability to simulate realistic emotional responses. Producing emotional responses requires both the ability to generate facial expressions, and a model that the agent can use for synthesizing appropriate emotions in a dynamic environment. There has been a great deal of work done on generating believable facial expressions [17]. A number of basic facial motions, often derived from unique muscle groups, have been identified, such as a specific twist of the eyebrow or pucker of the lips, which can be combined to form most facial expressions. These can be implemented as controls and attached to a 3D model to generate a wide range of facial deformations. Finally, the generation of expressions can be described as a mapping onto these controls, such that each expression consists of a unique, generally linear, combination of the settings for the underlying controls that determine how to render the appropriate face.

While research on facial expressions has led to the development of sophisticated algorithms for producing reasonable simulations of various emotional states, this work generally does not specify how to model the emotional response of an agent in a dynamic environment. For believable interactive applications, we need to connect facial expression generation with a process that produces believable behaviors, given the inputs to the system (i.e. the actions of the user). There has been considerable work in the Intelligent Agents' community (with contributions from psychologists) on devising computational models of emotions [1]. Some models have demonstrated how to map events into emotional responses based on how they affect the agent's

*Correspondence may be addressed to: Magy Seif El-Nasr, 301 H.R. Bright Building, College Station, TX 77845, (409)845-1298, magys@cs.tamu.edu

goals [15], while other models have focused more on internal motivational states, such as pain, thirst, and tiredness or fatigue [2]. It has also been shown that learning is a critical component for generating believable behaviors in a dynamic environment because agents are expected to notice patterns of actions that tend to occur in sequences, to make associations with good or bad events, to become conditioned to repeated actions, and in general to adapt to their environment. In our recent work, we have developed a new computational model of emotions in intelligent agents that can be used to generate believable behavior in a dynamic environment using several learning algorithms [5, 6].

In this paper, we integrate our model of emotional response with a model for facial expressions to produce a simulated agent that is capable of both generating and expressing believable emotions in real time. For demonstration purposes, our system is centered around simulating interactions with a human baby, since babies have fairly well-known and recognizable emotional behaviors. As a sequence of actions on the baby is selected by the user, the program uses the emotional model to update the emotional state of the baby and generate behaviors consistent with its cumulative experiences (e.g. Sadness when a toy is taken away, which gradually changes to Happiness when it is hugged). The outputs of the emotional model are continuously supplied to a module that renders the appropriate expression by manipulating a 3D model of the baby's face. This involves mapping the levels of the simulated emotional state into the appropriate control units governing facial expression.

This work parallels the process that computer animators use in developing characters for a series or for a story. A great deal of effort is put into developing protocols for a character's typical emotional response to situations, and then designing the typical actions and appearance of the character while experiencing these emotions. By coupling a tunable algorithm for encoding emotional response with a tunable model of facial expression, we are taking the first step in developing a very powerful character design methodology. This could potentially be developed into an intelligent tool capable of capturing the intent of the character designer, and allowing the character's personality to be expressed as the designer intended, even in complex interactive situations such as in a video game.

2 Modeling Emotional Dynamics

2.1 Believable Agents

Intelligent agents are software programs that exhibit autonomous, goal-oriented, adaptive behavior [9]. Generally, agents receive inputs from the external environment and take actions to achieve their goals. Intelligent agents are particularly suitable for building interactive systems because they can be designed to make their own decisions, without intervention from a user, based on an internal representation of their goals and the state of the world. Some agents have additional characteristics, such as the ability to collaborate with other agents. A particularly important capability for systems that interact with humans is *believability*. Creating the illusion that an agent is actually intelligent or alive is useful for a wide range of applications, from computer games to public information kiosks to computer-aided training systems [20, 7].

Emotions are a very important component in enhancing the believability of intelligent agents [1]. To create a believable agent, the agent needs to be able to simulate emotional responses to its environment. For example, character modeling studies in theater often focus on developing ways of acting that are effective at conveying emotions [22]. Thus, to be truly convincing, the agent must have the capability of generating and expressing emotions. Many studies have been done on visualizing static expressions, which are sufficient for many applications in the entertainment industry, such as movies. However, emotional behavior in these applications is often hard-wired to be triggered at a specific time in a script. In contrast, interactive systems (such as games with artificially-generated characters) require immediate responses to a dynamically changing environment, and some situations may combine conditions in ways that were not entirely anticipated beforehand. In such applications, for the intelligent agent to simulate having believable emotions, it needs to have algorithms for generating them dynamically.

