

Modeling Learner Affect with Theoretically Grounded Dynamic Bayesian Networks

Jennifer Sabourin, Bradford Mott, and James Lester

Department of Computer Science, North Carolina State University,
Raleigh, North Carolina, USA 27695
{jlobiso, bwmott, lester}@ncsu.edu

Abstract. Evidence of the strong relationship between learning and emotion has fueled recent work in modeling affective states in intelligent tutoring systems. Many of these models are based on general models of affect without a specific focus on learner emotions. This paper presents work that investigates the benefits of using theoretical models of learner emotions to guide the development of Bayesian networks for prediction of student affect. Predictive models are empirically learned from data acquired from 260 students interacting with the game-based learning environment, CRYSTAL ISLAND. Results indicate the benefits of using theoretical models of learner emotions to inform predictive models. The most successful model, a dynamic Bayesian network, also highlights the importance of temporal information in predicting learner emotions. This work demonstrates the benefits of basing predictive models of learner emotions on theoretical foundations and has implications for how these models may be used to validate theoretical models of emotion.

Keywords: Affective modeling, Intelligent tutoring systems, Dynamic Bayesian networks.

1 Introduction

Affect has begun to play an increasingly important role in intelligent tutoring systems. The intelligent tutoring systems community has seen the emergence of work on affective student modeling [1], detecting frustration and stress [2,3], modeling agents' emotional states [4,5], detecting student motivation [6], and diagnosing and adapting to student self-efficacy [7]. All of this work seeks to increase the fidelity with which affective and motivational processes are understood and utilized in intelligent tutoring systems in an effort to increase the effectiveness of tutorial interactions and, ultimately, learning.

This level of emphasis on affect is not surprising given the effects it has been shown to have on learning outcomes. Student affective states impact problem-solving strategies, the level of engagement exhibited by the student, and the degree to which he or she is motivated to continue with the learning process [8,9]. All of these factors have the potential to impact both how students learn immediately and their learning

behaviors in the future. Consequently, the ability to understand and model affective behaviors in learning environments has been a focus of recent work [1,10,11].

Correct prediction of students' affective states is an important first step in designing affect-sensitive learning systems. Knowledge of a student's current state is necessary to guide specialized feedback aimed at improving learning and motivation. However, the detection and modeling of affective behaviors in learning environments poses significant challenges. On the one hand, many current approaches to affect detection make use of a variety of physical sensors in order to make affective predictions (see Calvo et al. [12] for a review). Reliance on these types of sensors when building affect-sensitive learning environments severely limits how the systems can be delivered to students, reducing overall impact. On the other hand, systems that attempt to model emotion without the use of physiological sensors typically do so by incorporating theoretical models of emotion, such as appraisal theory, which is particularly well-suited for computational environments [4,13]. These models specify how individuals appraise events and actions along specific dimensions (e.g., desirability or cause) to arrive at emotional experiences. While there are a variety of appraisal-based theories of emotions, few models have been proposed that focus specifically on the emotions that typically occur during learning [9]. The lack of a widely accepted and validated model of learner emotions poses a challenge for the development of affect-detection systems using only contextual and goal-based features. This is especially true for learning environments where interpreting goals or measures of success or failure is non-trivial, such as exploratory environments or those focusing on ill-formed domains.

In this paper we investigate empirically derived models of student affect based on an appraisal-based theory of learner emotions that considers the different goals students may have during learning. Student self-reports of emotion were collected in an exploratory game-based learning environment, CRYSTAL ISLAND. Given the expected uncertainty of students' goals or appraisals, Bayesian techniques were used to develop models for prediction of student affect.

2 Background

As noted above, despite a large body of work, there is no single uniformly accepted theoretical model of emotion. However, appraisal theory (specifically the OCC model) has been typically favored by the affective computing community [9,12]. The OCC model proposes 22 emotions that occur as a result of an individual's appraisal of events, objects and the actions of others as well as oneself. Appraisal occurs across several dimensions including desirability, likelihood, control, and many others [13]. While this model has proven useful in several computing applications [1,4,5], it does not include many emotions that are believed to be important during learning situations [14,15].

