

Inferring User Goals from Personality and Behavior in a Causal Model of User Affect

Xiaoming Zhou, Cristina Conati
Computer Science Department
University of British Columbia
2366 Main Mall
Vancouver, BC, V6T1Z4, Canada
+604-8224632
{xzhou,conati}@cs.ubc.ca

ABSTRACT

We present a probabilistic model, based on Dynamic Decision Networks, to assess user affect from possible causes of emotional arousal. The model relies on the OCC cognitive theory of emotions and is designed to assess student affect during the interaction with an educational game. A key element of applying the OCC theory to assess user affect is knowledge of user goals. Thus, in this paper we focus on describing how our model infers these goals from user personality traits and interaction behavior. In particular, we illustrate how we iteratively defined the structure and parameters for this part of the model by using both empirical data collected through Wizard of Oz experiments and relevant psychological findings.

Categories and Subject Descriptors

I.2.3 [Artificial Intelligence]: Deduction and Theorem Proving – *uncertainty, “fuzzy”, and probabilistic Reasoning*;
K.3.1 [Computers and Education]: Computer Uses in Education – *computer-managed instruction (CMI)*.

General Terms

Human Factors

Keywords

User Modeling, Affective Computing, Dynamic Decision Networks, Educational Games.

1. INTRODUCTION

As researchers seeking to devise intelligent interfaces began to realize the importance of emotions in attention, planning, learning, memory, and decision-making, there has been increasing interest in building emotionally intelligent interactive systems that can express emotions and respond intelligently to human emotions [18].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

IUI'03, January 12-15, 2003, Miami, Florida, USA.

Copyright 2003 ACM 1-58113-586-6/03/0001...\$5.00.

One of the challenges in building emotionally intelligent systems is recognizing user emotional states. Humans use different sources of information to assess a person's emotions, including causal information on both context and the person's relevant traits, as well as symptomatic information on the person's visible bodily reactions. However, this information is often incomplete and even contradictory, making emotion assessment a task riddled with uncertainty. A computer attempting to recognize users' emotions will inevitably face the same problem.

This paper describes an affective user model designed to assess a variety of student emotions during the interaction with an educational game, Prime Climb. The long term goal is for the model to be used by an emotionally intelligent pedagogical agent, to respond to and possibly regulate student emotions during the interaction with the game. The model is based on an existing cognitive theory of emotions, the OCC theory [17] and uses Dynamic Decision Networks (DDNs) [6] to combine different sources of possibly ambiguous information on the causes of a user's emotional state.

One of the contributions of our model is its ability to perform on-line recognition of a variety of specific emotions. To date, the only work on recognizing multiple emotions focuses on the off-line assessment of emotions generated on demand by a professional actress [19]. All other relevant work focuses either on detecting a specific emotion in a fairly restrictive situation (e.g., anxiety in combat pilots [14], stress in car drivers [12]) or dimensions of emotional reactions such as valence (whether the emotion is positive or negative) and level of arousal. For instance, Ball and Breeze [1] present a Bayesian model to assess valence and arousal of user affect, as well as the dominance and friendliness of the user's personality, from linguistic and vocal expression, posture and facial expressions.

A second contribution is toward assessing user affect in the context of a tutorial interaction. Although student emotions play a key role in the learning process, very little work has been done on affective student modeling. Murray and Vanlehn [16] present a decision theoretic tutor that takes into account both student learning and morale in deciding how to act, but the authors do not discuss how student morale is assessed. Kaapor, Mota, and Picard [15] present preliminary results on how to monitor eyebrow movements and posture to provide evidence on students' engagement during the interaction with a computer based tutor, but these results were yet to be integrated in a computational student model. Del Soldato and du Boulay [7] and de Vicente and

Pain [5] have developed tutoring systems to assess and enhance student motivation, a variable related but not congruent to affective states.

A third contribution of our affective model is that it gives the first probabilistic representation of the OCC cognitive theory of emotions [17]. The OCC theory (named after its developers, Ortony, Clore and Collins) formally defines how emotions causally arise from the interaction of one’s goals and preferences with given states of the world. Although this theory has been used as a basis for several computational models of emotions (e.g., [8], [9]), none of these models deals with the uncertainty inherent to applying this theory to emotion recognition. Much of this uncertainty arises both because users can have multiple and often conflicting goals, and because these goals can be hard to assess. Thus, in this paper we focus on describing how our affective model manages the additional ambiguity these factors introduce in the modeling task by relying on Dynamic Decision Networks [6].

