A Web of Emotions

Magy Seif El-Nasr Thomas R. Ioerger John Yen
Texas A&M University
Computer Science Department
College Station, TX 77844-3112
{magys, ioerger, yen}@cs.tamu.edu

ABSTRACT

The fast growing technology of agents and the rise of emotional intelligence opened the doors to many new and interesting ideas. Among these ideas, one that has captured the interest of many researchers within the AI community is the development of emotional agents. While research on emotional agents is in a very early stage, many models have been developed to simulate emotions in agents. Some of the most recent computational models of emotions attempt to account for the inevitable links between emotions and many other processes residing in the mind, such as planning, cognition, personality, and motivation. Most of these models, however, simulate an incomplete picture of the links involved. Furthermore, most of these models tend to simulate emotional triggered expressions or behaviors according to one emotion, which is an invalid assumption, since emotions, realistically, occur in a mixture. We have taken a relatively new approach in emotional agents' development. Rather than defining the emotional process as a single independent process, we define the emotional process as a web whose different emotional states act as strings that pass through many other processes or states affecting or becoming affected by them. Thus, instead of triggering emotions and then defining a behavior according to one emotion, we use other processes, such as learning and experience, to trigger mixtures of emotions. These mixtures then propagate through the web affecting or becoming affected by other processes and states, including motivations and behaviors. In this paper, we describe the design and implementation of the emotional web. Additionally, we demonstrate empirically through a computer simulation of a pet that simulating the link between emotions and learning is crucial to users' assessments of the believability of an agent's behaviors.

1. INTRODUCTION

Emotions have long been recognized to be an important aspect of the human mind. However, the role that emotions play in our thinking and behaviors has often been misunderstood. Many theories and hypothesis have recently emerged to define and explain the role that emotions play in human intelligence and social abilities [5, 17, 18].

As a result, many researchers within the AI field started developing computational models of emotions. Computational models of emotions are very useful and essential to many applications. For instance, emotions and emotional expressions were identified as necessary substrates for producing believable characters [46], and thus many computational models of emotions were employed by researchers within the agents' community to develop believable agents [3, 36, 35]. Believable agents have been used in a wide variety of applications including video games [16], educational software [25, 9], and training environments [38]. Additionally, the role of emotions in communication has long been acknowledged [18]. A computational model of emotions may be used to develop software agents that can effectively communicate and collaborate among themselves, and thus the emergence of 'Social Agents [6, 20].' Social agents are being developed with emotional and social capabilities to enable them to communicate and collaborate on several tasks [6].

As it can be seen from the discussion above, many computational models of emotions are being developed for a wide variety of applications. There is no doubt, however, that these models are still in their early stages. The emotional process is a very complex process that is not well understood. The complexity of emotions derives from their ability to affect and be affected by many processes within the mind [14]. A wide variety of evidences have shown that emotions lead a major role in thinking and recall [4], learning [24, 23], decision-making [5], motivation [19] and judgments [14, 15]. Thus, in order to simulate a realistic picture of the emotional process, one will have to simulate the links between the emotional process and other processes or states including learning, decision-making, motivations, etc. With the exception of a few recent models, most of the previous computational models of emotion do not account for the links between emotions and other processes, and thus most of them develop the emotional process as a single confined process whereby emotions are triggered independently. In the next section, we will attempt to describe some of the most recent models and outline their approach in simulating the emotional process.

We are presenting a relatively new approach to modeling emotions. Instead of triggering emotions and simulating their effects on behaviors independently, we use the learning process to activate mixtures of emotions that affect and become affected by a number of processes within the cognitive system including motivations and behaviors. In this way we simulate a rather realistic picture of the emotional process, which can be used to enhance applications where a computational model of emotions may be used. We hypothesize that simulating a realistic picture of such a dependant structure of emotions will provide a better illusion of the emotional process, and thus may achieve better believable agents. Additionally, the model may potentially provide better functionality if used in applications, such as decision-making models, that rely on a realistic model of emotions.

Thus, in an effort to build a significantly more believable and realistic emotional experience, we present our model. We will first present the architecture of the system, which simulates the link between emotions, learning, motivations, and behaviors. In order to evaluate the approach presented in the paper, we developed a prototype agent, a pet named PETEEI (a PET with Evolving Emotional Intelligence). We devoted a section in this paper to describe PETEEI and illustrate some observations that emerged from the simulation. We will further explore the results of an evaluation that was performed using 21 subjects to judge the believability of the agent simulated.

2. PREVIOUS MODELS

Many computational models of emotions have been proposed and used in many applications. Although some of the computational models of emotions attempt to address links between emotions and other processes, their simulation tend to be limited in many ways. In the next several paragraphs, we will detail some of the most prominent models in the field and discuss the differences between their approaches and ours.

An important effort was made by the OZ project [3, 36]. The OZ project simulates believable emotional and social agents; each agent initially has preset attitudes towards certain objects in the environment. Furthermore, each agent has some initial goals and a set of strategies that it can follow to achieve each goal. The agent perceives an event in the environment. This event is then evaluated according to the agent's preset goals, standards and attitudes. After an 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 [32]. These emotions are then used as preconditions to the planning process by which an action is triggered [36].

Although the OZ project seems to explore the link between emotions and planning, the planning process was simulated as a very simple selection process. Thus, the influence of emotions on the planning process, which generally tends to include higher level assessments, was not extensively explored. Furthermore, the emotion activation phase was simulated as an independent process where emotions were triggered by the occurrence of events. One way to expand the model is to allow for a dynamic representation of standards, goals, attitudes and expectations. In this way the model can be expanded to handle the dependency between learning or experience and emotions. We will further address this issue in the next section.