2.2 Emotions in Agents

A number of computational models of emotions have been proposed, especially based on psychological research. Most models treat emotions as internal states of an agent which can be activated to varying degrees. Event-appraisal models [15, 21] study the effect of events on emotions by evaluating how they impact the agent's goals. Other models link emotions to motivational states, such as pain or hunger [2]. Still other models have examined the connection between emotions and physiological conditions, which is useful for various applications such as affective wearable computers [19], or studied emotions from the perspective of neurobiology [3].

Within the Intelligent Agents community, these models of emotions have been implemented in various agent-based systems. One of the best-known systems that includes emotional agents is the OZ project developed at CMU [1, 20]. The OZ project simulated believable emotional and social agents; each agent initially has some predetermined goals, different strategies to achieve each goal, and attitudes towards certain objects in the environment. When the agent perceives an event in the environment, it is then evaluated according to the agent's goals, standards and attitudes. After the event is evaluated, special rules are used to produce an emotion with a specific intensity. These rules are based on Ortony et al.'s event-appraisal model [15]. The emotions triggered are then mapped, according to their intensity, to a specific behavior. This behavior is then expressed using text or animation [20]. One of the limitations of OZ is that it treats the outward behaviors in a crisp fashion, such that, with sufficient activation, the agent either becomes happy or sad, without intermediate states. Velasquez [24] has also implemented an agent that simulates emotions, based on the event-appraisal model by Roseman et al. [21], in a robot named Yuppy.

One aspect that is not frequently addressed in these computational models of emotions is how to adapt to an environment dynamically. For example, in both OZ and Yuppy, the degree of the impact of events on goals is fixed and does not change with the user or the environment. While it is often possible to determine how an agent should behave in a static situation, realistic interactive agents must be able to adapt to their environment, which requires learning. If an event is repeated several times, emotional responses in real agents tend to decrease in intensity (due to desensitization or attenuation). If several events typically happen in sequence, real agents will eventually develop expectations for what is about to happen next. And if certain signals in the environment frequently co-occur with certain emotions, real agents would be expected to become conditioned to these patterns. To generate a believable behavior for an interactive agent, the model of emotions must take these types of adaptation into account. In our previous work, we have developed a model of emotions called FLAME, for Fuzzy Logic Adaptive Model of Emotions [5]. FLAME uses fuzzy logic to represent the relationships among events, goals, and emotions in a flexible way that helps produces smooth transitions in behavior. FLAME also incorporates several learning algorithms that it uses for the agent to adapt to its environment. In this section, we briefly describe the architecture and major components of the FLAME model.

2.3 The FLAME Model of Emotions

The FLAME model consists of three major components: an emotional component, a decision-making component, and a learning component. Figure 1 shows the architecture of the model. The agent first perceives various events in the environment, as depicted from the right hand side of the diagram. These perceptions are then passed to both the emotional component and the learning component (on the left-hand side of the diagram). The emotional component assesses the relevance of each event and then computes new levels for each of the internal emotional variables. The as-

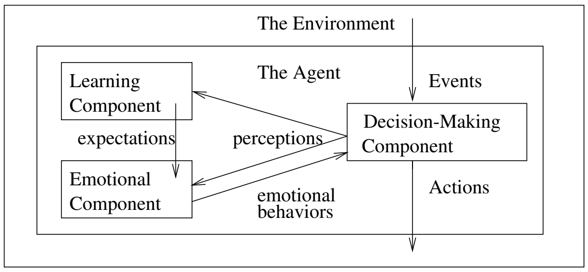

Figure 1. FLAME System Architecture.

sessment of events depends on expectations and other things learned from prior experience. Therefore, the learning component keeps track of associations and patterns in the history of events and provides useful inputs to the emotional process. Finally, the emotional levels are filtered to produce a coherent mixture, and an appropriate emotional behavior is selected and returned to the decision-making component to influence what specific actions are taken.

The emotional component is detailed in Figure 2. The evaluation of an event perceived in the environment is carried out in two steps. First, the experience model determines which goals are affected by the event and the degree of its impact. Second, a desirability level of the event is computed according to the measure calculated by the first step, along with the importance of the goals involved. In other words, the event desirability measure is calculated from two major factors: the impact of the event on the agent's goals and the importance of these goals. Fuzzy rules are used to calculate the desirability measure given these two values as shown in the rules below:

```
IF Impact(G1,E) is A1
AND Impact(G2,E) is A2
  ...
AND Impact(Gk,E) is Ak
AND Importance(G1) is B1
AND Importance(G2) is B2
  ...
AND Importance(Gk) is Bk
THEN Desirability(E) is D
```

where k is the number of goals involved. This rule reads as follows: if the goal G_1 is affected by an event E to a degree A_1 (where A_1 is a fuzzy descriptor like $HighImpact$, $MediumImpact$, or $LowImpact$, for which we have defined triangular membership functions), and the importance of the goal G_1 is B_1 (using fuzzy membership functions like $HighImportance$, $MediumImportance$, and