In the rest of the paper, we first give a brief introduction to the OCC theory of emotions. Then, we describe Prime Climb, the educational game that is the testbed for this research. Next, we illustrate how we refined the goal inferring part of a preliminary affective model described in [3], through both data collected from a formal user study and relevant psychological findings.

2. THE OCC COGNITIVE THEORY OF EMOTIONS

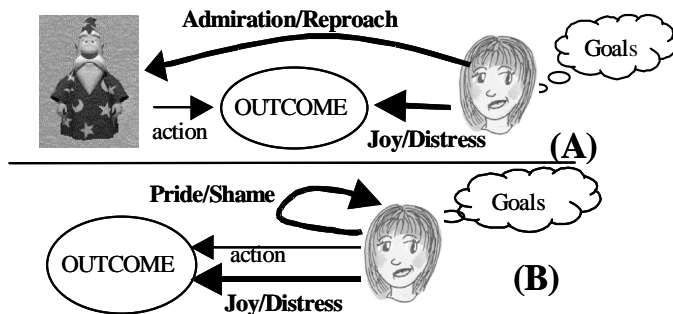


Figure 1: Example emotions in the OCC theory

According to the OCC cognitive theory of emotions [17], emotions derive from cognitive appraisal of the current situation consisting of events, agents, and objects. The outcome of the appraisal depends on how the situation fits with one’s goals and preferences. For instance, depending upon whether the current event (e.g., the outcome of an action in Figure 1) does or does not fit with one’s goal, that person will feel either *joy* or *distress* toward the event (see Figure 1, A and B); correspondingly, if the current event is caused by a third-party agent, that person will feel *admiration* or *reproach* toward the agent (see Figure 1A); if that agent is oneself, the person will feel either *pride* or *shame* (see Figure 1B). Based on this structure, the OCC theory defines 22 different emotions.

The OCC theory assumes that there is only one active goal during the cognitive appraisal process and, thus, the resulting emotional reaction is always deterministic. But, in reality, a person can have multiple goals or even conflicting goals (e.g., a student may have both the goals of having fun and to learn when playing an

educational game), and the goals might have different weights. The OCC theory does not describe how to deal with multiple and conflicting goals, nor how to assess these goals in the first place.

The affective model presented here applies the OCC theory to assess student emotions during the interaction with the educational game Prime Climb. As we will see in a later section, students do have multiple goals when playing this game, and these goals cannot be precisely assessed by questioning the students before they start playing the game. Thus, the model relies on DDNs to infer goals from a student’s interaction patterns and personality traits. It also models how multiple goals influence the student’s appraisal of the game interaction, and consequently the student’s emotions. Since most of the relations between student personality, interaction patterns, goals, game states and emotions are not deterministic, using DDNs allows our model to rely on the sound foundations of probability theory to cope with the resulting uncertainty. The model currently includes only the six emotions depicted in Figure 1 (*joy/distress*, *admiration/reproach*, and *pride/shame*) because they are the most relevant to our target domain. Before proceeding to describe the model, we describe the educational game, Prime Climb.

3. THE PRIME CLIMB EDUCATIONAL GAME




Figure 2: Part of the Prime Climb Interface

Figure 2 shows a screen shot of the Prime Climb interface. Prime Climb has been designed by the EGEMS group at the University of British Columbia to help 6th and 7th grade students practice number factorization. It is a two-player game, in which the players need to cooperate to get to the top of a series of mountains. The mountains are divided in numbered sectors. Each player can only move to a numbered sector that does not share any factors with the sector occupied by his partner. When a player makes a wrong move, she falls and starts swinging from the climbing rope. For example, in Figure 2 the player at the bottom fell because she tried to move to number 42, which shares factor 3 with number 9, occupied by the partner.

Prime Climb provides two tools to help students during climbing. The first is a *magnifying glass* the student can use to see the factor decomposition of any number on the mountain. The second is a help box that allows a player to communicate with the pedagogical agent we are building for the game (see agent on the

left in Figure 2). This tool is especially useful during a phase of Prime Climb called “Practice Climb”. During this phase, instead of playing with a peer, a student climbs with the pedagogical agent, which acts as a climbing instructor and provides help both unsolicited and on-demand. The affective model described here assesses student emotions during this phase. We currently focus on this phase because modeling student affect during the interaction with a peer would entail modeling emotions toward this peer in addition to the emotions toward the agent and the player himself. Given the many challenges that already exist in modeling the interaction with the agent only, we decided to address them first, before moving to a more complex interaction.