Elliot [8] developed another multi-agent computational model, called the *Affective Reasoner*. The model is an extension of Ortony et al.'s psychological model [32]. The *Affective Reasoner* project involved agents that are capable of producing twenty-four different emotions, including joy, happy-for, gloating, resentment, and can generate many different emotional expressions. Each agent included a representation of the self (agent's identity) and the other (identity of other agents involved in the situation). During the simulation, agents judge events according to their pleasantness and status (unconfirmed, confirmed, disconfirmed). Joy, for example, is triggered if a confirmed desirable event occurs. Additionally, agents account for the other agents' responsibility of the occurring events. For example, gratitude towards agent *x*, can be triggered in agent *y*'s emotional state if agent *y*'s goal was achieved by the agent *x*. In addition to the emotion generation and action selection phases, the model presents another dimension to the emotional problem, which is social interaction. During the simulation, agents, using their own knowledge of emotions and actions, can infer the other agents' emotional states from the situation, emotional expressions, and actions. These inferences can potentially enhance the social interaction process [8].

While Elliot's model presents an interesting simulation describing emotion generation, emotional expressions, and their use in interactions, the model still faces some difficulties. The model does not address several issues, including conflicting emotion resolution in case of emotional mixtures, and the impact of learning on emotions and expectations [8]. The model presents a good starting point for addressing the link between emotions and social interaction. However, it does not address other extensions of the emotional process, such as the effect of emotions on motivations or learning.

A model that extensively explores the links between emotions and motivations was presented by Velasquez [47] to simulate emotions using a multi-agent architecture. The model captures several aspects of the human mind. It captures the neurophysiology aspect, which simulates effects of neurotransmitters, brain temperatures, etc. It also captures the sensorimotor aspect, which simulates facial expressions, body gestures, postures and muscle action potentials. Additionally, the model presents a simulation of motivational states and emotional states using event appraisals, interpretation of events, comparisons, attributions, beliefs, desires, and memory.

Clearly, the model presents and simulates many aspects of the different processes within the brain. It further accounts for the link between motivations and emotions. For example, some emotional states, such as fear or sadness, may suppress some drives or motivations such as thirst or hunger. Our model adds yet another dimension to the Velasquez's model by extending the independent nature of the emotional triggering process. We have attempted to add a dependency model that simulates the link between the emotion activation process and the learning or experience processes. We will discuss this extension in more detail in the next section.

As it can be seen, most of models discussed above attempt to simulate links between emotions and motivation or social interaction or planning. However, if we look more closely at some of these models, we find that the emotional activation process was simulated as an independent process whereby an emotion or some emotions are triggered according to some rules. Furthermore, the link between emotions and the behavior selection process was incompletely explored. Generally, a behavior is selected according to one emotion which is typically singled out by a function, such as *max(e)* (returns the emotion with the maximum intensity).

Our approach attempts to expand on the many ideas presented in these models. We treat the emotional process as a web of emotions that spreads its influence to many processes within the mind. Emotions are activated through an emotional activation phase that subsequently interacts with the learning process to trigger a mixture of emotions with different intensities. Once a mixture of emotions is triggered, emotions go through different parallel paths affecting many states and processes, including but are not limited to, motivational states, facial expressions and actions. We, additionally, explore the link between emotions and behaviors. Instead of triggering a behavior according to a single emotion, we blend behaviors by selecting a behavior according to the mixture of emotions triggered through the emotional activation phase.

3. PROPOSED MODEL

We developed a model that simulates the emotional process as a web-like structure which interacts with different processes and states, including learning, motivations and behaviors. The model is illustrated in Figure 1. The figure shows the links (shown in gray arrows) between the emotional process on the one hand and learning, motivations and behaviors on the other.

Emotions are first activated using two main factors: the events' desirability and events' expectations. These factors represent the main link between the emotional process and learning as shown in the figure. As an outcome of this activation phase, a mixture of emotions is triggered. These emotions are triggered with different intensities. Emotional intensities are calculated as a function of the desirability degree of events, expectations, standards and values.

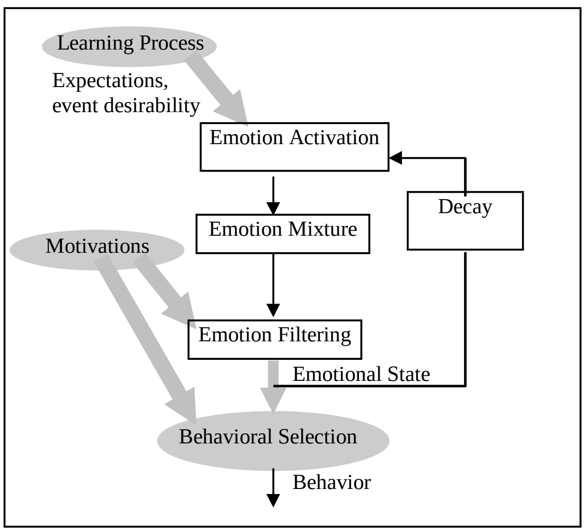

Figure 1. The Emotional Process

Once a mixture of emotions is triggered, these emotions then go through a filtering process where the motivational states and emotional states are examined. As a result, some emotions are suppressed and some others are enhanced. The resultant emotional state is then passed to the behavioral process which in turn interacts with motivational states to select an action. The emotional state is then decayed linearly using a constant *t*. It should be noted that the choice of a linear decrement as a decay function was selected because of its simplicity and the appropriateness of its performance when tested in the prototype.

