A Probabilistic Framework For Recognizing and Affecting Emotions

Cristina Conati, Xiaoming Zhou

Department of Computer Science University of British Columbia,
Vancouver, BC, V6T1Z4 Canada
conati@cs.ubc.ca

Abstract

We present a framework for affective user modeling that deals with the high level of uncertainty involved in recognizing a variety of user emotions by relying un Dynamic Bayesian Network. We summarize how we used this framework to build a model of player affect to be used by a socially intelligent agent during the interaction with an educational game.

Introduction

One of the challenges in building emotionally intelligent interactive systems is recognizing user emotional states.

How detailed an understanding of a user's emotional state an agent should have depends on the task and the application. Sometime, detecting overall valence and level of arousal is sufficient to improve the quality of the interaction (Ball and Breeze 2000). In other situations, a single specific emotion is the factor that mostly influences user's performance (e.g., Hudlicka and McNeese 2002). However, there are interactions that trigger a variety of user emotions, and that can benefit from an agent who can detect these emotions and act upon them. Clearly, the more detailed the emotional information required, the more challenging the modeling task becomes.

The long-term goal of our research is to devise a framework for emotion modeling that can detect multiple specific emotions in addition to lower level information such as valence and arousal. We also want the framework to enable an agent to make principled decisions about how to react to and possibly influence the interlocutor's emotional states, given the agent's goals and the potential uncertainty of how agent actions affect user emotions.

Such comprehensive affective modeling has been addressed only at a very high level in general architectures for emotions. Elliot (1993) describes how other agents' emotions could be assessed through "several levels of defaults" in his architecture for emotion reasoning based on the OCC theory of emotions (Ortony et al. 1988). Gratch (2000) discusses how his plan-based emotion architecture could more accurately model others' emotions with the addition of plan-recognition functionalities. Marsella et al., (2000) mention how their framework for Interactive Pedagogical Drama can explicitly select agent actions to affect other agents' emotional states. However, the authors do not specify how they handle the uncertainty in the outcome of these emotion-affecting actions.

In this paper, we present an approach to comprehensive emotion modeling that explicitly deals with the uncertainty in emotion recognition by relying on probabilistic reasoning. In particular, we use Dynamic Decision Networks (Dean and Kanazawa 1989) to represent in a coherent framework the probabilistic dependencies among emotional states, their causes and their effect.

There are several reasons for using DDNs for modeling emotions. First, not only can DDNs generate as accurate an assessment as possible by leveraging any existing information on the user emotional state; they can also explicitly express the uncertainty of their predictions when limited information is available. Second, DDN allow representing the temporal evolution of emotion related variables at the same level as any other factor influencing the user's emotional states. Thus, they do not require to resort to ad hoc mechanisms to represent the dynamic of emotions, as it is done for example in (Gratch 2000). Third, they provide formal mechanisms based on decision theory to model how an agent can rationally chose among actions with uncertain outcomes.

In spite of these advantages, using any probabilistic framework based on Bayesian networks involves a very complex modeling task: specify the structure of the model and the associated prior and conditional probabilities.

We specify the structure and the parameters of our probabilistic model for emotion recognition by combining relevant findings from psychology with empirical data. We discuss our model in the context of the interaction with pedagogical agents designed to improve the effectiveness of computer-based educational games. However, the proposed approach is general and applies to a variety of tasks in which recognizing multiple specific emotions can greatly improve the interaction..

A Dynamic Decision Network for Emotion Recognition

Figure 1 shows two time slices of our DDN model of affect. The nodes in the figure represent classes of variables in the actual DDN. As the figure shows, the network can combine evidence on both causes and effects of emotional reactions, to compensate for the fact that often evidence on causes or effects alone is insufficient to accurately assess the user's emotional state.

The part of the network above the nodes *Emotional States* represents the relations between possible causes and emotional states, as they are described in the OCC model. To apply the OCC theory to emotion recognition, our DDN includes variables for goals that a user may have during the interaction with an intelligent agent, (nodes *Goals*[1] in Figure 1). The object of the user's appraisal is the outcome of any event caused by a user's or an agent's action (nodes *Student Action Outcome* and *Agent Action Outcome* in Figure 2) . Agent actions are represented as decision variables in the model indicating points where the agent decides how to intervene in the interaction. The desirability of an event in relation of the user's goals is represented by the node class *Goals Satisfied*, which in turn influences the user's *Emotional States*.