The main goal of the Prime Climb pedagogical agent is to provide tailored support to help students learn number factorization during climbing, while maintaining their level of engagement in the game. To achieve the latter, it is important that the agent recognize specific student emotions. For example, in Figure 2 the pedagogical agent is trying to help the player who just fell down the mountain. If the agent can monitor a student’s emotional state, it may decide not to provide this help to a player who has previously shown reproach for the agent’s interventions; or, it may opt to phrase the help differently to cheer up a player who seems to be ashamed of having made a mistake.

In all the empirical studies described in the rest of the paper, the pedagogical agent was controlled by a researcher through a Wizard of OZ interface.

4. THE PROBABILISTIC MODEL OF USER AFFECT

Figure 3 shows two time slices of a high level description of the DDN underlying our affective model for Prime Climb. A new time slice is created after every student’s or agent’s action that can influence the student’s emotions. The *Goals* nodes in Figure 3 represent high level goals that students may have when playing the game. The emotion eliciting situations represented in the model are events resulting from the outcome of either a student’s action (*Student Action Outcome* node in slice t_i) or the agent’s action (*Agent Action Outcome* node in slice t_{i+1}). The *Agent Action Outcome* node is represented as a decision variable in the model, indicating the different actions available to the agent. Because this paper focuses on the affective model, the utility node describing the agent’s preferences is not shown in Figure 3. The *Goals Satisfied* nodes in Figure 3 represent the outcome of the student’s situation appraisal, which is affected by both the student’s goals and the current situation. This appraisal in turn affects the student’s emotions (*Emotions for Outcome*, *Emotions for Self* and *Emotions for Agent* nodes in Figure 3), as specified by the OCC theory.

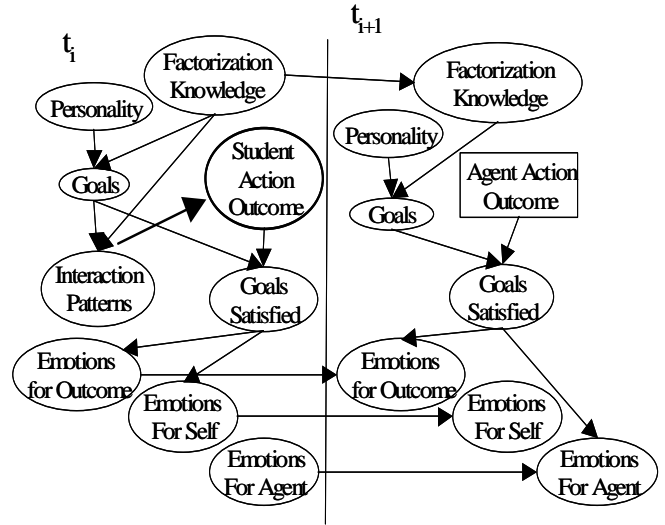


Figure 3: High level description of the affective model

Since students’ goals are a key element in the application of the OCC theory, the model also includes nodes to facilitate their assessment. Students’ goals can depend on both the students’ personality traits [4] and their number factorization knowledge, as indicated by the links from the *Personality* and *Factorization Knowledge* nodes to the *Goals* nodes. The links from *Goals* and *Factorization Knowledge* nodes to *Interaction Patterns* nodes in Figure 3 reflect the influence of the students’ goals and knowledge on their game behavior. Interaction patterns can be inferred from specific features of the students’ individual actions. These actions are dynamically added to the network as binary evidence nodes as the interaction proceeds.

The links across time slices between the *Factorization Knowledge* nodes and between *Emotion* nodes reflect the temporal dependencies for the corresponding student states. For instance, a satisfying event is more likely to make a student feel joy if the student already was in that emotional state. Our model currently includes the assumption that student high level goals do not change over time, as indicated by the lack of a link between the *Goals* node at t_i and the *Goals* node at t_{i+1} in Figure 3.

An initial detailed structure for the model in Figure 3 was defined by relying on two user studies we performed with grade six students in two Vancouver public schools. As we described in [3], these two studies (involving a total of 23 students) gave us the following information.