In the next subsections, we will detail the different phases that the emotional process passes through. Through these discussions, we will illustrate the points of interaction between the emotional process and the other processes and states shown in the figure.

3.1 Emotion Activation Phase

We trigger emotions according to some predefined rules. These rules were defined using an event-appraisal's model developed by Ortony et al. [32]. The rules define emotions in terms of events' desirability and expectations. For example, hope was defined as an emotion that is triggered when the subject expects a certain desirable event to occur. The intensity of hope will depend on the likelihood degree of the event occurrence and the desirability of the event expected. Thus, as it can be seen from this example calculating desirability degrees and expectations represent a very important step in triggering emotions such as hope, fear, joy, and distress. However, it should also be noted that expectations and desirability of events are both dynamic entities that change through experience and time. Therefore, we have employed different inductive learning algorithms to emulate their dynamic nature.

Desirability measures: Learning about events

To trigger an emotion, such as joy (defined as the occurrence of a desirable event), the agent needs a mechanism by which it can assess whether a particular event is desirable and to estimate the degree of its desirability. To achieve such a mechanism one can develop an algorithm that estimates an event's impact to a specific goal. However, it is often the case that a given event does not have any impact on any specific goal directly, but some sequence of events may eventually have an impact on some goals. Thus, in order to measure the desirability of a specific event, we will need to identify the link between an event or a sequence of events and the corresponding goals. Moreover, an agent does not normally know the consequences of a given event until a complete sequence of events is finished. Therefore, the agent faces the problem of temporal credit assignment (defined as determining which of the event in its sequence are responsible for producing the eventual rewards). Fortunately, the agent can use a learning algorithm such as reinforcement learning [21, 30] to learn the association between rewards and events. Additionally, it can also learn to relate sequences to events to a specific reward or rewards. We will briefly outline one of the reinforcement algorithms, namely the Q-learning algorithm. The reader is referred back to [21, 30] for more details.

To illustrate the solution that Q-learning offers to this problem, we will look at Q-learning in more detail. The agent represents the problem space using a table of Q-values in which each entry corresponds to a state-action pair. The table can be initially filled with default values. The agent will begin from a state *s*. He will take an action, *a*, which takes him to a new state s'. The agent may obtain a reward, *r*, for his action [30]. A reward in our simulation is defined as a goal achievement. Thus, if a certain action causes the success of a goal, then the agent will receive a reward, *r*, which is equivalent to the value or the importance of the goal. When the agent receives such a reward it updates the table above using the following formula:

$$Q(s,a) \leftarrow r + \gamma \max_{a'} Q(s',a') \quad [30]$$

where *r* is the immediate reward, γ is a discount factor, s' is the new state, and a' is an action from the new state s'. As the formula shows, the Q-value of the previous state-action pair depends on the Q-value of the new state-action pair.

This algorithm is guaranteed to converge only for deterministic Markov Decision Processes (MDP)[1]. However, in our application there is a one-to-one alternation between the agent and the user. It is almost impossible to predict the reward given a state and an

[1] A Markov Decision Process refers to an environment in which rewards and transition probabilities for each action in each state depend only on that state, and are independent of the previous actions and states.

action, because it may change according to the user and environment. For example, at a state s_0 the agent may take an action x and as a consequence the user gives it a positive reward. At a later time step, the agent is in the same state s_0 so it takes action x again but this time the user rewards it negatively. Therefore, the user introduces a non-deterministic response to the agent's actions. In this case, we treated the reward as a probability distribution over outcomes based on the state-action pair. Figure 2 illustrates this process in more detail. The agent starts off in state s_0 and takes an action a_0 then the user can take either action a_1 that puts the agent in state s_1, or action a_2 that puts the agent in state s_2. The dotted lines show the nondeterminism induced by user's actions, while the straight black lines show the state-action transition as represented in the table. We use another learning model, which will be detailed in the next section, to determine the probability of user's actions.

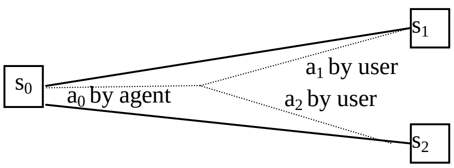

Figure 2. Nondeterministic Rewards

A new formula was formulated to calculate the Q-value as follows:

$$Q_n(s,a) \leftarrow E[r(s,a)] + \gamma \sum_{s'} P(s'|s,a) \max_{a'} Q(s',a')$$

where $P(s'|s,a)$ is the frequency count that was described above, and $E[r(s,a)]$ is the expected probability of the [30].

Using this type of learning the agent will be able to calculate the desirability of events using the Q-values of the state-action pairs. Therefore, if the agent is at situation s and did an action a, it will be able to look up the Q-table to identify the value of the new state s', which would define the desirability of the event or situation s'. Using this paradigm the agent will also be able to define, with some confidence, what events may occur in the future given the events that occurred, which defines the expectation measures that the agent needs to trigger emotions.

However, it should be noted that the agent is most likely to be interacting with a user. Since the agent is interacting with a user, the agent may further refine its expectations of user related events by using another learning mechanism which learns user patterns.