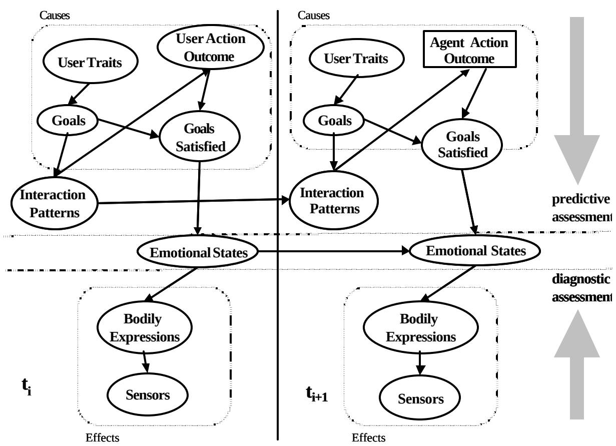

Figure 1: Two time slices of the DDN for affective modeling

User's goals are a key element of the OCC model, but assessing these goals is not trivial, especially when asking the user directly is not an option (as is the case in educational games, for example). Thus, our DDN also includes nodes to infer user goals from indirect evidence. User goals can depend on *User Traits* such as personality (Costa and McRae 1992) and knowledge. Also, user goals can influence user *Interaction Pattern*, which in turn can be inferred from the user individual actions. Thus, observations of both the relevant user traits and actions can provide the DDN with indirect evidence for assessing user goals.

In Figure 1, the links between nodes in different time slices indicate how the corresponding variables evolve over time. These links model, for example, the fact that a user is more likely to feel a given emotion at time t_{i+1} if the student felt it at time t_i. A new time slice is added to the network whenever either the user or the agent perform an action, but only the time slices that directly influence the current state need to be maintained in the model. We currently assume that maintaining two time slices is sufficient to capture the relevant temporal dependencies.

The part of the network below the nodes *Emotional States* represents the interaction between emotional states and their observable effects. *Emotional States* directly influence user *Bodily Expressions*, which in turn affect the output of *Sensors* that can detect them. Because in many situations a single sensor cannot reliably identify a specific emotional state, our model is designed to modularly combine any available sensor information, and gracefully degrade when the information becomes less reliable. In this paper, we focus on the causal part of the network. For details on the diagnostic part see (Conati et al. 2003). In the next section, we show how we are instantiating the general causal model described above to allow an intelligent pedagogical agent to model student's emotions during the interaction with the Prime Climb educational game.

Model Construction for the Interaction with an Educational game

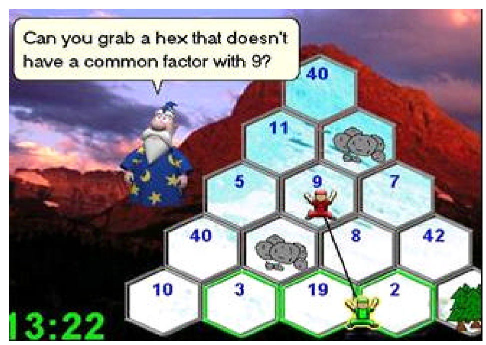

Figure 2: Prime Climb interface

Figure 2 shows a screenshot of the PrimeClimb game. The game is designed to teach number factorization. Two players must cooperate to climb a series of mountains. The mountains are divided in numbered sectors. Each player should move to a number that does not share any factors with his partner's number. Otherwise, the player falls and loses points. Prime Climb provides two tools to help students during climbing: a *magnifying glass* to see a number's factorization, and a *help box* to communicate with the pedagogical agent we are building for the game (see Figure 2). In addition to providing help when a student is playing with a partner, the pedagogical agent engages its player in a "Practice Climb" during which it climbs with the student as a climbing instructor. The affective model described here assesses student emotions during this phase.

[1] We currently represent players preferences in terms of goals, as suggested in (Elliot 1993)

We start by summarizing how we defined the sub-network that assess students' goals. For more details on the process see (Conati and Zhou 2002) and (Zhou and Conati 2003). We instantiated this sub-network with data collected through Wizard of Oz studies, during which an experimenter guided the pedagogical agent. Before the game, students were given a pretest on factorization knowledge, and a personality test based on the Five Factor personality theory (Costa and McRae 1992).