Student goals while playing the game. We identified five high level student goals: *Have Fun*, *Avoid Falling*, *Learn Math*, *Beat Partner*¹, and *Succeed By Myself*. We also found that students can have more than one of these goals at the same time.

Student interaction patterns. The interaction patterns we identified to be relevant for inferring student goals included: (1)

¹ This goal seems odd given that the two players should collaborate during climbing. However, it is consistent with findings showing that some personality types tend to be competitive even in collaborative situations.

tendency to make moves quickly or slowly; (2) tendency to use the magnifying glass often or not; (3) tendency to ask the agent for advice often or not; (4) tendency to follow the agent’s advice often or not; (5) tendency to fall often.

The studies also gave us an initial understanding of some of the relationships between goals and interaction patterns, described in [3] and shown in **Figure 4**. However, this understanding mostly relied on the experimenter’s subjective observations, because we recorded log files of the interaction in only one of the two studies (involving 10 students). Notice that **Figure 4** only shows the portion of the model that assesses student goals. For more details on the part of the model assessing emotions from goals see [3].

The studies also had additional limitations. First, they did not provide any information about the interaction between student goals and personality traits, because we were not able to administer a personality test to the participants. Thus, in the preliminary model we postulated an initial set of probabilistic dependencies (see Figure 4) by relying on the Five Factor Personality theory, which describes personality types using five domains: *Neuroticism*, *Extraversion*, *Openness*, *Agreeableness*, and *Conscientiousness* [4].

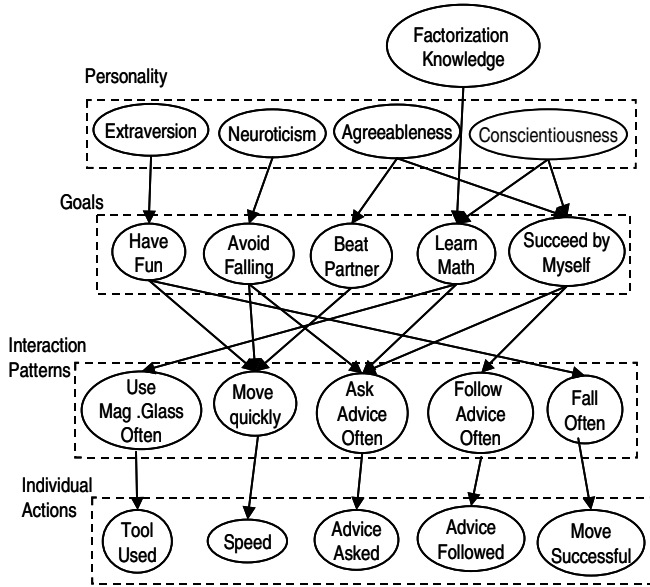


Figure 4: portion of preliminary model to assess goals

A second limitation of the studies was that, because of time limitations, we could only give students a very short pretest on factorization knowledge, including only three questions. This test could not give reliable information on how knowledge influences other variables in the model. Thus, the links from factorization knowledge in the preliminary model were also based on the researchers’ judgment (see Figure 4). Finally, the post-session questionnaire that we used to identify the students’ high level goals, was based on True/False questions. This limited the students’ choices and may have induced some students to exclude a goal solely because it had a much lower priority compared to the others.

To overcome the above limitations and obtain stronger quantitative information on the relevant relationships between

student personality traits, factorization knowledge, goals and interaction patterns, we ran a third user study, described in the next section.

5. A NEW STUDY TO REFINE THE AFFECTIVE MODEL

A total of 19 6th grade students volunteered for the study. Each student played for about 15 minutes with the pedagogical agent, who acted as a climbing instructor and was directed by an experimenter through the Wizard of Oz interface. Before the game, students were given a pretest including 15 questions on number factorization. After the game, the students filled out a questionnaire to assess their goals while playing the game. The questionnaire included questions to be answered using a scale of 1 to 5, and targeted the five goals identified in the previous studies. There were several questions targeting each goal, to increase the reliability of the students’ answers. There was also an open-ended question to see if the students had any other goal in addition to the five already established, but none were found.

To assess the participants’ personality, before the study we gave them a version of Goldberg’s 100 standard markers [10]. This test measures the five personality domains of the Five Factor theory and is suitable for children [11]. It consists of one hundred adjectives, twenty for each personality domain. The students were instructed to rate each adjective as an accurate description of themselves, from 1 (strongly disagree) to 5 (strongly agree). An observer recorded observations on the students’ interaction patterns. Log files were created for each student, to record relevant game events. Table 1 summarizes these events.