Expectation measures: Learning user patterns

To learn about user's patterns of behavior a probabilistic approach is used to learn patterns based on the frequency with which an action, a_1, is observed to occur given that previous actions a_2, etc. have occurred. We limited the length of patterns learnt to model three consecutive actions, due to several reasons. Since most actions are small, quick, and specific such as *throw(ball)*, a sequence of two actions pattern would not capture a meaningful pattern. Therefore, in order to capture a meaningful pattern, we set the length of the pattern to a lower bound of two. On the other hand, an upper bound of 7±2 was made under the assumption that the short-term memory in humans can hold only 7±2 chunks [28]. Therefore, the agent's mind may not be able handle more than a length of 7±2 consecutive action comparisons. Considering an agent such as a pet, we decided to limit the length of the patterns learnt to three consecutive actions.

A typical pattern is illustrated as follows: the owner goes into the kitchen (a_1), takes out the pet's food from the cupboard (a_2) and feeds the pet (a_3). These three consecutive actions led to the pet being fed. Thus, if the owner goes to the kitchen again, the pet would expect to be fed with some probability. In general, learning sequences of actions can be very useful in predicting the user's next action.

Tables of counts are kept. These tables are used to define the conditional probability $p(e_3| e_1, e_2)$ denoting the probability of an event e_3 to occur, given that events e_1 and e_2 have just occurred. 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. 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 occurring, given that two previous events X and Y occurred. The expected probability of event Z is calculated as follows:

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

Cases where experience is limited can be handled by reducing the probability to be conditioned only on one previous event. For example, if the sequence Y and Z was observed, then the probability of event Z occurring is:

$$P(Z \mid Y) = \frac{C[Y,Z]}{\sum_i C[Y,i]}$$

However, if the prior event Y has never been seen to precede the event Z, the probability of event Z occurring will be calculated as an unconditioned average (prior probability over all events):

$$P(Z) = \frac{C[Z]}{\sum_i C[i]}$$

These probabilities allow the agent to determine how likely the user is to take certain actions, given the events that have recently occurred, which is required for reinforcement learning, discussed above. It should be noted that the patterns discussed above only account for the user's actions, since the goal was to adapt to and learn about the tendencies of the user. A useful but more complex extension would be to interleave the agent's reactions in the patterns.

Standards: Learning values of actions

The learning algorithms described so far allow the agent to associate expectations and rewards to events. Even though learning about rewards and expectations is very important, it is not sufficient to simulate emotions such as pride, shame or admiration. Another learning algorithm is needed to learn immediate values of actions taken. To learn about values of actions, we adopted a new simple learning algorithm. We allow the user to give the agent feedback on its actions. For example, consider an agent who takes an action *YellAt*(User), the user may as a consequence take an action, such as *YellAt*(agent) or

Hit(agent). Each user action or feedback is rated within the agent's knowledge-base; for example, *YellAt*(agent) may have a value of -2 and *Hit*(agent) may have a higher negative value such as -6. Thus, consider a user action u_a, using the knowledge-base, we can define a function *evaluate*(u_a) that evaluates the given user action. Using this function the agent can then calculate the values of its own actions. The expected value of an action *a* is calculated as the sum of the values of user feedback given after the occurrence of action *a* over the number of occurrences of action *a*, which is formalized as follows:

$$E(v(a)) = \frac{\sum_t eval(u_a(t))}{n_a}$$

Where *a* denotes the agent's action, $eval(u_a(t))$ denotes user feedback at time *t*, n_a denotes the number of occurrence of action *a* during the simulation and *E(v(a))* denotes the expected value of a specific action *a*.

The agent can then use this formula to assess it's own actions and trigger emotions, such as pride and shame. Additionally, it can trigger emotions such as admiration or reproach by assessing the value of the other agent's actions. The agent may also use these values to determine actions that will be pleasing to the user or to other agents.

3.2 Emotions and Motivations Filtering Phase

Using the learning algorithms discussed above a mixture of emotions with different intensities will be triggered. This mixture will then propagate through different links affecting different processes, as it was asserted earlier. One of these links is the link between emotions and the motivational states.

Motivational states are internal states that promote or drive subjects to take a specific action. Some psychologists have explored the links between emotions and some motivational states. One such study was done by Higgins [19]. Higgins focused on exploring the link between affective states and self-regulatory motivational states [19]. Other psychologist have explored the link between motivational states such as fear and pain [12]. Fanslow's work provides some insight on the impact of emotions, such as fear, on some motivational states, such as pain. In his work they defined different states where pain may inhibit fear. He has additionally defined other states where fear inhibits pain [12].

In this paper, we consider hunger, fatigue, thirst and pain as motivational states. These states were linearly incremented during the simulation using some internal state measures, such as time, water loss, action (running, walking, etc.). Once these motivational states reach a sufficient level, they send a signal to the cognitive process indicating a specific need that has developed. These motivational states can then block the processing of the emotional component to produce a plan to enable the agent to appease its needs, whether it is water, food, sleep, etc. The plan was simulated using a number of simple rules, which depend on the agent's situation at the particular time. For example, if the agent already has access to water, then it will drink, but if it does not have access to water and knows its whereabouts, then it will form a plan to get it. However, if the agent must depend on another agent to satisfy its needs then it will use the model that it learned about the other agent to send it a signal.

According to Fanselow's model, fear inhibits pain if (1) the cause of fear is present and (2) the fear level is higher than the pain level. At a later time step, when the cause of fear disappears, pain will inhibit the fear. We have taken a similar approach in simulating the interaction between emotions and motivational states. Before inhibiting emotions, we make a situation assessment and an emotional versus motivational states assessment. A number of simple rules were used to measure the quality of a particular action that is selected according to the specific emotion or motivational state. These rules measure the quality of actions according to their goal's achievement and priority. The emotion or motivation with the action that has the maximum quality will take precedence, while the others will be inhibited.