Currently there are many theories and models of learner emotions, often called *achievement* emotions. Many models focus mainly on classifying emotions as they relate to the learning task. For example, the model proposed by Kort et al. considers four quadrants of emotions based on a dimension of learning and valence [8].

Alternatively, a model proposed by Csikszentmihalyi considers emotions along dimensions of individual skill and the challenge of the task, with high skill and high challenge corresponding to the optimal level of experience, a state termed *flow* [15]. While useful for classifying learner emotions, these theories do not take into account the many goals that students may have while learning.

Goal orientation is a term that has been used to describe a learner's primary focus when engaged in learning activities [16]. Students may either view learning in relation to performance or mastery. A performance approach would result in a student wishing to prove his competence and achieve better results than other students. A student with a mastery approach, however, views learning as an attempt to gain a skill, regardless of how her ability compares to others. This distinction between learning and performance goals forms the basis for the appraisal-based theory of learning emotions described by Elliot and Pekrun [17]. This model considers emotions in terms of learning and performance goals, along with evaluations of success and failure in these two categories. Additionally, they argue that certain individuals are more likely to focus on negative or positive valences of achievement emotions. For example, individuals with a positive (approach) disposition are more likely to experience positive feelings of enjoyment and pride, while those with negative (avoidance) dispositions are more likely to experience feelings of anxiety or shame [17]. This model of achievement emotions was used to inform the design of affect prediction models for the interactive learning environment, CRYSTAL ISLAND.

To date, many models of affect detection have been developed for use in computer-based learning environments. For instance, a model developed by D'Mello et al. considers facial expressions in terms of action units as well as students' posture and dialog acts to predict students' emotions as assessed by expert judges [11]. Similarly, Arroyo et al. have found benefit to multiple channels of physical evidence of affect [10]. By adding features such as facial expressions, skin conductivity, posture, and pressure they were able to account for much more variance over using contextual features of the tutoring environment alone. Conati and Maclaren have incorporated physical sensors into a complex model based on OCC theory [1]. Though they focus only on a subset of the emotions proposed by OCC they have used a dynamic Bayesian network to capture many of the complex phenomena associated with appraisal theories. This model estimates student goals based on personal traits and behaviors in the environment as well as evidence from physical feedback channels that further support the model's prediction. As in other environments the incorporation of physiological feedback data offered substantial improvement over models without this feature; however, the reliance on these sensors limits the ability to deploy the learning environments when sensors are unavailable or inappropriate.

The models explored in this paper extend previous work in the following ways. First, the models focus on incorporating features from learning-specific models of emotion in hopes of improving the accuracy of the predictive models. Second, the models are designed to achieve reasonable predictive accuracy without the use of physical sensors that would not be available during widespread distribution of the learning environment. Finally, the benefits of representing the dynamic nature of emotional experience will be demonstrated by comparing the performance of typical Bayesian networks with dynamic Bayesian networks.

3 Method

The predictive models of learner emotions were built using data from students' interactions with CRYSTAL ISLAND (Figure 1), a game-based learning environment being developed for the domain of microbiology that is aligned with the standard course of study for eighth grade science in North Carolina. CRYSTAL ISLAND features a science mystery set on a recently discovered volcanic island. Students play the role of the protagonist, Alex, who is attempting to discover the identity and source of an unidentified disease plaguing a newly established research station. The story opens by introducing the student to the island and the members of the research team for which her father serves as the lead scientist. As members of the research team fall ill, it is her task to discover the cause and the specific source of the outbreak. Typical game play involves navigating the island, manipulating objects, taking notes, viewing posters, operating lab equipment, and talking with non-player characters to gather clues about the disease's source. To progress through the mystery, a student must explore the world and interact with other characters while forming questions, generating hypotheses, collecting data, and testing hypotheses.

