Probabilistic Assessment of User's Emotions in Educational Games

Cristina Conati

Department of Computer Science
University of British Columbia
Vancouver, BC V6T1Z4
conati@cs.ubc.ca, +604-8224632

Abstract

We present a probabilistic model to monitor a user's emotions and engagement during the interaction with educational games. We illustrate how our probabilistic model assesses affect by integrating evidence on both possible causes of the user's emotional arousal (i.e., the state of the interaction) and its effects (i.e., bodily expressions that are known to be influenced by emotional reactions). The probabilistic model relies on a Dynamic Decision Network to leverage any indirect evidence on the user's emotional state, in order to estimate this state and any other related variable in the model. This is crucial in a modeling task in which the available evidence usually varies with the user and with each particular interaction. The probabilistic model we present is to be used by decision theoretic pedagogical agents to generate interventions aimed at achieving the best tradeoff between a user's learning and engagement during the interaction with educational games.

1. Introduction

In recent years, there has been an increasing interest in studying how to make computers more "sociable" by enabling them to both display their own emotions and react to the user's emotions. Building computers that display emotions in a natural and meaningful way is already a challenging endeavor, since it requires formalizing concepts and mechanisms that are often still under investigation in emotional psychology. But building computers that recognize a user's emotions is even more challenging, as is proven by the fact that even human beings are not always proficient in this task. The challenge is due to the high level of ambiguity that exists in the mapping between emotional states and the factors that can be used to detect them. For instance, different people can have different emotional reactions to the same stimulus, and the variability depends upon traits that are not always easily observable, such as a person's goals, preferences, expectations and personality. Emotions can be recognized because they often have observable effects on a user's behavior and bodily expressions. But the mapping between emotions and their observable effects also depends on often hidden traits of a person, as well as on the context of the interaction. Furthermore, observable effects of emotions are not always easily recognizable by a computer (i.e., subtle changes in facial expression and intonation).

Existing approaches have tackled the challenge of recognizing user's affect by trying to reduce the ambiguity in the modeling task. This has been achieved either by focusing on recognizing a specific emotion in a fairly constraining interaction (Healy and Picard, 2000; Hudlicka and McNeese, 2002) or by assessing only lower level dimensions of emotional reaction, such as its intensity and valence[1] (Ball and Breeze, 2000).

In this paper, we present an approach to modeling user affect designed to assess a variety of emotional states during interactions in which knowing the details of a user's emotional reaction can enhance a system capability to interact with the user effectively. Instead of reducing the uncertainty in emotion recognition by constraining the task and the granularity of the model, our approach explicitly encodes and processes this uncertainty by relying on probabilistic reasoning. In particular, we use Dynamic Decision Networks (Dean and Kanazawa, 1989; Russell and Norvig, 1995) to represent in a unifying framework the probabilistic dependencies between possible causes and emotional states (including the temporal evolution of these states), and between emotional states and the user's bodily expressions they can affect. Our goal is to create a model of user affect that can generate as accurate an assessment as possible, by leveraging any existing information on the user's emotional state, but that can also explicitly express the uncertainty of its predictions when little or ambiguous information is available.

We discuss our model in the context of the interaction with pedagogical agents designed to improve the effectiveness of computer-based educational games (which we will simply call *educational games* throughout the paper). In the rest of the paper, we first describe why detecting emotions is important for educational games. We then introduce Dynamic Decision Networks (DDN) and illustrate how they can be used to enable pedagogical agents for educational games to generate interactions tailored to both the user's learning and emotional state. Next, we describe in detail the DDN underlying our model of user

[1] Valence measures whether the emotion generated a positive or negative feeling

affect and how it integrates in a principled way different sources of ambiguous information on the user's emotional state. We end with an overview of related work, discussion and conclusions.

2. Emotionally Intelligent Agents for Educational Games

Several authors have suggested the potential of video and computer games as educational tools (e.g., Silvern, 1986; Malone and Lepper, 1987). However, empirical studies have shown that, while educational games are usually highly engaging, they often do not trigger the constructive reasoning necessary for learning (Conati and Fain Lehman, 1993; Klawe, 1998).

An explanation of these findings is that it is often possible to learn how to play an educational game effectively without necessarily reasoning about the target domain knowledge (Conati and Fain Lehman, 1993). Possibly, for many students the high level of engagement triggered by the game activities acts as a distraction from reflective cognition. This seems to happen especially when the game is not integrated with external activities that help ground the game experience into the learning one. Also, educational games are usually highly exploratory in nature, and empirical studies on exploratory learning environments have shown that these environments tend to be effective only for those students that already possess the learning skills necessary to benefit from autonomous exploration (e.g., Shute, 1993).

To overcome the limitations of educational games, we are working on designing intelligent pedagogical agents that, as part of game playing, can generate tailored interventions aimed at stimulating a student's reasoning if they detect that the student is failing to learn from the game. "*As part of game playing*" is the key point in the design of these agents. The main advantage of educational games versus more traditional computer-based tutors is that the former tend to generate a much higher level of students' positive emotional engagement, thus making the learning experience more motivating and appealing. In order not to lose this advantage, it is crucial that the interventions of pedagogical agents be consistent with the spirit of the game and consider the players' emotional state, in addition to their learning. On the one hand, these pedagogical agents need to make sure that a student learns as much as possible from the game. On the other hand, they also need to avoid interventions that make the student start seeing the interaction with the game more as an educational chore than as a fun activity. Thus, at any point during the player interaction with the game, a pedagogical agent may need to consider the tradeoff between the player's learning and entertainment when deciding how to act. The more information the agent has on the student's learning and emotional state, the more focused and effective its actions can be. We formalize this behavior by designing our pedagogical agents as decision theoretic agents (Howard and Matheson, 1977; Russell and Norvig, 1995) that select actions so as to maximize the outcome in terms of a student's learning and emotional engagement, as we describe in the next section.