It should be noted and recognized that the link between motivational states and emotions goes much more beyond what is described here. We are currently exploring the possibility of expanding the definition of motivational states to incorporate other types of states, such as the self-regulatory motivational states discussed by Higgins [19], that may bring about other interesting relations and interactions between motivations and emotions.

3.3 Behavioral Selection Phase

The behavior of an agent depends on its emotional state and the situation or the event that occurred. For example, consider the following rule:

If	Anger is High
	AND Event is dish-was-taken-away
THEN	behavior is Bark-At-user

The behavior, *Bark-At-User*, depends on what the user did and the emotional intensity of the agent. If the user did not take the dish away and the agent was angry for some other reason, it would not necessarily be inclined to bark at the user, because the user might not be the cause of its anger. Thus, it is important to identify both the event and the emotion. It is equally important to identify the cause of the event. In this case, we are assuming that non-environmental events, such as *dish-was-taken-away*, *throw-ball*, *ball-was-taken-away*, etc. are all caused by the user.

Rather than using one emotion to define the behavior of the agent, we are using the mixture of emotions and their intensities to formulate and blend emotional expressions or behaviors. Thus, each mixture results in a different behavior. Generally, we used rules in the form given below.

IF	emotion_1 is A
	AND emotion_2 is B
	….. AND emotion_k is C
THEN	BEHAVIOR is F

where *k* is the number of emotions in the system. *A*, *B* and *C* are fuzzy sets defining the emotional intensity as being *HighIntensity*, *LowIntensity* or *MediumIntensity*. Behaviors are simulated as singletons, including *Look-At-User* and *Play-With-User.*

To elaborate on how behaviors are selected in the model, we will present an example involving fear and anger. Consider, for instance, every time you take the food disk away from the dog you hit it to prevent it from jumping on you and barking at you after taking its food away. Taking the food dish away produces anger, because the pet will be experiencing both distress and reproach, anger is a compound emotion of both reproach and distress. The pet will feel reproach because, by nature, it disapproves of the user's action (taking the food dish away), and it will be distressed because the event by definition is unpleasant. Additionally, since the user hits the dog whenever he/she takes the dish away, fear will be produced as a consequence. Thus, taking the dish away will produce both anger and fear. Using generalized fuzzy rules described above, the rule fired will be as follows:

IF Anger is *HighIntensity*
AND fear is *MeduimIntensity*
AND Event is *dish-was-taken-away*
THEN BEHAVIOR is growl.

Therefore, as it can be seen, the behavior was much less aggressive than with anger alone. In effect, the fear dampened the aggressive behavior that might have otherwise been produced. While some emotions, such as fear and anger, may have converging effects on the behavior, other mixtures may have diverging effects. For example, a mixture of emotions that may have distress and joy both with high intensities may drive the agent to exhibit an oscillation between two extreme behaviors, as we shall discuss later.

4. RESULTS AND DISCUSSION

We have implemented the model discussed in the previous section in an agent called PETEEI — a PET with Evolving Emotional Intelligence. The interface of the simulation is shown in Figure 3. The interface is based on the verb-object metaphor discussed in [2]. The user can interact with the pet through various actions, including (1) introducing objects to the pet's environment, (2) taking objects away from the pet's environment, (3) hitting objects, including the pet, and (4) talking aloud.

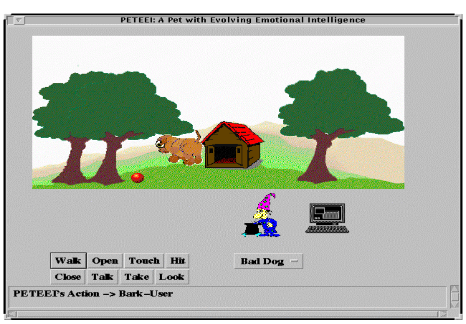

Figure 3. Model's Interface

4.1 Evaluation of PETEEI

We hypothesize that the simulation of the emotional process as a web or a dependant process may enhance believability, and thus we decided to validate the approach using three variations of the model discussed. The first model simulates behaviors and emotions that are chosen randomly during the simulation. The second model, simulates the emotional process as an independent process. Even though this model represents the emotional process as an independent entity, the inevitable link between the emotional process and the behavioral process exits, which is due to the choice of the evaluation procedure. The evaluation experiment requires the representation of behaviors as manifestations of the agent's internal states. The third model simulates part of the emotional web discussed in section 3, particularly it simulates the link between emotions and learning. We used the same interface, discussed above and shown in Figure 3, for all models.

Twenty-one users were assembled to evaluate the model and the effectiveness of the approach used. We asked the users to run the simulation and perform certain actions according to a scripted scenario given to them at the beginning of the session. As soon as they finished running the simulation, they were asked to answer a questionnaire. One of the questions asked the user to rate PETEEI's behaviors in terms of their believability in a scale from 1 to 10 (1: Not believable, 10: Extremely believable).

Answers were collected and an average was calculated for each question. The average is shown in Table 1 as the mean. We also calculated the confidence interval for a 95% confidence with a sample size of 21 users. This measure was calculated using the following formula:

$$\bar{x} \pm 1.96(\frac{\sigma}{\sqrt{n}})$$

where $\bar{x}$ is the mean, the σ is the variance and n is the size of the sample which is 21 in this case. The confidence interval is also shown in the table. As the table shows, the random model (model 1) did not convey a realistic or a convincing behavior to the user since the rating was on average 1. In contrast, the introduction of the model of emotions without the links (model 2) improved this measure by about 3.33 units on average. The model of the emotional process with the learning link (model 3) further improved this measure by 2.6 to get an average of 8.1 units. Therefore, we conclude that by simulating the link between emotions and learning, we were able to enhance the believability of the agent modeled. Subsequently, we hypothesize that by including the links between the emotional process and the other different processes and states, we could potentially further improve the believability of the agent simulated.