In order to empirically build and validate models of student affect, data from a study involving 296 eighth grade students from a rural North Carolina middle school was collected. After removing instances with incomplete data or logging errors, the remaining corpus included data from 260 students.

Pre-study materials were completed during the week prior to interacting with CRYSTAL ISLAND. The pre-study materials included a demographic survey, researcher-generated CRYSTAL ISLAND curriculum test, and several personality questionnaires. *Personality* was measured using the Big 5 Personality Questionnaire, which indexes student personality across five dimensions: openness, conscientiousness, extraversion, agreeableness and neuroticism [18]. *Goal orientation* was measured using a 2-dimensional taxonomy considering students' mastery or performance orientations along with their approach or avoidance tendencies [16]. Students' *affect regulation* tendencies were also measured using the Cognitive Emotion Regulation Questionnaire [19] though features from this survey were not included in the current models.

Students were given approximately 55 minutes to work on solving the mystery.



Figure 1. CRYSTAL ISLAND environment



Figure 2. Self-report device

Students' affect data was collected during the learning interactions through regular self-report prompts. Students were prompted every seven minutes to self-report their current mood and "status" through an in-game smartphone device (Figure 2). This report was described to students as being part of an experimental social network being developed for the island's research camp. Students selected one emotion from a set of seven options, which included *anxious*, *bored*, *confused*, *curious*, *excited*, *focused*, and *frustrated*. This set of cognitive-affective states is based on prior research identifying states that are relevant to learning [14, 17]. Each emotion label was accompanied by an emoticon to help illustrate the mood to students.¹ After selecting an emotion, students were instructed to type a few words about their current status in the game, similarly to how they might update their status in an online social network.

4 Results

In total, 1863 emotion self-reports were collected from 260 students, an average of 7.2 reports per student. These reports covered the range of available emotion choices with *focused* (22.4%) being the most frequent. Following this were reports of *curiosity* (18.6%), *frustration* (16.3%), *confusion* (16.1%), *excitement* (13.5%), *boredom* (8.5%) and *anxiety* (4.6%). Overall emotions with positive valence (*focused*, *curious*, and *excited*) accounted for 54.5% of emotion self-reports. These totals inform a baseline accuracy based on most frequent class against which the predictive models were compared: 22.4% for emotion prediction and 54.5% for valence prediction. These levels offer a more conservative estimate than a random model.

4.1 Predictive Modeling

Because of the inherent uncertainty in predicting student emotion, Bayesian networks were used to model the cognitive appraisal process. Bayesian networks are graphical models used to model processes under uncertainty by representing the relationship between variables in terms of a probability distribution [20]. In this study, each Bayesian network was specified using the GeNIe modeling environment developed by the Decision Systems Laboratory of the University of Pittsburgh (<http://dsl.sis.pitt.edu>). The variables and their dependencies were informed by the model of learner emotions described earlier. After the structure of the model had been specified, the parameters, or probability distributions of each dependency, were learned using the Expectation-Maximization (EM) algorithm provided within GeNIe. Each model was trained using 10-fold cross-validation, in which the model is trained on data from 90% of the students and is then tested for accuracy on the remaining 10%.

The models contained three types of variables:

¹ The emoticons were selected based on results from a validation study in which 18 graduate and undergraduate students rated the degree to which the emoticon images represented the desired emotional state.

- (1) **Personal Attributes.** These static attributes were taken directly from students' scores on the personal surveys prior to the interaction. Included were all four attributes for goal orientation and three personality attributes expected to be relevant to the student's appraisal: *conscientiousness*, *openness*, and *agreeableness*.
- (2) **Observable Environment Variables.** These dynamic attributes capture a snapshot of the student's activity in the learning environment up until the time of the self report. They provide a summary of important actions taken, such as *TestsRun*, *BooksViewed*, and *GoalsCompleted*. They also include information about how well the student is doing in the environment based on certain milestones, such as *SuccessfulTest* and *WorksheetChecks*.
- (3) **Appraisal Variables.** The values of the appraisal variables are not directly observable in the environment. Instead they are the result of the student's cognitive appraisal of many factors. The selected appraisal variables and their relation to observable variables are informed by the model of learner emotions.