After the game, students filled a questionnaire on their goals in playing the game. The probabilistic dependencies among goals, personalities and interaction patters were established through correlation analysis between student answer to the goal questionnaire and the personality test, and between goal questionnaire and student actions in the log files recorded during the interactions. The correlations we found between goals and personality are either consistent with the description of the different personality

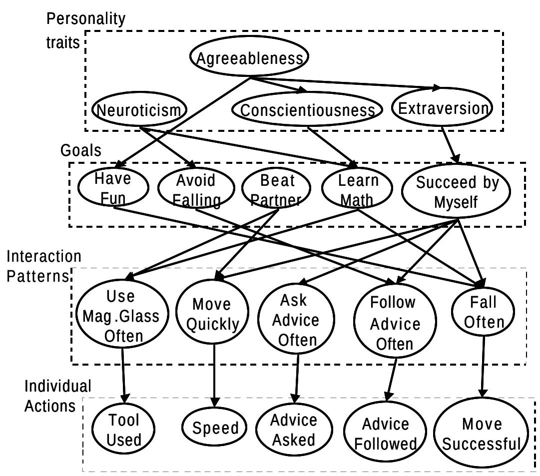

Figure 3: sub-network for goal assessment

domains in the 5 factor model or can be explained through the correlations existing between these domains (Graziano et al. 1997) . Figure 3 shows the resulting sub-network.

The bottom level of this sub network specifies how student interaction patters are recognized from sequences of individual actions. For instance, evidence of the student having a tendency to follow advice often comes from seeing a frequent number of instances in which the student follows the agent advice (Zhou and Conati 2003).

No significant correlation was found between the students' personality traits and the goal *Beat Partner*, despite the fact that the Five Factor model definition of *Agreeableness* would support a negative correlation. A possible explanation for this result is that, in the study, the students were playing not with a peer but with the agent acting as a climbing instructor, who played cooperatively and often provided help. Thus, even students with a more competitive personality might have felt compelled to cooperate with the agent.

Let's consider now the part of the network that represent the appraisal mechanism. We currently represent in our DDN only 6 of the 22 emotions defined in the OCC model: *joy* /*distress* for the current state of the game, *pride*/*shame* of the student toward herself, and *admiration*/*reproach* toward the agent. The three pairs of emotions are represented in the network by three two-valued nodes (see Figure 4), *emotion_for_event*, *emotion_for_self* and *emotion_for_agent*.

The links between *Goal* nodes, the outcome of student or agent actions and the *Goal Satisfied* nodes, as well as the corresponding conditional probabilities, are currently based on our subjective judgement. For some of these links, the connections are quite obvious. For instance, if the student has the goal of avoid falling and the outcome of her last action was a fall, this outcome will lower the probability that the goal is achieved. Other links are less obvious. For instance, what causes a student to have fun during game playing highly varies from student to student. We plan to refine these links through further user studies designed to obtain the actual student emotional states during the interaction.

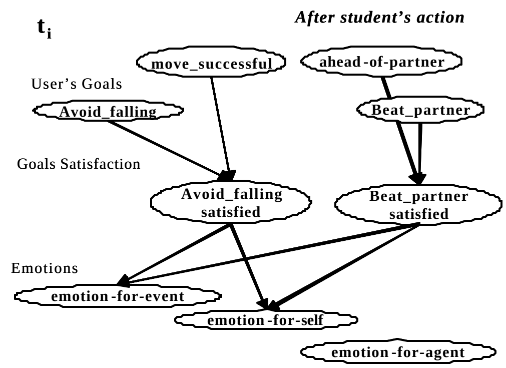

Figure 4: Sample sub-network for appraisal

The links between *Goal Satisfied* nodes and the emotion nodes are defined as follows. Because we assume that the outcome of every agent or student action is subject to student appraisal, every *Goal Satisfied* node in a given time slices influences *emotion-for-event* in that slice (see Figure 4, both slices). Whether a *Goal Satisfied* node influences *emotion-for-self* or *emotion-for-agent* in a given slice depends upon whether it was the student (slice t_i in figure 4) or the agent (example not shown for lack of space) who caused the current game state. The conditional probabilities of emotion nodes given goal-satisfied nodes are defined so that the probability of the positive emotion is proportional to the number of *Goal Satisfied* nodes that are true.

To validate the current version of the model, and to refine the links between action outcome and goal satisfaction, we are running a new user study with a version of prime Climb augmented with a mechanism that allows students to periodically input their emotions as they are playing the game. The model accuracy will be analyzed by comparing the model predictions with the student specified emotions.