3. Decision-theoretic Pedagogical Agents

In a decision-theoretic model (Howard and Matheson, 1977), an agent's preferences over world states S are expressed by a utility function U(S), which assigns a single number to

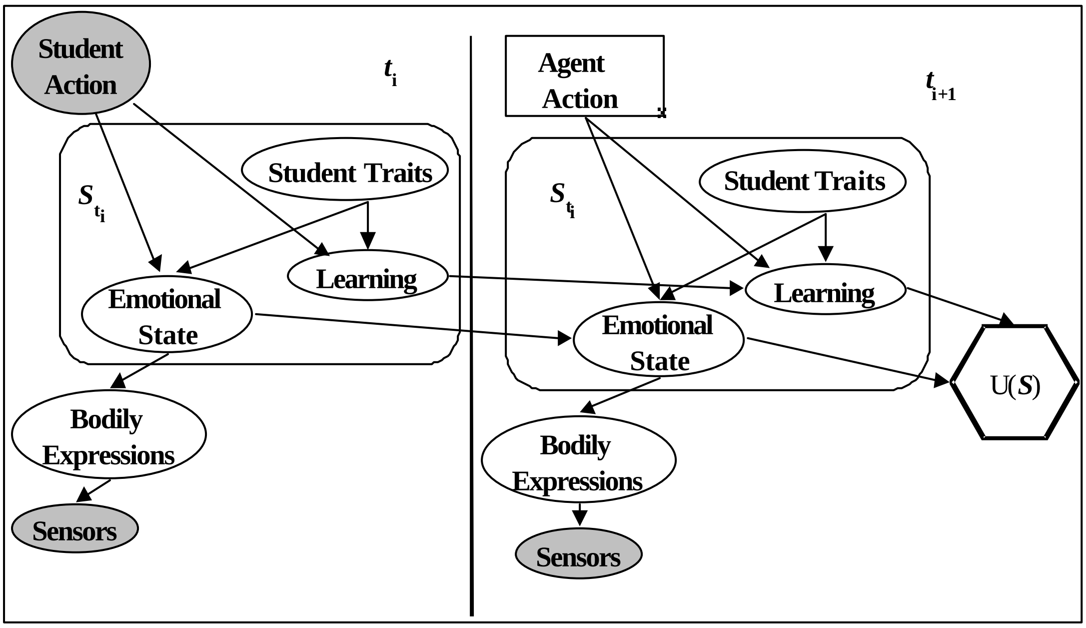

Figure 1: DDN to model the decision of a pedagogical agent

express the desirability of a state. Furthermore, for each action *a* available to the agent, and for each possible outcome state S' of that action, P(S'|E, *a*) represents the agent's belief that action *a* will result in state S', when the action is performed in a state identified by evidence E. The expected utility of an action *a* is then computed as

$$EU(A) = \Sigma_{S'} P(S'|E, a)U(S')$$

A decision-theoretic agent selects the action that maximizes this value when deciding how to act.

Decision Networks (DNs), or *influence diagrams* (Henrion, Breeze and Horvitz, 1991), are an extension of Bayesian Networks (Pearl, 1988) that allow modeling decision-theoretic behavior. In addition to nodes representing probabilistic events in the world, a DN includes nodes representing an agent's decision points and utilities. By relying on propagation algorithms for Bayesian networks, DNs allow computing the agent's action (or sequence of actions) with maximum expected utility given the available evidence on the current state of the world.

Dynamic Decision Networks (DDNs) add to DNs the capability of modeling environments that change over time. Figure 1 shows how a DDN can be used to define the behavior of pedagogical agents that take into account both the student's learning and emotional reactions when deciding how to act. This DDN models behavior over two *time slices*, to answer the question: given the student's state S_{ti} at time t_i, what is the agent's action that will maximize the agent's expected utility at time t_{i+1}, defined in terms of the student's learning and emotional state at that time?

In a DDN, the links between variables in different time slices indicate that the values of these variables evolve over time and that the value at time t_i influences the value at time

t_{i+1}. In Figure 1, this is the case for the random variables *Learning* and *Emotional State* representing a student's learning and emotional state, respectively. The links between *Learning* nodes, for example, model the fact that a student is likely to know a given concept at time t_{i+1} if she knew it at time t_i. The links between *Emotional State* nodes encode that a student is more likely to feel a given emotion at time t_{i+1} if something that can trigger that emotion happens and the student was already feeling that emotion at time t_i. The shaded nodes in Figure 1 represent random variables for which evidence is available to update the student model at a given time slice. In Figure 1, this evidence includes the student's game action at time t_i, as well as the output of sensors for monitoring the student's affective response at time t_i and t_{i+1} (we will say more about these sensors in a later section). The rectangular node in time slice t_{i+1} represents the agent's available actions at that time, while the hexagonal node represents the agent's utility. To compute the agent's action with highest expected utility in this time slice, the DDN computes the expected value of each action given the evidence currently available at time slice t_i. The agent's decision node is then set to the action with the highest expected utility, and new evidence on the student's emotional reactions in collected to assess what emotional state the agent's action actually generated.

The links from the Learning and Emotional State nodes to the utility node in Figure 1 indicate that an agent's utility function is defined over the student's learning and emotional states. By varying this utility function, we can define agents that play different roles in the game. So, for instance, the utility function of a *tutoring-oriented* agent will assign higher values to states characterized by high levels of student learning, giving less importance to the student's emotional engagement. In contrast, the utility function of a *game-oriented* agent will value more those states in which the student is positively engaged.

In the rest of the paper, we will concentrate on illustrating the part of the DDN that assesses the user's emotional state, to show how a probabilistic model can deal with the high level of uncertainty involved in this still largely unexplored user modeling task. For simplicity, we will ignore any relation between emotional state and learning, as well as details on how assessment of learning is performed.

4. A Dynamic Decision Network for Modeling Affect

Figure 2 shows two time slices of the DDN that forms our model of student affect. The nodes in Figure 2 represent classes of variables in the actual DDN. As the figure shows, the network includes variables that represent both causes and effects of emotional reactions. Being able to combine evidence on both causes and effects aims to compensate for the fact that often evidence on causes or effects alone is insufficient to accurately assess the user's emotional state, as we illustrate in the next subsection.

4.1 Uncertainty in Modeling Affect

Although emotions often visibly affect a person's behaviour and expressions, the effects of emotions are not always discriminating enough to allow a precise diagnosis of the emotional states that generated them. For example, some accentuated facial expressions and prosody features can be quite indicative of specific emotional states such as fear, joy or anger (Ekman, 1993; Murray and Arnott, 1993). However, whether these intense emotion expressions arise usually depends on the intensity of the emotion, on the user's personality and on the interaction context. For instance, an introvert person might have a tendency to control her display of emotions, especially in the presence of people she is not well acquainted with. Thus, in many situations, changes in facial expressions and prosody might be too subtle to be easily detected, especially if the detection is done by a computer. Emotional states can also affect biometric measures such as heart rate, blood pressure, skin conductance, color and temperature (Picard, 1997). A person usually has little control over these covert biometric measures, and therefore they could provide a more reliable source of information on a person's affect. However, information on a single biometric measure is usually not sufficient to recognize a specific emotion. For instance, skin conductivity is a very good indicator of general level of arousal, but cannot identify the valence of the emotion that caused the arousal (Picard, 1997). Emotions with negative valence tend to increase heart rate more than emotions with positive valence (Cacioppo, Berntson, Poehlmann and Ito, 2000), but heart rate provides little information about specific emotions (Ekman, Levenson and Friesen, 1983).