Table 1. Behavior Ratings

Model	Behavior	
	mean	interval
1	1	0.67 - 1.33
2	4.33	3.60 - 5.07
3	8.095	7.65 - 8.542

4.2 Oscillation of Emotions – an Observation

Consequently, we ran various experiments where we asked 10 users to run an interactive simulation for a thirty-minute interval using a scripted scenario in which various objects were given or taken from PETEEI. These scenarios were meant to show the possible ranges of emotional mixtures and behavioral blends that the agent can simulate. After the users ran these simulations, they were asked to comment on the agent's behaviors using a two paragraph written summary of their assessments. One of the most interesting observations that was noted by users during this simulation is the oscillation of extreme behaviors due to the oscillations of extreme emotions. To adequately explain his observation, we will detail the specific events and emotions that have resulted in such an observation.

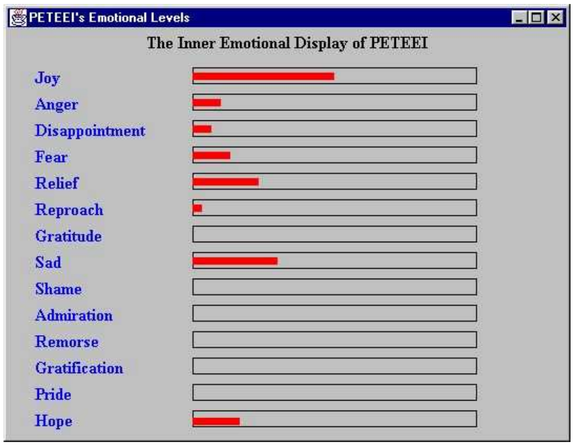

(a) Before event

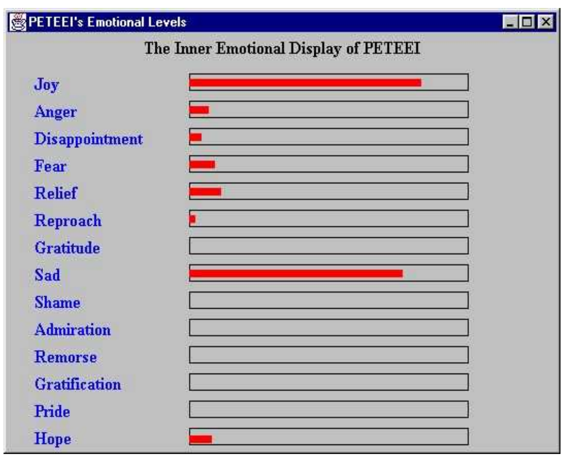

(b) After Event

Figure 4. Window of PETEEI's Emotions

At time t_0 = 1500 (simulation time steps, or about 2 minutes of real time) during one execution, the pet's emotional state was noted to be very conflicting. At time step t_0-1 (1499 simulation time), an event *give(drink)* was executed. This event was executed by the user according to the scripted scenarios. The event implies that the user gives PETEEI water to drink. This event triggered a very high intensity of both joy and distress as was projected in a window showing the emotional state of the agent, which is shown in Figure 4(b). Additionally, Figure 4 shows the emotional state of the agent (PETEEI) before and after the event was executed. It appears that due to the learning process and previous experience subjected by negative feedback from the user, PETEEI learned that water or drink is an undesirable event, and thus it causes distress. Using the primary desirability measures, however, water or drink was projected as a very desirable event, because it eliminates thirst. Therefore, PETEEI was facing an event that was evaluated as both desirable and undesirable, and thus producing emotional conflict.

PETEEI's behavior was even more intriguing. After event, *give(drink)*, was executed, PETEEI immediately acted on the joy emotion by wiggling its tail. Two time-steps later, however, it paused with a confusion look for 3 time-steps. It started barking, acting on the distress emotion, for another 3 time-steps. It then paused again with confusion for 6 time-steps, after which it resumed barking. PETEEI's behavior then oscillated between confusion and barking until the barking faded away. The intensity of distress and joy could also be deduced from the frequency of the barking and wiggling actions. For instance, it was observed in the simulation that PETEEI barked 11 barks per time-step when intensity of distress was at maximum then the number of barks slowly declined to zero as distress decayed through time. Furthermore, Figure 5 shows the intensities and the decay of the emotions joy and distress. In the figure, the dark gray line represents joy, while the light gray line represents distress.

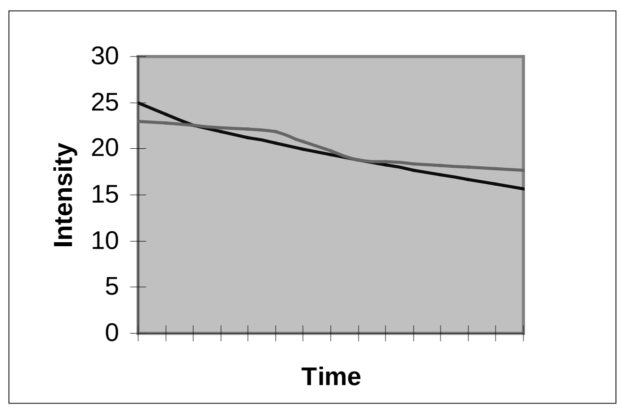

Figure 5. Emotional Decay

The oscillation or emotional confusion mechanism described here is dependent on the decay function used, as it can be seen from Figure 5. In our model, we set a larger decay value for negative emotions than for positive emotions, and thus the distress action prevailed over the joy action. However, the decay value could also be set to equal values for both emotions. In this case, one would not see the kind of oscillation shown here, but instead one would see a complete emotional confusion with no action, but a confusion look.