Discussion

We have presented a probabilistic framework for emotion modeling designed to allow an intelligent agent to recognize multiple user emotions. Because the framework is based on DDNs, it can be also extended to allow the agent to rationally decide how to react to user emotions. We argue that a formal probabilistic approach like DDNs is called for to deal with the high level of uncertainty involved in reasoning about others' emotions in fairly unconstrained environments. Fuzzy rules have been proposed as an alternative approach to deal with this uncertainty. However, because this approach has so far been applied only to recognize one anxiety in combat pilots (Hudlicka and McNeese 2002), it is not clear how it would scale up to deal with the additional uncertainty involved in recognizing multiple emotions.

We acknowledge that, like any other formalism, DDNs have several limitations. One often mentioned is having to quantify probabilities in a more precise manner than with more heuristic approaches. We are addressing this issue empirically, by setting the network parameter from data. This is possible only when the network variables are observable. Thus, we are trying to limit the number of unobservable variables in our model. For example, we represent only high level user goals, instead a more detailed belief system as proposed by (Carofiglio et at., in press) in a model also built on Dynamic networks. While this model provides a more accurate description of a user 's mental state, it may be more problematic to specify the network's conditional probabilities because of the high number of unobservable variables.

A second limitation of DDNs is that they can become intractable as their structure grows. This is another reason to limit the grain size of our model, as well as the number of dependencies it explicitly represents. Although the DDN we described in this paper has satisfactory real time performance, it is an open question how it will scale up to model additional emotions (e.g. a student emotion toward the climbing partner in the two-player stage of the game) or dependencies (e.g. direct dependencies between personality and emotions, or emotions and goals).

There are additional questions that we would like to explore in future research. One is to what extent our probabilistic framework for emotion recognition can be extended or integrated into a more general architecture for emotion reasoning. For instance, our framework allows encoding how an agent *deliberate* actions affect user emotions, but these actions need to be predefined as alternative values of the DDN decision nodes. We would like to investigate how to extend our framework to include more sophisticated mechanisms for deliberate action generation, as well as mechanisms that allow to model the agent's instinctive reactions to the interaction. A second question is what role does probability play in generating rather than recognizing emotions. For instance, in his plan-based architecture Gratch (2000) uses probability of goal attainment to define emotion intensity. We plan to explore whether we can use this approach, combined with an assessment of goal importance, to define in our framework emotion intensity for both the agent and the user.

References

Ball, G. and J. Breeze, (2000) *Emotion and Personality in a Conversational Agent*, in *Embodied Conversational Agents*, J. Cassel, *et al*., Editors., The MIT Press.

Carofiglio, V., F. de Rosis, and R. Grassano, (in press) *Dynamic models of mixed emotion activation*., in *Animating expressive characters for social interactions*, Canamero L. and R. Aylett, Editors. John Benjamins.

Conati, C., (2002) *Probabilistic Assessment of User's Emotions in Educational Games*. Journal of Applied Artificial Intelligent, 16 (7-8)

Conati, C., R. Chabbal, and H. MacLaren. (2003) *A study on using biometric sensors for minitoring user emotions in educational games*. in *Workshop on Modeling User Affect and Actions: Why, When and How. At UM '03, Int Conf. on User Modeling*. Johnstown, PA.

Conati, C. and X. Zhou. (2002) *Modeling Students' Emotions from Cognitive Appraisal in Educational Games*. in *ITS 2002, 6th Int Conf. on Intelligent Tutoring Systems*. Biarritz, France.

Costa, P.T. and R.R. McRae, (1992) *Four ways five factors are basic*. Personality and Individual Differences. **13**

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

Elliot, C. (1993) *Using the affective reasoner to support social simulations*. in *Proceedings of the Thirteenth Annual Joint Conf. on Artificial Intelligenc*. Chambery, France,: Morgan Kaufmann.

Gratch, J. (2000) *Emile: marshalling passions in training and education*. in *Agents 200: 4th Int Conf. on Autonomous Agents*. BArcellona, Spain.

Graziano, G.W., L. Jesen-Campbell, and J.Finch, (1997) *The Self as a mediator between personality and adjustment*. Journal of Personality and Social Psychology, (73)

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).

Marsella, S., W.L. Johnson, and C. LaBore. *Interactive Pedagogical Drama*. in *Agents 200: 4th Int Conf. on Autonomous Agents*. 2000. BArcellona, Spain.

Ortony, A., G. Clore, and A. Collins (1998) *The cognitive structure of emotions*: Cambridge University Press.

Zhou, X. and C. Conati (2003) *Inferring User Goals from Personality and Behavior in a Causal Model of User Affect*. in *IUI 2003, Int Conf. on Intelligent User Interfaces*. Maimi, FL, U.S.A.