Predicting emotions from possible causes is also not always easy. Although there are psychological theories that define the mapping between causes and emotional states, in practice information on possible causes does not always provide unequivocal indication on the actual affective reaction. Consider, for instance, the cognitive theory of emotion developed by Ortony Clore and Collins and known as the OCC model (Ortony, Clore and Collins, 1988). This theory defines emotions as valenced (positive or negative) reactions to situations consisting of events, actors and objects. The valence of one's emotional reaction depends upon the desirability of the situation for oneself, which in turn is defined by one's goals and preferences. The OCC theory clearly defines twenty two emotions as the result of situation appraisal, thus making it quite straightforward to predict a person's emotions if the person's goals and perception of relevant events are known. The problem is that this information is not always easily available when assessing a user's emotion.

The above factors make emotion recognition a task permeated with uncertainty. Most of the existing research on modeling users' affect has tried to reduce this uncertainty either by considering tasks in which it is relevant to only monitor the presence or absence of a specific emotion (Healy and Picard, 2000; Hudlicka and McNeese, 2002) or by focusing on monitoring lower level measures of emotional reaction, such as the intensity and valence of emotional arousal (Ball and Breeze, 2000). In educational games, neither of these approaches is appropriate, for two main reasons. First, educational games do tend to arouse different emotions in different players. For instance, the exploratory nature of a game can be very exciting for students that mainly want to have fun, while it may cause frustration or anxiety in students that want to learn from the game but tend to prefer more structured pedagogical activities. Second, detecting the student's specific emotions is important for an agent to decide how to correct possibly negative emotional states or leverage the positive ones. For example, if the agent realises that the student is ashamed because she keeps making mistakes during the game, it can try to take actions that make the student feel better about her performance. Or, if the agent realizes that the student enjoys its character but is distressed with the game at a particular point in time, it can initiate an interaction with the student with the sole purpose of entertaining her.

In the next subsection we describe how we use a DDN to explicitly represent the uncertainty underlying the relationships between a student's emotional states, their causes and effects during the interaction with educational games.

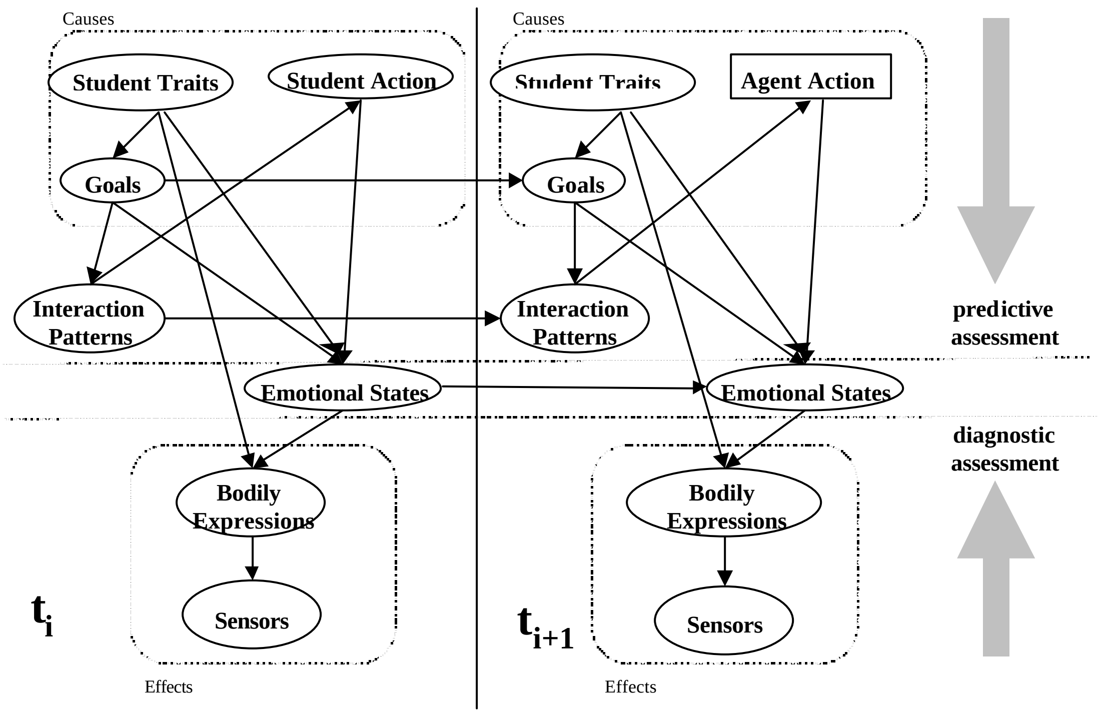

Figure 2: Two time slices of the DDN model of user affect

4.2 Probabilistic Dependencies Between Emotions, Their Causes and Their Effects

In our DDN, the causes of emotional arousal are modeled following the OCC cognitive theory of emotions described in the previous section. To apply this theory to the assessment of emotions during the interaction with educational games, our DDN includes variables for goals that students may have when playing one of these games, summarized in Figure 2 by the nodes *Goals*[2]. The subject of the student's appraisal is any event caused by either a student's game action (node *Student Action* in Figure 2, time slice t_i) or an agent's action (node *Agent Action* in Figure 2, time slice t_{i+1}). The probabilistic dependencies between student's goals, game states and emotional reactions are summarize in the DDN of Figure 2 by the links connecting the nodes *Goals* and *Student Action* (or *Agent Action*) to the node *Emotional States.*

User's goals are a key element of the OCC model, but it is often unfeasible to identify these goals with certainty, for example by asking the user. Thus, our DDN also includes nodes that can help the model infer the student's goals from indirect evidence. What goals a student has depends on the student's traits such as *Personality* (Matthews, Derryberry and Siegle, 2000) and *Domain Knowledge*, as represented by the links connecting the nodes *Student Traits* with the *Goals* nodes in Figure 2. Also, the student's goals can directly influence how a student plays the game, as modeled by the links between the

---

[2] We currently do not explicitly represent the player's preferences in our model .

nodes *Goals* and *Interaction Patterns* in Figure 2. In turn, interaction patterns can be inferred from specific features of the student's individual actions at each time slice. Thus, observations of both the relevant student's traits and game actions can provide the DDN with indirect evidence for assessing the student's goals.

The part of the network below the nodes *Emotional States* represents the interaction between emotional states and their observable effects. The node *Emotional States* directly influences the node representing the class of bodily expressions that are affected by emotional arousal. In turn, this node directly influences the node *Sensors*, representing devices that can detect the bodily expressions of interest. In recent years, there have been encouraging advances in the development of such devices, which include, among others, software for face and prosody recognition (Mozziconacci, 2001; Bianchi-Berthouze and Lisetti, 2002), as well as sensors to capture biometric signals (Picard, 1997). However, none of these devices, by itself, will always reliably identify a specific emotional state. By explicitly representing the probabilistic relationships between emotional states, bodily expressions and techniques available to detect them, our DDN can combine and leverage any available sensor information, and gracefully degrade when such information becomes less reliable.

In the rest of the paper, we describe an example application of the above model in the context of Prime Climb, the game we are using as a test-bed for our research.