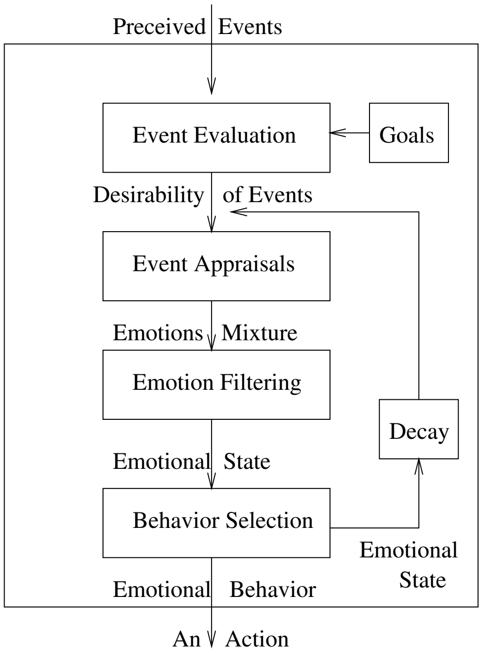

Figure 2. The Emotional Component.

$LowImportance$), and the other goals of the agent are taken into similar consideration, then the desirability of the event E will be of degree D. The various antecedents may be combined using the *SupMin* rule, in which the satisfaction of the rule is determined by the minimum over all the antecedents, and strength of a conclusion is determined by taking the maximum over all relevant rules [26].

The calculated desirability measure will be passed to an event appraisal model to determine the emotional state of the agent. The event appraisal model used in FLAME is based on Ortony et al.'s model [15]. The appraisal model is essentially a set of rules that produces emotions according to the agent's expectations and the desirability of expected or perceived events. To illustrate how the appraisal model works, we use Joy as an example. Joy is an emotion that depends on the desirability of a perceived event, and thus the model will produce the emotion of Joy if the desirability of a particular event is positive. The intensity of the emotion will depend on the desirability level and the degree of expectation of that particular event (if the event is unconfirmed). Another example is Hope, which is defined as the occurrence of an unconfirmed desirable event. Thus, to trigger an emotion, we need both the event desirability measure and the expectation measure.

The mixture of triggered emotions will then be filtered to produce a coherent emotional state. The filtering process uses a variation of Bolles and Fanslow's [2] model to produce a consistent set of emotional intensities that co-exist. Filtering the triggered emotions involves a variety of complex suppression and enhancement relationships involving motivational states. For example, sadness tends to override happiness when the intensities are similar, and pain tends to (temporarily) inhibit the degree of fear [2]. The filtered emotional state is used to select a consistent emotional behavior, such as laughing or crying. These are determined according to the emotional state and the current situation using fuzzy rules. The selected emotional behavior is passed to the decision-making component to further determine what specific actions to take. In FLAME, we also include feedback mechanisms to model time-dependent effects (such as moods), and we decay emotional intensities over time to improve the realism of the simulation. For more details on the emotional component of the FLAME model, see [5].

2.4 Learning Algorithms in FLAME

Learning and memory have a major impact on emotional behavior [11]. In FLAME, we incorporated algorithms for learning several different types of things about the environment, including the likelihood of certain events to occur in any given situation, and the desirability of various actions.

Forming a User Model At any given time, an agent will need to know what to expect, i.e. how certain it is that a particular event will happen. Many emotions depend on whether an event was expected (as in the case of relief) or unexpected (as in surprise). In many environments, the frequency of occurrence of certain events is not always known *a priori* and must be determined on-line. This is especially true of the user's tendencies, which must be learned dynamically since reactions may vary from person to person. Thus the agent needs to be able to anticipate or predict the user's actions based on prior experience. We use a simple probabilistic approach to learn associations among events or actions that are taken in sequence by the user, which we call "patterns." We keep a table of counts of observed sequences of events up to some fixed length. We focused on patterns of length three, since this was adequate to produce believable learning behavior in our application (i.e. a baby). We use these counts to define the probability of an event e_3 to occur, given that events e_1 and e_2 have just occurred, which we notate as $p(e_3 \mid e_1, e_2)$.

When a pattern is first observed, an entry is created in the table that indicates the sequence of three events, with a count of 1: $c[e_1, e_2, e_3] = 1$. Then, every time this sequence is repeated, the count is incremented. We can use these counts to calculate the expected probability of a new event Z happening, given the previous two events X and Y:

$$p(Z \mid X, Y) = \frac{c[X, Y, Z]}{\sum_i c[X, Y, i]}$$

To handle cases where experience is limited, if the total number of observations of the sequence of events X and Y is zero (or too small to be reliable), then we reduce the probability to be conditioned only on one previous event (Y in this case), which can be calculated by summing over more counts in the table. These probabilities allow the agent to determine how likely the user is to take certain actions, given the events that have occurred in recent history, which facilitates the determination of believable emotional responses.