Table 1: Events captured in the study log files

Game event
Student move, including: <ul style="list-style-type: none"> – The time spent between two successive moves – Whether the move is successful or not – Whether the student is ahead of partner after the move
Student usage of the magnifying glass.
Student usage of the help box and the content typed.
Student responses to agent advice (i.e., whether the advice was followed)
Agent interventions, including: <ul style="list-style-type: none"> – Advice to use the magnifying glass – Advice to ask questions to the agent – Specific advice on where the student can move – Generic hint to reason about a fall or a successful move

6. STUDY RESULTS AND INFLUENCE ON THE MODEL

When sufficient data is available, a possible approach to define the structure of a Bayesian network is to use existing algorithms to learn the network from data (see for instance, [2]). Unfortunately, because only 19 students volunteered for the study, our data is insufficient to adopt this approach. Instead, we used a greedy approach. We relied on the data to incrementally generate a few variations of the model in Figure 4. At each step of this incremental process, we used the log marginal likelihood scoring method ([13]) to select the structure that best explained the data.

In particular, we first used the data from the post questionnaire and from the personality test to decide which dependencies to represent between student personality and goals. Next, we used the data from the post-questionnaire and from the log files to generate alternative sets of links from goals to interaction patterns. The rest of this section describes this process.

6.1 Relationships between Personality and Goals

To decide which dependencies between personality traits and goals to insert in the model, we ran a correlation analysis between the goal and personality scores from the study, and represented in the model only those correlations that were substantial and statistically significant. Table 2 summarizes the correlations we found. We now explain them in more detail, in relation to each goal, along with the changes they brought to the model.

Table 2: Correlations between personality and goals

	Extr.	Agr.	Consc.	Neur.	Open.
HaveFun	.107	.589*	.528*	-.097	.424
Succeed by Myself	-.646*	-.482*	-.314	.434	-.312
Learn Math	.650*	.634*	.550*	-.480*	.621*
Avoid Falling	-.315	-.008	-.087	.529*	-.120
BeatPartner	.154	.270	.262	-.317	.344

*: Correlation is significant at the 0.05 level (2-tailed).

Have-Fun. As Table 2 shows, only *Agreeableness* and *Conscientiousness* were significantly correlated to the goal *Have Fun*. The correlation between *Conscientiousness* and *Have Fun* may seem surprising, but it can be actually explained by the significant correlation between *Agreeableness* and *Conscientiousness* ($r = 0.810$, $p < 0.01$). This correlation, as well as all the other correlations we found between personality traits (discussed later in the section) was also found in [11]. Given these results, we need to remove the initially postulated link from *Extraversion* to *Have Fun*, and add links to represent the interactions between this goal and *Agreeableness* and *Conscientiousness*. Adding these links is not trivial, because various link configurations could be suitable representations for the correlations found (see Figure 5). We describe how we select the most appropriate configuration after we discuss all the other correlations we found.

Succeed By Myself. Both *Agreeableness* and *Extraversion* were found to be significantly and negatively correlated with this goal. The negative correlation with *Agreeableness* supports the hypothesis we embedded in the initial model because of the definition the Five Factor Model gives for this personality trait [4]: “sympathetic to others and eager to help them, and believes that others will be equally helpful in return”. The fact that an extravert is “sociable, liking people, and preferring large groups and gatherings” provides a plausible explanation for the negative correlation between *Extraversion* and *Succeed By Myself*. However, this correlation could also be explained by the positive correlation we found between *Extraversion* and *Agreeableness*. This generates the same problem of multiple plausible models illustrated in Figure 5.




Figure 5: Possible models for the correlations between Agreeableness (A), Conscientiousness (C) and HaveFun (HF)

Learn Math. All five personality traits were found to be significantly correlated with the goal *Learn Math* (see Table 2). This result is quite surprising, but it can be explained by the correlations existing between personality traits. These include, in addition to the ones discussed above, the correlation of *Openness* with every personality trait, and the negative correlation between *Neuroticism* and *Extraversion*

Avoid Falling. The only personality trait that is significantly correlated with the goal *Avoid Falling* is *Neuroticism*. This confirms the assumption we made in the first model – the more neurotic a student is, the more she is afraid of falling and making mistakes, because the *neurotic* personality is defined as “prone to worry, easy to become upset in stressful situations”.