5. The Prime Climb Educational Game

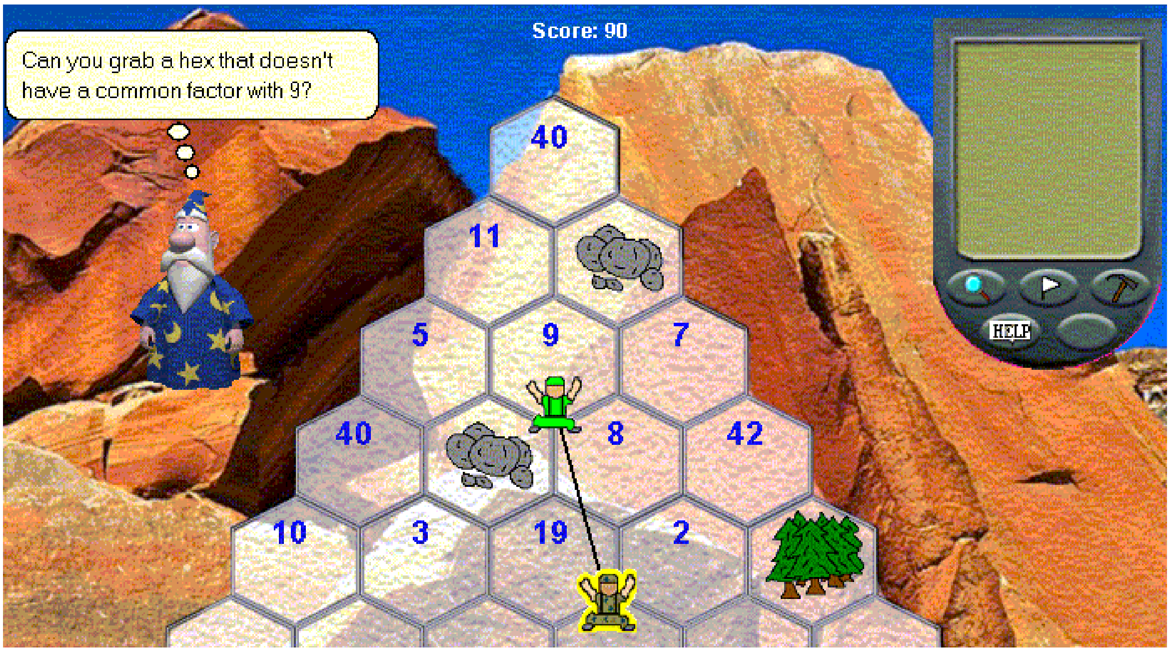

Figure 3: The Prime Climb Interface

Prime Climb is an educational game designed by the EGEMS (Electronic Games for Education in Math and Science) group at the University of British Columbia to help students learn number factorization. In Prime Climb, teams of two players must climb ice-faces divided into numbered sections (see Figure 3). Each player can only move to sections with numbers that do not share any factors with that occupied by the other team member. When a player moves to a section that does not satisfy the above constraint, the player falls and the team looses points. For instance, the player at the bottom in Figure 1 fell because she tried to move to section 42, which shares the factor 3 with section 9, where the other player is. To help the students understand factorization, Prime Climb includes tools to inspect the factorizations of the numbers on the mountain. These tools are accessible by clicking on the icons representing a magnifying lens and a flag on the PDA shown at the top-right corner of Figure 3.

An informal study of this version of Prime Climb showed that, while some students used and benefited from these additional tools, other ignored them even when they kept falling. Furthermore, many of the students who had very weak math knowledge and accessed the tools did not seem to gain much from their use. In light of these findings, we are designing pedagogical agents that, as part of Prime Climb, aim at stimulating a student's reasoning when they realize that the student is not learning from the game. One of the agents is a climbing instructor that can provide tailored help, both unsolicited and on demand, to help the student better understand number factorization as she is climbing, and that can do so without compromising the player's level of engagement. The actions that this agent can perform include stimulating the student to think about the reasons that caused a fall, giving more specific advice on how to recover from a fall (see Figure 3), suggesting and helping with the usage of the available tools, and deciding the level of difficulty of the climbing task.

We now show an illustrative example of how the general model in Figure 2 can be instantiated and used to allow the Prime Climb climbing instructor to monitor a player's emotional state and react adequately to it.

6. Sample Affective Model for the Interaction with Prime Climb

6.1 Model Variables and Structure

For the sake of simplicity, the model described in this example (shown in Figure 4) covers in detail only slice t_{i+1} of the general model shown in Figure 2, and includes only a subset of the variables that are necessary to completely specify this time slice. We chose this subset to give the reader a sense of how the model is built and of its workings, but several additional variables should be included to accurately model a real interaction.

All the variables and links in the model have been derived from findings described in relevant literature, from observations of students playing Prime Climb, and in a few occasions from our intuition. The conditional probabilities are currently based mainly on our estimates of relevant qualitative findings described in the literature, but we are working on revising them empirically.

Student's goals. By observing and interviewing students playing Prime Climb, we have derived a set of common high level goals that students may have when interacting with the game. We use three of these goals to exemplify the role of these variables in our model:

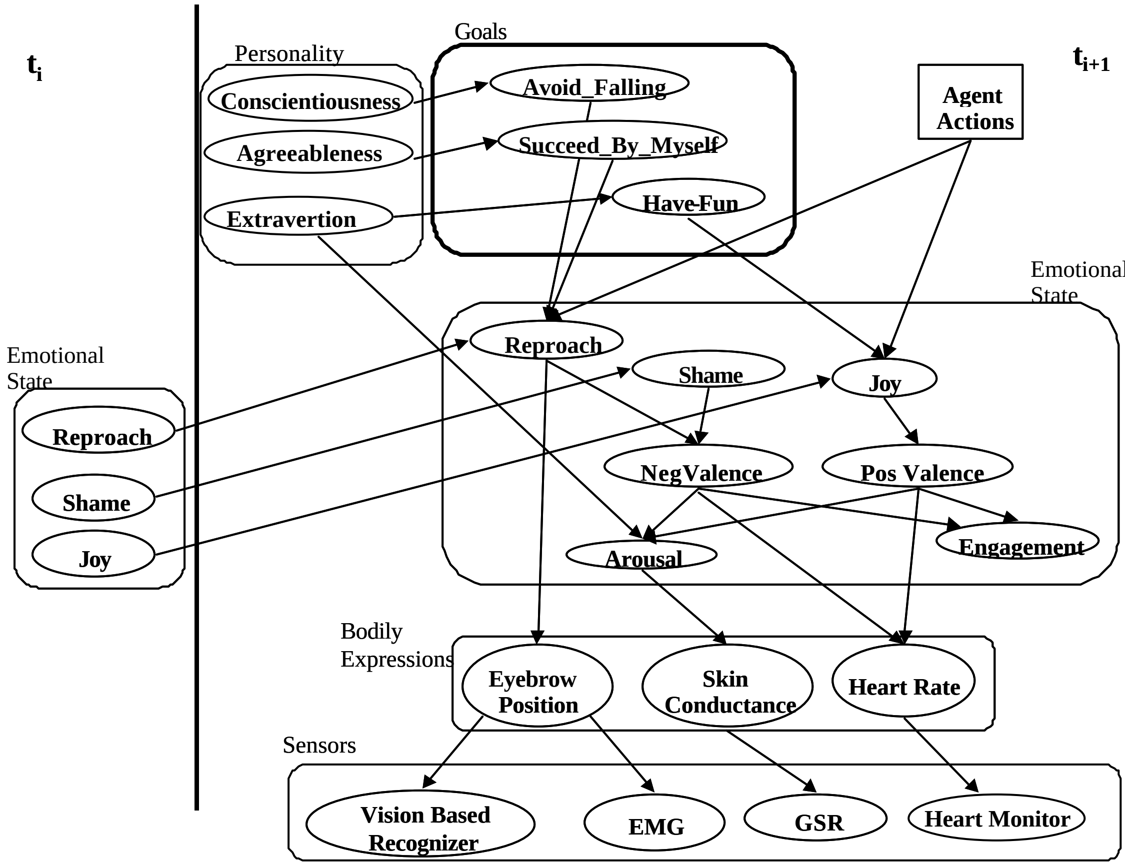

Figure 4: Sample portion of the affective model for Prime Climb

having fun (node *Have_Fun* in Figure 4), succeeding without the agent's help (node *Succeed_by_Myself*) and not falling (node *Avoid_Falling*).

Variables describing the student's personality traits. We consider three personality traits in this example, taken from the Five Factor Model of Personality (Costa and McCrae, 1992): *extraversion, agreeableness*, and *conscientiousness* (the two other personality types that are part of the Five Factor Model are *openness* and *neuroticism*). Each of these traits is represented by a node that has as values the two extremes of the personality type (e.g. *extrovert* and *introvert* for the node *extraversion*) Personality traits can directly influence what goals a student has (Matthews et al., 2000). The links between personality nodes and goals can be derived from the definition of the different personality types. For instance, the definition of an agreeable person includes the following statements "*…. eager to help…and believes that others will be equally helpful in return*". By contrast, the disagreeable person is "*egocentric, skeptical of others' intentions, and competitive rather than cooperative.*" This definition indicates that *agreeableness* can directly influence a player's goal to succeed in the game without any external help, and this influence is modeled in the network by a link between the node representing the *agreeableness* personality type and the goal *Succeed-by-Myself*. In addition, the conditional probability table (CPT) for *Succeed-by-Myself* is defined so that the probability of this goal is high for a disagreeable person, and low for an agreeable one. Similarly, the CPT for the node *Have_Fun* indicates that this goal is likely for an extravert player, while the CPT of the goal *Avoid_Falling* indicates that this goal is more likely for a person that is conscientious. Although in this example we have a one-to-one mapping between personality traits and goals, in reality, when additional goals and personality traits are considered, the mapping can be many-to-many. For instance, it is plausible for a conscientious person to have both the goal to avoid falling and the goal to learn math from the game. The goal to avoid falling is also compatible with a person belonging to the *neuroticism* personality type.

Personality traits can also directly influence emotional reactions. For instance, psychological studies have shown that introverts tend to reach a higher level of emotional arousal than extroverts, given the same stimulus (Kahneman, 1973). This is encoded in our network by linking the node for the *extraversion* personality type with the node representing the level of emotional arousal (see Figure 4), which we will describe later in the section.

Agent's actions. For this example, we will consider only two of the possible actions that the Prime Climb agent can generate: (1) provide help when the student makes a mistake, and (2) do nothing. These actions are represented as two different values of the decision node *Agent Actions* in Figure 4.

Variables describing the user's emotional state. Following the OCC cognitive model of emotions, we model the user's emotional state as the result of the user's appraisal of the current interaction event in relation to her goals. In our model, a new interaction event corresponds to either a student's or an agent's action and generates the addition of a new time slice in the DDN. To keep things simple, in this example we only consider a time slice corresponding to an agent's action (see Figure 4). The appraisal mechanism is explicitly modeled in the network by conditioning the nodes representing emotional states to both nodes representing user's goals and nodes representing interaction events (the node *Agent Actions* in this case). The nodes representing emotional states are also defined following the OCC theory of emotions. Out of the twenty two emotions that the OCC theory describes, we currently represent six that related to the appraisal of the direct consequences of an event for oneself[3]. These emotions include: *joy* and *distress* toward the event that is appraised by the user; *reproach* and *admiration* toward the entity that caused the event; *pride* and *shame* toward the entity that caused the event when the entity is oneself . For illustrative purposes, we'll consider only three of these emotions in our example (see *Emotional State* cluster in Figure 4): (i) *Reproach,* which arises when the behavior of the Prime Climb agent interferes with a player's goals; (ii) *Shame*, which is felt when the player is disappointed with the outcome of her actions in the game; (iii) *Joy* which arises in response to any interaction event that satisfies the student's goals.

Notice that in Figure 4 the node *Agent Actions* is linked only to the emotion nodes *Reproach* and *Joy*, not to the node *Shame.* This is because shame is an emotional reaction to the student's actions, not to the agent's actions, and therefore can be directly involved in the appraisal process only in the DDN time slices representing student's actions. When an emotion node is not directly involved in the appraisal process at a given time slice, its probability depends only upon the probability of the corresponding emotion node in the

[3] Other emotions relate for, instance, to the consequences of an event for others or to the evaluation of objects rather than events.

previous time slice and its CPT represents the fact that an emotional state persists over brief periods of time, but it slowly decays if no new event revives it.

Because we are interested in assessing the student's level of engagement in the game, a corresponding variable is inserted into the model, along with links representing how this variable is influenced by the valence of a user's emotions (represented in Figure 4 by the nodes *Pos_Valence* and *Neg_Valence*). The corresponding conditional probabilities are defined to express the rather simplifying assumption that emotions with positive valence increase the level of engagement, while emotions with negative valence decrease it. In a more complete model, we may want to explicitly represent how specific emotions affect engagement. A node representing the level of *arousal* is also included in the model, because information on the level of arousal can be relevant to judge how much a given emotional state influences the user's behavior. As shown in Figure 4, the node *Arousal* has as parents the two nodes representing the valence of the emotional state and the node representing the personality type *Extraversion*. Conditioning arousal to valence is slightly misleading, since these are two orthogonal dimensions of emotional states. However, in our network the *valence* nodes are linked to the *Arousal* node for the practical purpose of summarizing that an emotional reaction does exist, without having to link every single emotion node to *arousal*. Combined with the input coming from the node for *Extraversion*, the links from the *valence* nodes allow us to compactly define a CPT representing the finding that an introvert reaches higher levels of arousal than an extravert given the same stimulus (Kahneman, 1973). Directly linking the emotion nodes to the arousal node may become necessary if the model needs to represent the influence that specific emotions have on the intensity of the arousal.