Learning the Desirability of Actions In addition to expectations of events, the agent needs to be able to learn the desirability of actions based on their impact on a set of goals. In realistic applications, a given event often does not have a direct impact on any specific goal, but instead a combination of events may eventually have an impact on a given goal. For example, consider the goal of becoming rich. Getting fired from your job might not affect this goal directly; however, if you are fired, you will no longer receive your salary, and thus it becomes harder to get rich, though not impossible if other actions are taken. Therefore, the link between events and the goals they affect is no longer a simple one-to-one relationship, but rather we need to identify more general relationships between events and goals through a sequence of actions. Identifying this connection has often been considered to be a very difficult task [20]. However, in FLAME, we use reinforcement learning [10] to learn how various events impact the agent's goals, based on experiences accumulated over sequences of events.

In reinforcement learning, the agent represents the problem space using a table of numbers called Q-values which are associated with each state-action pair. The table is initially filled with random values. Suppose the agent begins in a state S. It can take a variety of actions A_i that put it into a new state S_i. The agent may obtain a reward R_i from the environment for this action (such as being told "Good Baby"). After the agent makes the transition, it updates the corresponding Q-value in the table using the following formula [14]:

$$Q(S, A_i) = R_i + \gamma \cdot max_{A_j}\ Q(S_i, A_j)$$

to reflect the immediate reward plus the maximum expected reward achievable from the new state based on its Q-values, discounted by a constant factor γ. After a series of experiences, the values in the table will eventually converge so that the agent is able to make optimal decisions (decide which actions will lead to the highest expected reward in the long run). Variants of these formulas can be used to handle the inherent non-determinism introduced by the human in the interaction [6]. Using this approach, the agent will be able to formulate its expectations of rewards according to the situations it faces. Thus, the agent will be able to associate an event with a particular goal down the path, which is important for performing the impact assessment.

Figure 3. An example three-dimensional face model.

3 Facial Emotion Visualization

3.1 The Concept

The emotional state of the baby can be visually communicated using an expressive computer-based model of the baby's face. The development of expressive visual facial models has progressed over the last 25 years and is now quite well understood [17].

The basic idea is to manipulate a geometric representation of the face so that it mimics the expressive feature postures of the human face. The facial geometry is usually a three-dimensional surface or set of surfaces that conform to the face of interest. Figure 3 shows an image of such a three-dimensional face model. Facial models based on simpler geometry, such as the two-dimensional feature lines used by a cartoonist, are also possible.

The shape of the facial surfaces or lines are then controlled to take on the feature postures which represent the various facial expressions. This control is usually in terms of a set of facial parameters. These parameters may mimic the anatomical actions of the facial muscles [25] or may directly manipulate feature postures using pseudo-muscle [12] or *ad hoc* parameterizations [16].

The control parameterizations used for these facial models do not usually correspond directly to the emotional state of the character. Multiple facial parameters contribute to the visual expressions associated with each emotional or physical state. A mapping is required from emotional parameters into the facial expression control parameters as discussed below.

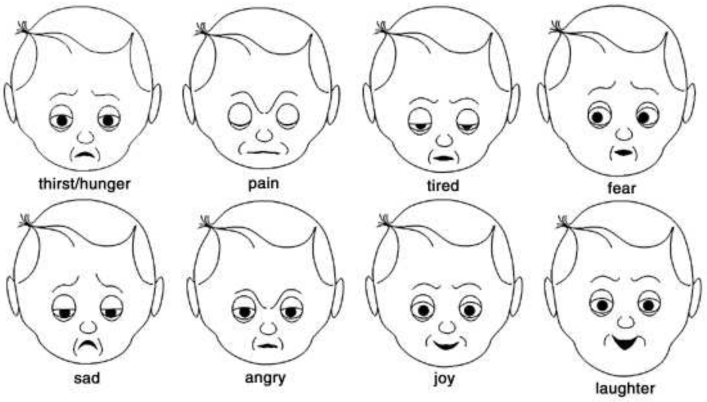

Figure 4. Facial expressions of the eight emotional and physical states created using the simple baby face model.

3.2 The Face Model

As introduced above, computer-based expressive facial models may create various visual representations of faces ranging from simple cartoons to very realistic three-dimensional renderings. For this project a fairly simple two-dimensional line drawn cartoon-style baby face was selected. The expressive capabilities of this style are a good match to the relatively simple emotional model for the baby.