4.2 Bayesian Networks

In order to provide an additional baseline of comparison, a naïve Bayesian network was learned. A naïve Bayesian network operates under the "naïve" assumption that all variables are directly related to the outcome variable but are conditionally independent of each other [20]. The learned naïve Bayesian network achieved a predictive accuracy of 18.1% on emotion label and 51.2% on valence. This performance is less accurate than the most frequent label baseline model, but provides an additional baseline measure. By comparing carefully constructed Bayesian networks against the naïve assumption we can determine the degree to which affective models benefit from theoretically informed structure.

Next, a Bayesian network (Figure 3) was designed with the structure informed by the proposed relationships described within Elliot and Pekrun's model of learner emotions. The design of the structure focused on the appraisal of learning and performance goals and how these goals were being met based on the status of the game environment. For example, learning-focused activities such as book reading or note-taking are expected to impact how much a student's learning goals are being met, while performance appraisals are more likely related to achieving important milestones such as running a successful test. Meanwhile, goal focus and valence tendencies are considered to be dependent on their personal attributes as described by the model. For example, students with approach orientations are expected to have generally more positive temperaments and emotional experiences than students with avoidance orientations. Similarly, personality traits such as agreeableness and openness are expected to contribute to an individual's overall temperament.

After the structure was designed, the parameters of the model were learned using the EM algorithm. Evaluation of the model showed that the Bayesian network could predict the emotion label with 25.5% accuracy and could predict the valence of the emotional state with 66.8% accuracy (Table 1). Both of these predictions offer a

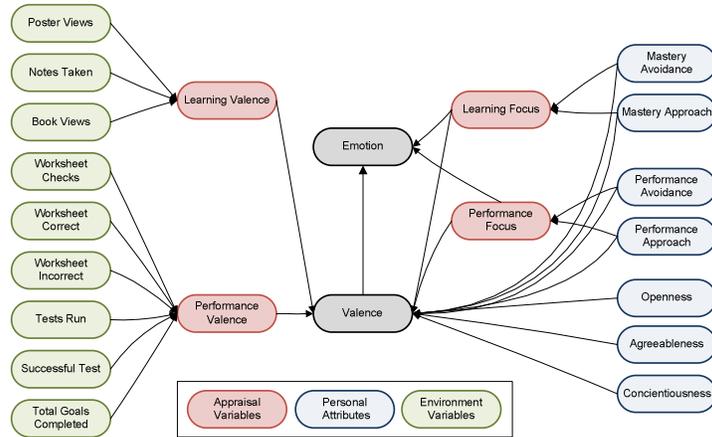


Figure 3. Structure of static Bayesian network

significant gain over the most frequent baseline and the naïve Bayesian network ($p < 0.05$). This improvement highlights the benefits of using a theoretical model of learner emotions to guide the model's structure.

However, the simple Bayesian network has no explicit representation of how emotions change over time. For instance, while poor performance at a task may merely be *frustrating* early in the interaction, for highly performance-oriented students this could turn into *anxiety* as more and more time passes. In order to capture the dynamic nature of emotions as they occur over time, the structure of the simple Bayesian network was used as the foundation of a series of dynamic Bayesian networks.