Variables describing bodily expressions and sensors. Let's suppose that we have sensors to detect three types of bodily expressions: (i) *eyebrow position*, by using, for instance, software to detect facial expression and an electromyogram sensor (EMG) to detect muscle contraction; (ii) *skin conductance*, through a sensor that detects galvanic skin response (GSR); (iii) *heart rate*, through a heart rate monitor. All these sensors can already be donned in a fairly non-intrusive manner (Picard, 1997), and considerable research is being devoted to make these kinds of devices increasingly wearable. Each bodily expression B is linked to each sensor S that can detect it, as shown in Figure 4, and if multiple sensors are available, the DDN propagation algorithms can automatically integrate evidence data coming from all of them. By encoding the probability of a sensor's value S given each value of bodily expression B, the conditional probability P(S|B) specifies the reliability of each sensor. Because this measure can be independently specified for each sensor and for the bodily expression that it detects, the model allows one to easily include new sensors as they become available.

Likewise, each conditional probability $P(B|E_1,..,E_n)$, indicates how a set of emotional states $E_1,..,E_n$ affects a given bodily expression B. As information on a bodily expression not yet considered in the model becomes available, a new variable for this expression can be added to the model and linked to the emotion variables that influence it, thus increasing the amount of evidence that can be used to detect the corresponding emotions. The conditional probabilities linking emotions and bodily expressions in our sample model represent the following findings (Picard, 1997):

1. Frowning eyebrows are a very good indicator of negative emotions in the anger range, including reproach[4].
2. Skin conductivity is a very good indicator of the level of arousal.
3. Heartbeat increases more in the presence of emotions with negative valence.

6.2 Sample Assessment

As we mentioned earlier, DDNs provide a flexible framework for reasoning under uncertainty. Given evidence on any subset of the random variables in our affective model, propagation algorithms compute the conditional probability of any other random variable in the model. Furthermore, if the agent needs to decide how to act at time t_{i+1}, the DDN computes the expected utility of every available action at that time and allows the agent to choose and execute the action with maximum expected utility.

We now give an example of how the propagation of available evidence allows our model in Figure 4 to incrementally refine the assessment on the user's emotional state as more relevant user data become available, thus providing the Prime Climb agent with increasingly accurate information to decide how to act in order to improve the user's interaction with the game.

Let's suppose that, at some point during the interaction with Prime Climb, the player falls and the agent decides to provide help. Let's also suppose that the only sensor signal available at this time comes from the heart rate monitor and indicates high heart rate. When this evidence is inserted in the model in Figure 4 and propagated, it increases the probability that the player's heart rate is high. High heart rate in turn increases the probability that the player is in an emotional state with negative rather than positive valence, because the conditional probabilities for the *Heart_Rate* node represent the finding that heart rate increases more in the presence of emotion with negative valence. Although the available evidence cannot discriminate between the player feeling reproach or shame, high probability of negative valence is sufficient to raise the probability that the player's engagement is low. At the next decision cycle, this probability may influence the model so that the agent's action with the highest expected utility is one designed to bring the level of engagement back up.

Let's now suppose that, in addition to high heart rate, we also detect high GSR. When propagated in the model, this evidence increases the probability of a high level of arousal and consequently the probability that our player is an introvert. This is because the CPT for *arousal* is set up to encode the finding that introverts reach a higher level of arousal than extraverts given the same stimuli. Although the resulting assessment does not add any information on the player's specific emotional state, it does give more information on the player's personality. At the next decision cycle, this information might result in having the action with maximum expected utility be one that deals specifically with overcoming a user's negative affective state when the user is an introvert (provided, of course, that such action is available to the agent).

---

[4] Other kinds of facial expressions are generally good indicators of valence, if not of individual emotions. In our sample model, eyebrow position contributes indirect information on valence through the *reproach* variable.

Lastly, if our sensors also detect that the user is frowning, the probability of the player feeling reproach rather than shame increases (because of the conditional probability representing the finding that frowning is a good indicator of emotions in the *anger* range). Indication that the player feels reproach also increases the probability that the player has the goal of succeeding by herself. This is because the conditional probabilities for *Reproach* give a high probability for this emotion if the player has the goal to succeed by herself and the agent provides unsolicited help (as it was the case in this example). Thus, in addition to giving an assessment of the user's emotional state, the DDN also assesses *why* the player is in that state. This information can further improve the capability of the decision model to select an adequate action. For instance, if the DDN assesses that the student feels reproach toward the agent because its interventions interfere with her goal to succeed by herself, the appropriate agent's behavior to revive the player's positive engagement in the game may be to refrain from giving further advice even if the student falls. A completely different cause of reproach toward the agent might be that the agent does not provide any help to a student that has the goal *Avoid Falling* but actually falls. A high probability for this particular configuration of the user's goal and emotion may influence the decision cycle so that providing help, not withdrawing it, is the action with the maximum expected utility.

Notice that the model would have generated a high probability of the user feeling reproach even if, instead of having evidence about the user frowning, it had evidence about the user having a disagreeable personality type (see top of Figure 2). This is because evidence of this personality type would increase the probability of having the goal *Succeed_by_Myself*, which is impaired by the agent's provision of help and therefore causes the user's reproach.

If contradictory evidence arises, such as evidence that the player has the goal to avoid falling but frowns when the agent provides help on how to recover from a fall, the model assessment of the user's affect will depend on the relative strength assigned to the different kinds of evidence by the model CPTs. However, in general the model probabilities will reflect a higher level of uncertainty on the user's emotional state. This also represents valuable information that the agent can use to decide how to act. The agent might decide, for instance, to explicitly ask the player how she is feeling or how she wants the agent to behave. Without a model of affect, explicit inquiries would be the only way the agent has to assess engagement, and might easily become annoying if they were too frequent. The model of affect allows the agent to explicitly interrogate the user only when the available evidence is insufficient to generate a reliable assessment.

6.3 Model Specification

One of the major difficulties in using probabilistic frameworks based on Bayesian networks is defining the required prior and conditional probabilities. In the model in Figure 4, the only prior probabilities to be specified are those for variables representing user traits, which can be defined through existing statistics, specialized tests, or set to indicate lack of specific information. The conditional probabilities for the model have been defined by the author to encode the general qualitative information available in the literature, and can be refined for our particular application and user population (students in grade 6 and 7) through empirical evaluations.