This facial model supports the range of expressions that correspond to the physical and emotional states of the baby. There are specific facial feature postures that correspond to these states [8], as shown in Figure 4 and summarized below. The first five of these – fear, sadness, anger, joy, and laughter - correspond to four of the six "universal" facial expressions [4], laughter being an extreme expression of joy. The last four listed - pain, tiredness, thirst and hunger – are physical rather than emotional states. Pain and tiredness have specific corresponding facial expressions. The distresses of thirst and hunger do not have specific expressions, but are here manifest as a combination of mild fear and sadness – being miserable.

- Hunger and Thirst - eyebrows slightly up, with inner third bent upward; eyelids more open; mouth closed with corners slightly down; face slightly down.
- Pain - eyebrows down, especially inner corner; eyelids tightly closed; mouth wide and open with corners down.
- Tired - eyebrows relaxed or somewhat raised; eyelids drooping; mouth relaxed; face tilted down.
- Fear - eyebrows lifted up and pulled closer together with inner third bent upward; eyelids wide open; mouth opened and widened; face turned slightly to the side.
- Sadness - eyebrows inner third bent upward; eyelids slightly closed; eyes slightly downcast; mouth closed with corners pulled down.
- Anger - inner eyebrows pulled together and down; eyelids open; mouth tightly closed.
- Joy - eyebrows relaxed or slightly up; eyelids open; mouth wide with corners pulled up and back.
- Laughter - similar to joy but with eyelids almost closed and mouth wide open; face tilted up.

3.3 Controlling the Facial Expressions

Facial expression is controlled through the action of sixteen parameters. Five of these control the eyebrows. Four control the eyes. Five control the mouth and two control face orientation. These are summarized in the following table.

Eyebrows	right eyebrow height left eyebrow height right eyebrow shape left eyebrow shape separation of eyebrows
Eyes	right eyelid opening left eyelid opening eye gaze azimuth eye gaze elevation
Mouth	jaw rotation (mouth opening) mouth shape mouth width right upper lip left upper lip
Orientation	face azimuth face elevation

The collective actions of these parameters approximate the expressive actions of the important facial muscles.

3.4 Mapping the Emotions

The emotion model for the baby outputs state information in the form of four physical state parameters and five emotion parameters. Five of these – thirst, hunger, pain, tiredness and laughter – can be either true or false. The other four are real-valued and range between 0 and 10.

A mapping is required to convert these physical and emotional state variables into the expression parameters needed to drive the visual face model. The first step in this process is to normalize the state variables. The binary values are converted into 0.0 for false and 1.0 for true. The continuous values are scaled so that they range between 0.0 to 1.0. This results in a set of nine values in the 0.0 to 1.0 range.

The next step is to establish baseline values for each of the face expression parameters. These correspond to the face having a "neutral" expression. Deviations from these baseline values result in the face portraying various expressions. The expressive parameters for a particular emotional or physical state are obtained by computing adjustments to the baseline values. The general form of these adjustments is

$$parameter(j) = baseline(j) + \sum_i c(i,j) * emotion(i)$$

where $parameter(j)$ is one of the expression parameters, $c(i,j)$ are the coefficient terms for $parameter(j)$ and $emotion(i)$ are the emotional state values. Depending on the coefficient values, each emotional state variable may add to or subtract from each of the baseline parameter values.

4 Results

The models of emotional dynamics and of facial expressions, explained in the Sections above, were incorporated into a program that simulates interaction with a baby and produces appropriate facial expressions dynamically. The program is implemented in Java, C, and OpenGL, and runs on SGI workstations. The interface is divided into two parts; one window controls the user's actions toward the baby. The other window shows the baby's facial expression generated in real-time. The expressions of the baby change according to its emotional state, as governed by the FLAME model discussed above.

The user can interact with the agent through a variety of actions, including playing with the baby, hugging the baby, singing to the baby, comforting the baby, pinching the baby, shouting at the baby, giving the baby objects (e.g. a toy, candy, a drink or vegetables), taking the objects from the baby, and saying certain phrases to the baby (e.g., "Good Baby," "No, don't do that," or "Come Here").

In response to various sequences of actions, the baby can produce complicated mixtures of emotions in this environment. However, the model for facial expressions is flexible enough to generate corresponding expressions covering a broad range of situations. In the remainder of this section, we describe the baby's responses to several scenarios to demonstrate the complexity of emotional behaviors that can be dynamically simulated. Figure 5 provides snapshot views of the behaviors exhibited during the scenarios.