Beat Partner. No significant correlation was found between the students’ personality traits and the goal *Beat Partner*. The initial model assumed that *Agreeableness* would be negatively correlated with *Beat Partner*, based on the Five Factor Model definition of this personality trait. A possible explanation for the lack of data support on this link 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 climbing instructor. A new study will be needed to clarify the relationship between personality and *Beat Partner*, when the model will be extended to assess emotions during regular playing between two peers.

We also found no statistically significant correlation between factorization knowledge and goals. Thus, we removed both the link from *Factorization Knowledge* to *Learn Math* in Figure 4, and the *Factorization Knowledge* node itself, since it could no longer provide evidence on student goals.

In order to find the best structure that explains all the correlations described above, we proceeded as follows. We first scored all the relevant models that could explain the correlations between Agreeableness, Extraversion, Conscientiousness, Have Fun and Succeed By Myself, using their log marginal likelihood ([13]). The two models that generated the highest score are shown in Figure 6. To check the sensitivity of this outcome to the data, we run cross validations as follows. We first generated 19 new data sets by eliminating from the original set one data point at a time. Then, we generated additional datasets by eliminating from the original set each pair and each triplet of different data points. Finally, we computed the log marginal likelihood of each new dataset given each structure. The two models in Figure 6 performed equally well over most datasets.




Figure 6: Highest scored structures for the personality traits influencing Have Fun and Succeed by Myself. Nodes are labeled with the initials of the corresponding variable names.

Since the two structures have the same score, we made the final selection between the two by considering their plausibility in terms of the Five Factor Model. This led to the selection of the structure to the left, because the direct connection between *Conscientiousness* and *Have Fun* in the model to the right is not obviously supported by the Five Factor Model.

From the selected structure, we then generated all the plausible models that include the additional correlations found between the personality traits and the goals *Learn Math* and *Avoid Falling*. Scoring these models resulted in the selection of the structure shown in Figure 7.




Figure 7: Final structure for the relations between personality and goals

We realize that this “greedy” approach is not completely sound, because the log marginal likelihood measure is not additive over network subparts. However, our 19 data points are not enough to reliably learn the structure of the complete network at once. Therefore, a greedy approach that also exploits the semantics of the variables involved is the only way to make sensible use of our data.

6.2 Relationships between Goals and Interaction Patterns

By analyzing the log files from the study, we calculated for each student the following scores related to the corresponding interaction patterns:

- number of magnifying glass and help box usages;
- percentage of agent’s advice followed, of moves resulting in a fall, and of quick moves. We used 4 seconds as a threshold to define a quick move, because the average moving time for all the students was 4.4 seconds and this was the threshold that best matched the judgments generated by the study observers.

Pearson correlation analysis was first tried to verify if there were any significant relationships between goal and interaction pattern scores. However, no statistically significant correlation was found, due most likely to the insufficient amount of data available.

We then decided to resort to simple cross tables to analyze the data. We converted goal and interaction pattern scores into two-valued categorical variables (using more than two values would provide a more accurate representation of the corresponding entities, but the resulting tables would have too few elements per cell to provide meaningful information). Next, we scored the contingency tables for all the possible two-way combinations of goals and interaction patterns, using the Fisher exact test (the only one applicable to contingency tables that, like ours, have many cells with less than 5 expected elements). The Fisher’s exact test returns a value between 0 and 1, where 1 indicates the virtual absence of association between two categorical variables, and 0 indicates the strongest possible association. Thus, we selected for inclusion in the model only links connecting goals and interaction patterns that had a contingency table with a Fisher’s score less than 0.4.

The conversion of each interactions pattern into a categorical variable with values T (showing the tendency described by the pattern) and F (not showing the tendency) was done by computing the average score for that pattern over all the study participants. Each participant was then categorized as showing that pattern if her actual score for it was above the average, and as not showing the pattern otherwise. For instance, the average of *Percentage of Quick Moves* for all the students was 60%. If more than 60% of a student’s moves were quick moves, the student was considered to tend to move quickly. Otherwise, the student was considered to move slowly.