An alternative approach for building a model of affect that combines multiple sources of ambiguous evidence would be to specify heuristic rules to define how the available evidence should be integrated. But defining these rules still requires quantifying at some point complex probabilistic dependencies, because not explicitly using probabilities does not magically get rid of the uncertainty inherent to the modeling task. The advantage of a formal probabilistic approach is that the model designer only needs to quantify local dependencies among variables. The sound foundations of probability theory define how these dependencies are processed and affect the other variables in the model. In contrast, heuristic approaches require defining both the dependencies and ways to process them. This task is not necessarily simpler that defining conditional probabilities and entails a higher risk of building a model that generates unsound inferences. Furthermore, the DDN graphical representation provides a compact and clear description of *all* the dependencies that exist in the domain, given the direct conditional dependencies that the model designer has explicitly encoded. This helps to verify that the postulated conditional dependencies define a coherent model and to debug the model when it generates inaccurate assessments.

7. Related Work

Although affective user modeling is a field still in its infancy, an increasing number of researchers have started investigating the problem of how to make a software agent aware of a user's emotional state and able to react appropriately to it.

The work that is more closely related to what we propose in this paper is the probabilistic model described in (Ball and Breeze, 2000). This model relies on a Bayesian network to assess valence and arousal of user's affect, along with the dominance and friendliness aspects of a user's personality, during the interaction with an embodied conversational agent. The assessment relies on evidence from the user's linguistic behavior, vocal expression, posture and facial expressions, thus combining information from multiple bodily expressions to more accurately detect valence, arousal, dominance and friendliness. The main differences between the Ball and Breese's model and the model we propose in this paper are the following: (i) our model leverages evidence on the causes of emotional reactions in addition to evidence on bodily expressions; (ii) it explicitly represents the temporal evolution of emotional states; and, (ii) it allows assessing specific emotions in addition to valence and arousal, when sufficient evidence is available.

A substantial amount of research on how to use bodily expressions to assess a user's affect has been done at the MIT Medialab. Healy and Picard (2000) have used input from electromyogram, electrocardiogram, respiration and skin conductance sensors to detect stress in a car driver. Kaapor, Mota and Picard (2001) discuss how to monitor eyebrow movements and posture to provide evidence on students' engagement while they interact with a computer based tutor. Vyzas and Picard (1999) have shown how physiological data on jaw clenching, blood volume pressure, skin conductance and respiration can quite accurately recognize eight different emotional states, when a single subject intentionally expresses them.

Hudlicka and McNeese (2002) propose a framework that, like our model, combines information on relevant bodily expressions with other factors that can help assess a user's affect. They focus on identifying and combining factors to detect anxiety in combat pilots during a mission. These factors include general properties of the mission at hand (such as difficulty and risk level), events that happen during the mission (e.g., the detection of an enemy plane), pilot's traits (such as personality, experience and expertise) and real-time information on the pilot's heart rate. The framework includes heuristic fuzzy rules specifying the weight that each of the above factors has in predicting anxiety, as well as a mechanism to integrate the different factors. The framework also includes rules that specify how the pilot's level of anxiety affects his beliefs and performance, as well as strategies to counteract the possible negative effects of anxiety on performance.

Elliott, Lester and Rickel (1999) discuss how the Affective Reasoner, a rule-based framework to build agents that respond emotionally, could also be used to model user's affect. Like part of our DDN, the Affective Reasoner is based on the OCC cognitive theory of emotions, but relies on deterministic rules to model the appraisal process. Elliot et al., describe these rules in the context of assessing a student's affect during the interaction with the pedagogical agent for Design_a_Plant, a learning environment for botany. In their discussion, the authors assume that the user's goals and preferences necessary to define the outcome of the appraisal are known.

Although we are not aware of other user models designed specifically to assess emotional states in addition to cognitive states, both Del Soldato (1995) and de Vicente (2000) have developed tutoring systems that assess and try to enhance student motivation, a variable closely related to affective states. In both works, student motivation is assessed by comparing how the tutorial interaction relates to student traits that are known to influence motivation. These variables include degree of *control* that the student likes to have on the learning situation, degree of *challenge* that the student likes to experience, degree of *independence* during the interaction and degree of *fantasy based situations* that the student likes the instructional interaction to include. Murray and VanLehn (2000) developed a decision theoretic tutor that takes into account both student learning and morale in deciding how to act. However the authors do not discuss how student morale is assessed in their system.

Other researchers have been investigating the decision theoretic approach to guide the behavior of adaptive interactive systems. Mayo and Mitrovic (2001) apply decision theory to guide the actions of a computer-based tutor, solely based on student's learning. Horvitz (1999a; 1999b), presents intelligent desktop assistants that use a decision theoretic approach to decide when and how to provide unsolicited help to the user. Finally, Jameson et al. discuss how to apply decision theoretic methods to automatically provide the user with a sequence of tailored recommendations and instructions (Bohnenberger and Jameson, 2001; Jameson, Großmann-Hutter, March, Rummer, Bohnenberger and Wittig, 2001).

8. Conclusions and Future Work

We have presented a probabilistic model of a user's affect that integrates information on the possible causes of the user's emotional state (e.g., stimuli from the environment and personality traits) as well as the behavioral effects of this state (e.g., the user's bodily expressions). The model relies on a Dynamic Decision Network (DDN) to explicitly represent the probabilistic dependencies between causes, effects and emotional states, as well as their temporal evolution. By taking into account different kinds of possibly ambiguous evidence on the user's emotional state, our probabilistic model aims at reducing the uncertainty that pervades the assessment of user's affect in situations in which a variety of emotions can arise in relation to a variety of user's features.

We have shown how our model of user's affect can be used by decision-theoretic pedagogical agents designed to improve the effectiveness of educational games. In particular, we have described an instantiation of the model for the interaction with the pedagogical agent of Prime Climb, an educational game to help students learn number factorization.

The current version of our model DDN has been defined by relying on various theories and findings on the psychology and physiology of emotions. The part of the model that defines the dependencies between emotional states and possible causes is based on the OCC cognitive theory of emotions, which links emotional reactions to a person's goals, preferences and how they are matched by the current situation. We have integrated the basic structure suggested by the OCC theory with variables that provide indirect evidence on a person's goals, such as a player's personality and interaction patters. The part of the model that encodes the dependencies between emotional states and their observable effects has been defined by relying on existing findings on how emotions generate changes in one's bodily expressions and how these changes can be captured by specialized software and sensors.

We are currently working on refining the structure and conditional probabilities in the model with data derived from observations of players interacting with Prime Climb. We are especially interested in gathering more accurate statistics on the relations between players' goals, task knowledge and interaction behavior, as well as in understanding what bodily expressions are more easily detectable in this kind of interaction.

We also plan to investigate the issue of if and how emotional reactions influence the players' goals and situation appraisal. There is increasing evidence that affective states can impact performance by altering the perceptual and cognitive processes that define how a given situation is perceived, as well as the cognitive and motor skills that influence behavior selection and actuation. However, it appears that what these influences are is very much task dependent, and we currently have no clear sense of what role they play during the interaction with educational games. Our current intuition is that in educational games the influences of emotional states on situation appraisal may not be strong enough to warrant being explicitly represented in the affective model, but this intuition needs to be verified empirically.