Initial Impressions We began by selecting actions such as playing with the baby or hugging the baby. However, as can be seen from Figure 5a, these actions appeared to irritate the baby. While hugging or playing with a baby should produce Joy, Anger and Fear were produced by the simulation instead (see the emotional vector in Figure 5a). However, our model provides a reasonable explanation for such reactions. We start each simulation by assuming that the baby does not know the user. This causes the baby to initially fear and dislike the unknown user. Since the baby initially does not know what to expect from the user, the emotional state is dominated by Fear and Anger, rather than Happiness. However, as we will see, these initial emotions are gradually suppressed as the user interacts with the baby.

Motivational and Emotional States Interaction Next, we selected actions to give the baby some candy and a toy, to reassure the baby that the user is a nice person. Happiness, as is shown in Figure 5b, prevails. However, there is still some Anger in the baby's emotional state, which was carried over from the previous situation, but with decaying intensity (eventually disappears over time). The introduction of the new objects to the baby, which were assumed to be initially unfamiliar, caused the Fear to increase.

In contrast to decaying emotional intensities, Hunger and Thirst build up over time and eventually come to take over the emotional state of the baby. In order to show this effect, we waited for some time until we observed that the baby was very thirsty. At that time, as can be seen in Figure 5c, Anger, Fear and Happiness were all diminished and overshadowed by Thirst. Even though the emotion vector in Figure 5c shows that the baby still has some degree of Fear and Anger in its emotional state, the facial expression of the baby shows nothing but Misery. This observation illustrates one of the most important factors in the model, namely the interaction between emotional and motivational states. As is discussed in the description of the FLAME model, sometimes emotions or motivational states are inhibited or reinforced by other emotions or motivational states. Thirst, in this case, was inhibiting Fear and Anger. Since Thirst and Hunger are both expressed by a miserable appearance, Misery was the dominant feature shown in the facial expression of the baby at the time.

Rewards and Punishment Now that the baby can trust the user somewhat, the user can take actions to condition the baby for certain responses. For example, we took away the candy from the baby. As a result the baby became very angry, as shown by the facial expression in Figure 5d. Then we gave the baby some vegetables, followed by pinching the baby. The baby formed a connection among *take(Candy)*, *give(vegetables)*, *pinch(Baby)*, and identified it as a negative pattern. However, after repeatedly giving the baby vegeta-

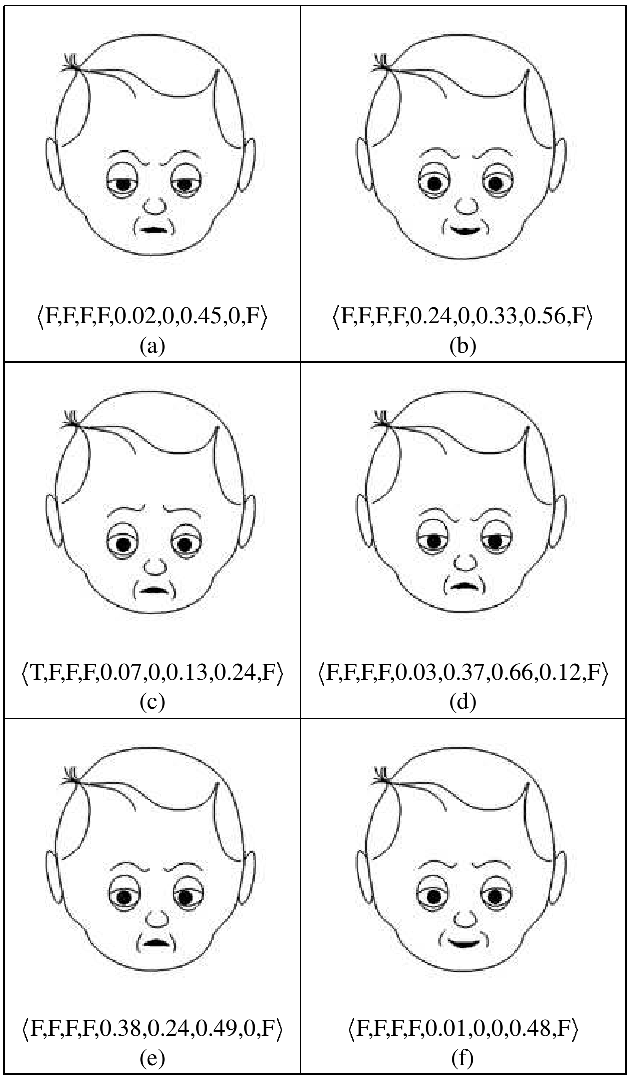

Figure 5. Results of the simulation. Facial expressions of the baby generated for scenarios producing various mixtures of emotions. For each situation (described in the text), the emotional vector output by the FLAME model is given, along with a depiction of the baby's face with the corresponding expression, generated by the facial expression model. The components of the emotion vector consist of the following binary and continuous attributes: *Thirst*, *Hunger*, *Pain*, *Tiredness*, *Fear*, *Sadness*, *Anger*, *Joy*, *Laughter*. **The first four components, along with the last, are binary states with values of True (*T*) or False (*F*). The other components have continuous values normalized to range between 0 and 1.**

bles and pinching the baby, the baby learned to specifically associate vegetables with getting pinched, using the reinforcement learning algorithm as described above. Eventually, when we gave the baby vegetables, the baby's fear increased in anticipation of being pinched, as shown in Figure 5e.