Table 3: Contingencies between Goals, and Interaction Patterns with Fisher's score < 0.4

Goal	Interaction pattern	Fisher’s exact score
Have Fun	Fall Often	0.26
Learn Math	Fall Often	0.177
	Use Mag Glass	0.35
Avoid Falling	Follow Advice	0.304
Succeed By Myself	Fall Often	0.181
	Follow Advice	0.074
	Move Quickly	0.352
	Ask Advice	0.38
Beat Partner	Move Quickly	0.38
	Use Mag Glass	0.057

We also converted the goal scores into categorical variables with values High and Low, indicating whether the corresponding goals have high or low priority for a student. The score of each goal question in the post-questionnaire ranges from 1 to 5. Thus, if there are n questions related to a goal, $3*n$ is used as a threshold to convert the scores into the two categories. For example, since there are three questions assessing the goal *Have Fun*, a student was considered to have that goal if her post-questionnaire score for it was more than 9.

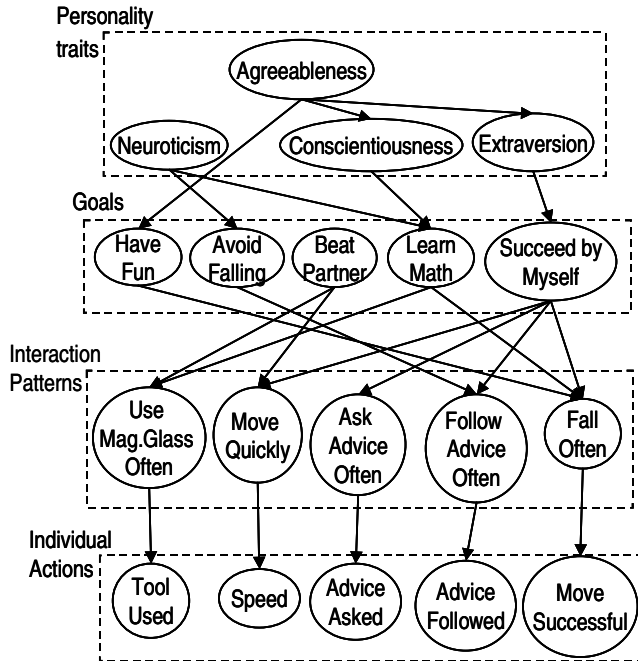


Figure 8: New structure for goal inference

Given the obtained Fisher's scores, only the links corresponding to the contingencies shown in Table 3 were considered as candidates for inclusion in the revised model. To verify that the new links actually generate a better model, we compared the log marginal likelihood of

- the new complete structure to assess student goals from personality and interaction patterns (see Figure 8), consisting of the structure in Figure 7 plus the links between goals and interaction patterns from Table 3.
- the initial model in Figure 4.
- an additional model we obtained from the new model by removing some of the links between goals and interaction patterns that seem less intuitive to us.

To test the reliability of the log marginal likelihood we ran cross validations as previously described. As shown in Figure 9, the new model with all the links suggested by the Fisher's test scores outperformed the other two models in most of the runs (Each run represents the computation of the log marginal likelihood with one dataset. Only the runs with datasets obtained by removing pairs of data points from the original set are shown in Figure 9).

7. THE MODEL CONDITIONAL PROBABILITIES

Since all the variables in the new portion of the affective model shown in Figure 8 are observable, the relevant conditional probabilities can be computed simply in terms of the relative frequencies in the data.




Figure 9: Cross validation of the three selected models

To obtain the CPTs for the goals and personality nodes, the personality scores were converted into categorical variables with T and F values, representing the two ends of each personality domain (e.g., extravert vs. introvert). Using three or more values for the personality nodes would be a more accurate representation of a student's personality, but it would also generate more sparse CPTs. Thus, we decided to start with two-valued nodes, and increase the number of values when more data will be available and if the model proves to be too inaccurate.

Binary evidence nodes representing student game actions are dynamically added to the model as the interaction proceeds. The CPTs between interaction patterns and individual actions (e.g., between the pattern Ask for Advice Often and an Ask for Advice action), were again computed by appropriate counting from the data.

Unfortunately, because of the limited amount of available data, using pure frequencies to generate the CPTs for the new model generates several sparse tables. As it is common practice in statistics, we currently address this problem by applying smoothing techniques.

8. CONCLUSIONS AND FUTURE WORK

In this paper, we have illustrated a causal model of user affect designed to infer a variety of user emotions during the interaction with an educational game. The model relies on the OCC cognitive theory of emotions. Following this theory, it represents in a Dynamic Decision Network how to infer user emotions from the evaluation of how the current game state fits with user goals. In addition, the model includes nodes and links to assess user goals from personality traits and interaction patterns. In the paper, we have focused on describing how we derived these nodes and links starting from an initial model built mostly from subjective observations of real interactions, and by refining this model using data from a formal user study.