This observation can be compared to the observations of real emotions. People often cry and laugh almost at the same time in occasions such as marriages and graduation. This oscillation of behaviors may have occurred because of the intensity of a single emotion involved or it may be due to the coexistence of different or opposing emotions. We have also seen people in an emotional confusion state where they look as if they don't know whether to laugh or cry. Many other examples can be observed or extracted from real life emotional experiences. However, the issue of whether this observation presents a useful component that could potentially enhances the believability of a character or the functionality of agents within their environments is undeniable. The utility of such mechanism can be foreseen within both domains. For example, we could easily foresee the advantages of such scheme for creating a believable character, because people often relate more to states of emotional confusion than clear emotions.

According the users' both qualitative and quantitative assessments of the agent's behaviors, we can conclude that triggering mixtures of emotions using the learning process, and simulating the dependency between these mixtures, motivational states and behaviors enhanced the believability of the agent simulated. Additionally, the simulation of such an emotional structure facilitated the emergence of more complex and realistic behavior.

5. FUTURE WORK

The approach presented and discussed in this paper is mainly another view of the emotional process. A view by which many ideas that have been presented through the previous models (psychological and computational models) can be further expanded. Our evaluation indicates that incorporating the link between learning and emotions have improved the believability of the agent simulated and produced dynamic emotional responses. Nevertheless, the model still falls short in several aspects.

The model simulates motivational states as drives that interrupt the cognitive process and may inhibit emotions in some situations. While this may be true of some motivational states such as pain, some others have a more profound role especially in relation to attention and emotions [7]. Higgins, for example, has explored the different roles that non-emotional motivational states, particularly those concerned with self-regulatory orientations, can play in relation to affective states [19]. Even though our model simulated and attempted to explore the link between emotions and motivations, the model does not represent a complete picture of this link. We are currently exploring the possibility of extending the model to handle other direct or indirect interactions between emotions and motivations.

Among future research issues that should be presented and simulated in the model is personality. Many researchers within the agents community have explored the inclusion of personality in an agent's model [40]. Many psychologists have explored the relationship between emotions or moods and personality [37]. Therefore, we argue that personality should not be presented as an independent entity, but instead it should be presented and discussed within an emotional web architecture whereby the link between emotions and personality can be simulated. We are currently investigating the feasibility of expanding our model to include personality as another process that interacts with the emotional process.

6. CONCLUSION

Emotion is a very complex concept that affects many processes within the mind. In an attempt to simulate such a complex structure, we describe an architecture that simulates the emotional process as a web-like process, where the links between emotions, learning, motivational states and behaviors are incorporated and further discussed. To validate the utility of such an architecture, we have developed and described a prototype agent called PETEEI. PETEEI was tested and evaluated using twenty one users. Some emergent behaviors resulting from the simulation was discussed. Additionally, results of quantitative evaluation experiments were discussed. These results suggest that exploring the link between emotions and learning produced a significant improvement in the believability of PETEEI as perceived by the users. Thus, leading to the conclusion that building the emotional process as a web-like structure is a worthwhile a goal to pursue in an attempt to enhance the believability of agents. Additionally, treating emotions as a web rather than a simple independent process may prove a successful technique to many other applications where a realistic emotional model is needed, such as decision-making models.

7. REFERENCES

[1] Aron, A. and Aron, N. Self-expansion motivation and including other in the self. In S. Duck (ed.), *Handbook of personal relationships: Theory, research, and interventions (*2nd ed., pp. 251-270). Chichester, England:Wiley, 1997.

[2] Baecker, R. M., Grudin, J. S. Buxlon, W. A., and Greenberg, S. Eds. *Readings in Human-Computer Interaction: Toward the Year 2000, 2nd. Ed*. San Francisco, CA: Morgan Kaufmann Publishers, 1995.

[3] Bates, J. The role of emotion in believable agents. *Communications of ACM,* Vol. 37, No. 7, pp. 122-125, 1992.

[4] Bower, G. H. and Cohen, P. R. Emotional influences in memory and thinking: Data and theory. In M. Clack and S. Fiske (ed.), *Affect and Cognition* (pp. 291-331). London, England:Lawrence Erlbaum Association Publishers, 1982.

[5] Damasio, A. Descartes' error: Emotion, reason, and the human brain, New York: G. P. Putnam, 1994.

[6] Dautenhahn, K. The Art of Designing Socially Intelligent Agents: Science, Fiction and the Human in the Loop. *Applied Artificial Intelligence journal, Special Issue on "Socially Intelligent Agents"*, Vol. 12, No. 7-8, pp. 573-619, 1998.

[7] Domjan, M. *The principles of learning and behavior* (4th ed.). Boston: Brooks/Cole Publishing Co., 1998.

[8] Elliot, C. The Affective Reasoner: A Process model of emotions in a multi-agent system. Institute for the Learning Sciences, Evanston, IL: Northwestern University, Ph.D. Thesis, 1992.

[9] Elliott, C., Rickel, J., and Lester, J. Integrating Affective Computing into Animated Tutoring Agents, *Fifteenth International Joint Conference on Artificial Intelligence '97 —Animated Interface Agents Workshop*, pp. 113-121, 1997.