We also tried to positively reinforce certain patterns and observe the baby's reaction. We hugged the baby, sang to it and then told it "Good Baby." We repeated this pattern a number of times. Then we gave the baby a toy and hugged it. There are two important reactions that can be pointed out in the results, shown in Figure 5f. The baby acquires a positive attitude toward singing because, over previous experiences, the baby has come to associate singing with the reward of being told "Good Baby." Also, the baby exhibits learned expectations. The baby now anticipates that singing will follow after it has been hugged, since the pattern of *hug(baby) - SingTo(Baby) - Tell("Good Baby")* had been reinforced many times. Thus the hugging has a suppressing effect on Anger and Sadness from earlier scenarios, and the baby now has an emotional state in which Joy prevails.

Summary The results from these simulations demonstrate that: 1) the baby can generate believable emotional responses using the FLAME model, 2) the learning algorithms in the model allow the baby to dynamically adapt to the user, and 3) the facial expression model is capable of generating plausible visualizations for emotional states involving complex mixtures of emotions.

5 Discussion

Our efforts to combine a model of emotions with a facial model, by providing a mapping from complex emotional states onto facial expressions, is a first step along the way towards developing the technology of a truly believable interactive agent. Our aim is to develop an agent that responds in a believable way to a complex series of events with a memory of what has transpired, and an emotional state that reflects that experience. Thus, not only must the actions of the agent be appropriate to its experiential context, but its attitude should be appropriate for the same context as well. The techniques demonstrated here also have potential for use outside of the field of intelligent agents. For example, they might be highly useful in the development of computer-assisted tools for character design for games and animation.

Much effort in computer graphics has been expended in developing tools to assist animators. These have ranged from tools to assist in traditional key-framing and ink-and-paint, to the emergence of the new genre of "three-dimensional" computer animation. However, very little work has been done in the very important area of character design. In traditional animation, animators spend a great deal of time developing both the look and the personality of each major character. This involves the production of hundreds of drawings showing the character in numerous situations, working out the details of how the character's personality would play out in its response to these situations. Painstaking attention is paid to how situations would affect the character's emotional state, and how the character would react in terms of facial expression, body language and overt actions [22]. For instance, compare the mild manner and philosophical reactions of Mickey Mouse to the hair-trigger highly emotional reactions of Donald Duck. The system that we are developing, which allows the design and testing of both emotional responses and related physical appearance in situations taking place over time, could be developed into an excellent tool for character design.

For the same reasons that such a tool would have application in the development and design of agents for computer interfaces, it would be especially good for designing animated characters who would be expected to respond to events occurring in interactive games. Not so directly, such a tool would be an excellent design aid for developing characters to appear in animated films. It would be particularly well suited to the development of characters who would appear as extras, or for example in crowd scenes, where appropriate actions are needed but it is not necessary for the character to receive the full attention of an animator in developing its motion. Similarly, such a system would be very useful in designing characters for animated television series, where there is a need to develop scenes quickly to meet a high-volume, short-deadline schedule, and where it is acceptable to trade animation quality for production efficiency.

There have been a number of other attempts to automate the animation of facial (and other) expressions, such the *Improv* project by Perlin [18]. These systems often provide the capability of generating a range of expressions or behaviors, but do not specify a controller for dynamically generating believable responses to events in real-time (other than through simple scripts and triggers). The work reported in this paper is more similar to the work of Velasquez [23]. Velasquez describes the simulation of an interactive baby, called Simón, with various facial expressions. The implementation of Simón is based on Roseman et al.'s [21] event-appraisal model, which is used to generate emotions and then select a corresponding face image to display. However, we extend Velasquez's approach by replacing his single-mode emotional response with one that allows for a multi-dimensional complex of physical and emotional states. This means that an agent can produce a much wider variety of emotions and can do this in ways that vary from subtle to overt. Since we map components of this complex state onto multiple facial controls, our faces should also have a much wider range of expression than those used by Valasquez. Also, Simón has static, pre-determined responses and cannot adapt its behavior to different users or environments, which was the motivation behind incorporating learning algorithms into the FLAME model.