The obvious next step for this work is to evaluate the accuracy of the model. This requires assessing the actual user affective states during the interaction with the educational game. We are currently working on an extension of the game interface that allows students to indicate their emotional states at various stages of the interaction. The challenge is to design this extension so that it elicits reliable student responses without itself becoming a factor

that affects student emotions. We are also working on implementing a decision theoretic pedagogical agent that uses this model to improve student interaction with the game. Part of this work involves adding to the affective model mechanisms to assess the evolution of student knowledge and how this knowledge interacts with student affect.

9. REFERENCES

1. Ball, G. and J. Breeze, *Emotion and Personality in a Conversational Agent*, in *Embodied Conversational Agents*, J. Cassel, et al., (Eds.). 2000, The MIT Press. p. 189-219.
2. Buntine, W., *A guide to the literature on learning graphical models*. IEEE Transactions on Knowledge and Data Engineering, 1996. **8**(195-210.).
3. Conati, C. and X. Zhou. *Modeling Students' Emotions from Cognitive Appraisal in Educational Games*. in ITS 2002, 6th International Conf on Intelligent Tutoring Systems. 2002. Biarritz, France.
4. Costa, P.T. and R.R. McCrae, *Four ways five factors are basic*. Personality and Individual Differences 1, 1992. **13**: p. 653-665.
5. de Vicente, A. and H. Pain. *Informing the detection of the students' motivational state: an empirical study*. in ITS 2002, 6th International Conf on Intelligent Tutoring Systems. 2002. Biarritz, France..
6. Dean, T. and K. Kanazawa, *A Model for Reasoning about Persistence and Causation*. Computational Intelligence, 1989. **5**(3): p. 142-150.
7. Del Soldato, T. and B. du Boulay, *Implementation of motivational tactics in tutoring systems*. Journal of Artificial Intelligence in Education, 1995. **6**(4)
8. Elliot, C. *Using the affective reasoner to support social simulations*. in Proc. of the 13th Annual Joint Conference on Artificial Intelligence. 1993. Chambery, France,; Morgan Kaufmann.
9. Elliott, C., J. Rickel, and J. Lester, *Lifelike Pedagogical Agents and Affective Computing: An Exploratory Synthesis*, in *Artificial Intelligence Today, Lecture Notes in Computer Science 1600*, M. Wooldridge and M. Veloso, (Eds.). 1999, Springer Verlag. p. 195-212.
10. Goldberg, L.R., *The development of markers of the Big Five factor structure*. Psychological Assessment, 1992. **4**: p. 26-42.
11. Graziano, W.G., L.A. Jensen-Campbell, and J.F. Finch, *The Self as a Mediator Between Personality and Adjustment*. Journal of Personality and Social Psychology, 1997. **73**: p. 392-404.
12. Healy, J. and R. Picard. *SmartCar: Detecting Driver Stress*. in 15th International Conf on Pattern Recognition. 2000. Barcelona, Spain.
13. Heckerman, D., *A tutorial on learning with Bayesian networks*, in *Learning in Graphical Models*, Jordan.M., (Ed.). 1998.
14. Hudlicka, E. and M. McNeese, *Assessment of User Affective and Belief States for Interface Adaptation: Application to an Air Force Pilot Task*. User Modeling and User-Adapted Interaction, 2002. **12**(1): p. 1-47.
15. Kaapor, A., S. Mota, and R. Picard. *Toward a learning companion that recognizes affect*. in AAAI Fall Symposium: Emotional and Intelligent 2, the tangled knot of social cognition. 2001: AAAI Press.
16. Murray, C. and K. VanLehn. *DT Tutor: A decision-theoretic dynamic approach for optimal selection of tutorial actions*. in ITS 2000. 2000. Montreal, Canada.
17. Ortony, A., G.L. Clore, and A. Collins, *The cognitive structure of emotions*. 1988: Cambridge University Press.
18. Picard, R., *Affective Computing*. 1997, Boston: MIT Press.
19. Vyzas, E. and R. Picard. *Offline and Online Recognition of Emotion Expression from Physiological Data*. in Workshop on Emotion-Based Agent Architectures, 3rd International Conference on Autonomous Agents. 1999. Seattle, Wa.