[10] Elliot, C. and Brzezinki, J. Autonomous agents as synthetic character. *AI Magazine,* AAAI press, pp. 13-30, 1998.

[11] Eibl-Eibesfeldt, I. *Human ethology,* New York: Aldine de Gruyter, 1989.

[12] Fanslow, M. A Neural Organization of the defensive behavior system responsible for fear, *Psychonomic Bulletin & Review,* Vol. 4, pp. 429-438, 1994.

[13] Fiske, A. P. The four elementary forms of sociality: Framework for a unified theory of social relations, *Psychology Review,* Vol. 99, pp. 689-723, 1992.

[14] Forgas, J. P. Mood and judgment: The affect infusion model (AIM), *Psychological Bulletin*, Vol. 117, pp. 39-66, 1995.

[15] Forgas, J. P. Sad and guilty?: Affective influences on the explanation of conflict in close relationships, *Journal of Personality and Social Psychology*, Vol. 66, pp. 56-68, 1994.

[16] Frank, A., Stern, A. and Resner, B. Socially Intelligent Virtual Petz, *AAAI Fall Symposium*, 1997.

[17] Gardner H. *Frames of Mind,* New York: Basic Books, 1983.

[18] Goleman, D. *Emotional Intelligence,* New York: Bantam Books, 1995.

[19] Higgins, Tory E. Promotion and Prevention Experiences: Relating Emotions to Non-Emotional Motivational States, in J. P. Forgas (Ed.) *Handbook of Affect and Cognition.* In press.

[20] Huhns, M. and Singh, M. P. *Readings in Agents.* USA:Morgan Kaufmann, 1998.

[21] Kaelbling, L. P., Littman, M. L., and A. W. Moore Reinforcement Learning: A Survey, *Journal of Artificial Intelligence Research,* 4, 237-285, 1996.

[22] Kim, J. and Moon, J. Y. Designing towards emotional usability in customer interfaces - trustworthiness of cyber-banking system interfaces. *Interacting with computer: The interdisciplinary journal of human-computer interaction,* Vol. 10, 1998.

[23] Konev, A., Sinisyn, E., and Kim, V. Information Theory of Emotions in Heuristic Learning, *Zhurnal Vysshei Nervnoi Deyatel'nosti imani*, Vol. 37, No. 5, pp. 875-879, 1987.

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

[25] Lester, J. C. and Stone, B. S. Increasing Believability in Animated Pedagogical Agents, *Autonomous Agents Conference 97*, 1997.

[26] Loyall, A. B. and Bates, J. Personality-based believable agents that use language, *Proc. of the first autonomous agents conference,* 1997.

[27] Maes, P. Agents that reduce work and information overload, *Communications of ACM,* Vol. 1, No. 7, pp. 37-45, 1995.

[28] Miller, G. A. The magical number seven, plus or minus two: Some limits on our capacity for processing information, *Psychological Review,* Vol. 63, pp. 81-97, 1956.

[29] Minsky, M. *The society of the mind*, New York: Simon and Schuster, 1986.

[30] Mitchell, T. M. *Machine Learning,* New York: McGraw-Hill Co., 1996.

[31] Nye, J. L. and Brower, A. M. *What's social about social cognition?* London: Sage Publications, 1996.

[32] Ortony, A., Clore, G., and Collins, A. *The cognitive structure of emotions*. Cambridge: Cambridge University Press, 1988.

[33] Pfeifer, R. Artificial Intelligence Models of Emotions. *Cognitive perspective on emotion and motivation,* pp. 287-320, 1988.

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

[35] Reilly, W. S. and Bates, J. Building emotional agents. Pittsburgh, PA: Carnegie Mellon University, Technical Report CMU-CS-92-143, 1992.

[36] Reilly, W. S. Believable social and emotional agents. School of Computer Science, Pittsburgh, PA: Carnegie Mellon University, Ph.D. thesis, 1996.

[37] Revelle, W. Individual differences in personality and motivation, Non-cognitive determinants of cognitive performance. In A. Baddeley & L. Weiskrantz (Eds.), *Attention: Selection, awareness and control: A tribute to Donald Broadbent,* pp. 346-373, Oxford: Oxford University Press, 1993.

[38] Rickel, J. and Johnson, W. L. Animated Agents for Procedural training in Virtual reality: Preception, Cognition, and Motor Control. To appear in *Journal of Applied Artificial Intelligence*.

[39] Riso, D. R. and Hudson, R. *Personality Types*. USA: Houghton Mifflin, 1996.

[40] Rousseau, D. and Hayes-Roth, B. Interacting with Personality-Rich Characters Stanford Knowledge Systems Laboratory Report KSL-97-06, 1997.

[41] Rusbult, C. E. and Arriaga, X. B. Interdependence theory. S. Duck (ed.), *Handbook of personal relationshps: Theory, research, and interventions* (2nd ed., pp. 221-250). Chichester, England:Wiley, 1997.

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

[43] Simon, H. Motivational and emotional control of cognition. *Psychological review,* 74, 29-39, 1967.

[44] Simon, H. *The science of the artificial.* Cambridge, MA: MIT Press, 1996.

[45] Suchman, L. *Plans and situated actions,* Cambridge University Press, 1981.

[46] Thomas, F. and Johnston, O. *The illusion of life,* New York: Abbeville Press, 1981.

[47] Velasquez, J. Modeling emotions and other motivations in synthetic agents, *Proc. of the AAAI conference*, pp. 10-15, 1997.