Acknowledgments

We would like to thank the EGEMS group for designing and implementing the Prime Climb game, and for their support in running preliminary user studies on the game. We also thank Xiaoming Zhou for his valuable contributions to the definition and implementation of the affective model, and Giuseppe Carenini for his comments on an earlier version of this paper.

References

Ball, G. and J. Breeze (2000). Emotion and Personality in a Conversational Agent. *Embodied Conversational Agents*. J. Cassel, J. Sullivan, S. Prevost and E. Churchill, The MIT Press: 189-219.

Bianchi-Berthouze, N. and C. L. Lisetti (2002). Modeling Multimodal Expression of User's Affective Subjective Experience. *User Modeling and User-Adapted Interaction* 12(1): 49-84.

Bohnenberger, T. and A. Jameson (2001). When Policies Are Better Than Plans: Decision-Theoretic Planning of Recommendation Sequences. *IUI 2001: International Conference on Intelligent User Interfaces*, New York: ACM.

Cacioppo, J. T., G. C. Berntson, K. M. Poehlmann and T. A. Ito (2000). The psychophysiology of Emotions. *Handbook of Emotions (2nd edition)*. M. Lewis and J. M. Haviland-Jones, The Guilford Press.

Conati, C. and J. Fain Lehman (1993). EFH-Soar: Modeling education in highly interactive microworlds. *Lecture Notes in Artificial Intelligence*. New York, Springer-Verlag.

Costa, P. T. and R. R. McCrae (1992). Four ways five factors are basic. *Personality and Individual Differences 1* 13: 653-665.

de Vicente, A. and H. Pain (2000). A Computational Model of Affective Computational Dialogue. *AAAI Fall Symposium: Building Dialogue Systems for Tutorial Applications,* North Falmouth, MA., AAAI Press.

Dean, T. and K. Kanazawa (1989). A Model for Reasoning about Persistence and Causation. *Computational Intelligence* 5(3): 142-150.

Del Soldato, T. and B. du Boulay (1995). Implementation of motivational tactics in tutoring systems. *Journal of Artificial Intelligence in Education* 6(4): 337-378.

Ekman, P. (1993). Facial Expression and Emotion. *American Psychologist* 48(4): 384-392.

Ekman, P., R. V. Levenson and V. W. Friesen (1983). Autonomic nervous system activity distinguishes among emotions. *Science* 221: 1208-1210.

Elliott, C., J. Rickel and J. Lester (1999). Lifelike Pedagogical Agents and Affective Computing: An Exploratory Synthesis. *Artificial Intelligence Today , Lecture Notes in Computer Science 1600*, Springer Verlag.

Healy, J. and R. Picard (2000). SmartCar: Detecting Driver Stress. *15th International Conference on Pattern Recognition,* Barcelona, Spain.

Henrion, M., J. Breeze and E. Horvitz (1991). Decision Analysis and Expert Systems. *AI Magazine*, 12(4):64-91

Horvitz, E. (1999a). Attention-Sensitive Alerting. *UAI '99, Conference on Uncertainty and Artificial Intelligence,* San Francisco, CA, USA.

Horvitz, E. (1999b). Principles of mixed initiative interaction. *CHI '99, ACM SIGCHI Conference on Human Factors in Computing Systems,* Pittsburgh, PA.

Howard, R., A. and J. E. Matheson, Eds. (1977). Readings in Decision Analysis, Decision Analysis Group, SRI International, Menlo Park, California.

Hudlicka, E. and M. McNeese, D. (2002). Assessment of User Affective and Belief States for Interface Adaptation: Application to an Air Force Pilot Task. *User Modeling and User-Adapted Interaction* 12(1): 1-47.

Jameson, A., B. Großmann-Hutter, L. March, R. Rummer, T. Bohnenberger and F. Wittig (2001). When Actions Have Consequences: Empirically Based Decision Making for Intelligent User Interfaces. *Knowledge-Based Systems* 14: 75-92.

Kaapor, A., S. Mota and R. Picard (2001). Toward a learning companion that recognizes affect. *AAAI Fall Symposium: Emotional and Intelligent 2, the tangled knot of social cognition.*, AAAI Press.

Kahneman, D. (1973). Arousal and attention. *Attention and effort*. Englewood Cliffs, N.J., Prentice Hall**:** 2-49.

Klawe, M. (1998). When Does The Use Of Computer Games And Other Interactive Multimedia Software Help Students Learn Mathematics? *NCTM Standards 2000 Technology Conference,* Arlington, VA.

Malone, T. W. and M. R. Lepper (1987). Making learning fun: A taxonomy of intrinsic motivations for learning. *Aptitude, learning and instruction: Volume III Conative and affective process analyses*. R. E. Snow and M. J. Farr. Hillsdale, NJ, Erlbaum.

Matthews, G., D. Derryberry and G. J. Siegle (2000). Personality and emotion: Cognitive science perspectives. *Advances in personality psychology*. S. Hampson, E. (Ed.). London, Routledge.

Mayo, M. and A. Mitrovic (2001). Optimizing ITS behavior with Bayesian networks and decision theory. *International Journal of AI in Education* 12.

Mozziconacci, S. J. L. (2001). Modeling emotion and attitude in speech by means of perceptually based parameter values. *User Modeling and User-Adapted Interaction* 11: 297-326.

Murray, C. and K. VanLehn (2000). DT Tutor: A decision-theoretic dynamic approach for optimal selection of tutorial actions. *ITS 2000*, Montreal, Canada.

Murray, I. R. and J. L. Arnott (1993). Toward the Simulation of Emotion in Synthetic Speech: a review of the Literature on human vocal emotion. *Journal of the Acoustic Society of America* 93(2).

Ortony, A., G. L. Clore and A. Collins (1988). *The cognitive structure of emotions*, Cambridge University Press.

Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Mateo, CA, Morgan-Kaufmann.

Picard, R. (1997). *Affective Computing*. MIT Press.

Russell, S. and P. Norvig (1995). *Artificial Intelligence: A Modern Approach*. Los Altos, CA, Morgan-Kaufman.

Shute, V. J. (1993). A comparison of learning environments: All that glitters... *Computers as Cognitive Tools*. S. Lajoie, P. and S. Derry. Hillsdale, NJ, Lawrence Erlbaum Associates: 47-73.

Silvern, S. B. (1986). Classroom Use of Videogames. *Educational Research Quarterly*: 10(1):10-16.

Vyzas, E. and R. Picard (1999). Offline and Online Recognition of Emotion Expression from Physiological Data. *Workshop on Emotion-Based Agent Architectures, Third International Conference on Autonomous Agents,* Seattle, Wa.