We are anxious to extend this work, and there are several key areas that could be readily improved. Most simply, the facial model could be extended in a straightforward way to use three-dimensional technology. Clearly there is also room to attempt to model the neural control mechanisms that govern facial expression as a function of emotional state, and to use these mechanisms to control a more physically and physiologically realistic facial muscle model. Likewise, the emotional model should be extended to exhibit the sophisticated emotional structure of an adult. This would no doubt require the modeling of the more complex cognitive processes of an adult, stemming not only from physiology and innate responses but also from life experiences, culture, education, and other aspects that contribute to a fully developed personality.

6 Conclusion

The approach described in this paper, of combining a model for emotional dynamics with a model for facial expressions, yields a very believable interactive simulation. We believe that this approach can be applied to simulating other characters besides babies. Our results suggest that believable agents require both of these underlying components: a adaptive computational model for simulating emotional states and responses to events, and a computational model for mapping the resulting emotional states onto controls for rendering facial expressions from 3D models. The recent independent maturation of technologies for each of these tasks signals a new era in building interactive applications, where these algorithms may be effectively integrated to produce believable simulations of agents in real time.

References

[1] J. Bates. The role of emotion in believable agents. *Communications of ACM*, 37(7):122–125, 1992.

[2] R. Bolles and M. Fanselow. A perceptual defensive recuperative model of fear and pain. *Behavioral and Brain Sciences*, 3:291–301, 1980.

[3] A. Damasio. *Descartes' Error: Emotion, Reason, and the Human Brain*. New York: Putnam, 1994.

[4] P. Ekman. The argument and evidence about universals in facial expressions of emotion. In H. Wagner and A. Monstead, editors, *Handbook of Social Psychophysiology*, pages 143–146. Wiley, 1989.

[5] M. El-Nasr. Modeling emotion dynamics in intelligent agents. Master's thesis, Department of Computer Science, Texas A&M University, College Station, TX, 1998.

[6] M. El-Nasr, T. Ioerger, and J. Yen. PETEEI: A pet with evolving emotional intelligence. In *Third International Conference on Autonomous Agents*, 1999.

[7] C. Elliot and J. Brzezinki. Autonomous agents as synthetic character. *AI Magazine*, 19(2):13–30, 1998.

[8] G. Faigin. *The Artist's Complete Guide to Facial Expression*. Watson-Guptill, New York, 1990.

[9] N. Jennings, K. Sycara, and M. Wooldridge. A roadmap of agent research and development. *Journal of Autonomous Agents and Multi-Agent Systems*, 1:275–306, 1998.

[10] L. Kaelbling, M. Littman, and A. Moore. Reinforcement learning: A survey. *Journal of Artificial Intelligence Research*, 4:237–285, 1996.

[11] J. LeDoux. *The Emotional Brain*. New York: Simon & Schuster, 1996.

[12] N. Magnenat Thalmann, N. Primeau, and D. Thalmann. Abstract muscle action procedures for human face animation. *Visual Computer*, 3(5):290–297, 1988.

[13] N. Magnenat Thalmann and D. Thalmann. Computer animation. In A. Tucker, editor, *The Computer Science and Engineering Handbook*, pages 1300–1318. CRC Press, 1996.

[14] T. Mitchell. *Machine Learning*. New York: McGraw-Hill, 1997.

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

[16] F. Parke. A parametric model for human faces. Technical Report UTEC-CSc-75-047, University of Utah, Salt Lake City, 1975.

[17] F. Parke and K. Waters. *Computer Facial Animation*. A.K. Peters: Wellesley, Mass., 1996.

[18] K. Perlin and A. Goldberg. Improv: A system for scripting interactive actors in virtual worlds. In *Proceedings of SIGGRAPH '96*, 1996.

[19] R. Picard. *Affective Computing*. Cambridge: MIT press, 1997.

[20] W. Reilly. *Believable social and emotional agents*. PhD thesis, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, 1996.

[21] I. Roseman, J. P.E., and M. Spindel. Appraisals of emotion-eliciting events: Testing a theory of discrete emotions. *Journal of Personality and Social Psychology*, 59(5):899–915, 1990.

[22] F. Thomas and O. Johnston. *The Illusion of Life*. New York: Abbeville Press, 1981.

[23] J. Velasquez. Modeling emotions and other motivations in synthetic agents. In *Proceedings of the Fourteenth National Conference on Artificial Intelligence*, pages 10–15, 1997.

[24] J. Velasquez. When robots weep: Emotional memories and decision-making. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence*, pages 70–75, 1998.

[25] K. Waters. A muscle model for animating three-dimensional facial expressions. *Computer Graphics*, 21(4):17–24, 1987.

[26] J. Yen and R. Langari. *Fuzzy Logic: Intelligence, Control, and Information*. Upper Saddle River, NJ: Prentice Hall, 1998.