Dynamic Bayesian networks extend Bayesian networks by representing changes of the phenomena modeled over time. In this way, observations at time t_n are able to inform observations at time t_{n+1} [20]. A variety of representations of the dynamic

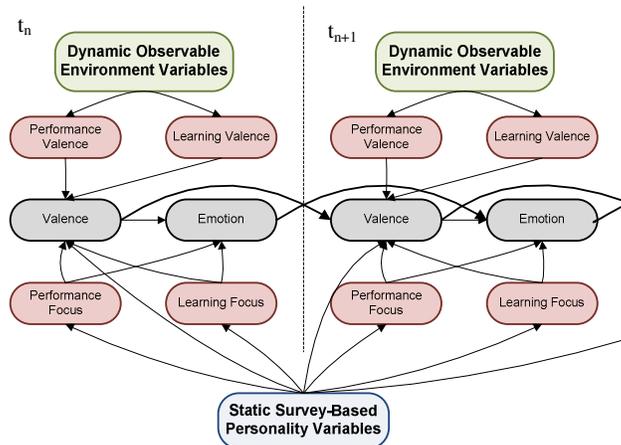


Figure 4. Structure of dynamic Bayesian network

Table 1. Predictive accuracies

	Emotion Accuracy	Valence Accuracy
Baseline	22.4%	54.5%
Naïve Bayes	18.1%	51.2%
Bayes Net	25.5%	66.8%
Dynamic BN	32.6%	72.6%

Table 2. Predictive accuracy by emotion

Actual Emotion	Correct Emotion Prediction	Correct Valence Prediction
anxious	2%	60%
bored	18%	75%
confused	32%	59%
curious	38%	85%
excited	19%	79%
focused	52%	81%
frustrated	28%	56%

Table 3. Valence confusion matrix

		Predicted Valence	
		Positive	Negative
Actual Valence	Positive	823	184
	Negative	326	512

nature of appraisal and the resulting affective states were tested. Of these, the model with the highest accuracy was able to predict emotional state with 32.6% accuracy and valence with 72.6% accuracy. This model (Figure 4) included a dynamic link between both emotion and valence, where the values of these two variables at t_{n+1} are partially informed by the emotion and valence at time t_n .

As expected, the predictive accuracy for *focused*, the most common self-report, was highest, with over half of the instances of *focused* being properly identified (Table 2). *Anxiety*, on the other hand, had the worst prediction with only 2% of instances being properly recognized. Positive affective states were recognized 81.7% of the time compared with 61.1% of negative affective states being correctly identified (Table 3). The predictive accuracies for specific emotions are particularly important for affect-sensitive learning systems that respond to detected emotions in light of recent work, which found that inappropriate responses can be detrimental to learners' emotional states [21].

5 Conclusion

This work presents Bayesian networks for predicting student affect with a structure informed by different models of learner emotions. By using empirical data to learn and validate the parameter of the models, it was found that the Elliot & Pekrun [17] model of learner emotions can successfully serve as the basis for computational models of learner emotions. The use of this model achieved performance beyond baseline measures as well as beyond a Bayesian network operating under a naïve assumption of the relationship of variables. This work also demonstrated a significant improvement in predictive power by extending the static model to a dynamic Bayesian network, which captured the changing nature of emotions across time. The models performed particularly well at recognizing positive emotional states. The negative affective states including *anxiety*, *boredom*, *confusion*, and *frustration* were often confused for each other. Improvements in predictions of these particular emotional states is an important area of future work for affective systems that intend to give feedback based on automated affect detection. Previous work has shown that experiencing these states has very different implications for students' behavior in the

environment [22] and how affective feedback may be received [21]. While the performance of the current models could likely be improved through the use of biofeedback sensors, developing reasonably accurate models that do not require invasive and expensive physical equipment is an important direction for providing affective support that can be used on a broad scale.

These findings suggest many interesting lines of investigation. For example, future work is needed to determine whether the proposed model performs well in other intelligent learning environments. Specifically, CRYSTAL ISLAND differs from most learning scenarios in that it is an open-ended exploratory environment without a clearly defined problem space. Additionally, it will be interesting to explore other theoretical models of learner emotions to compare how well these translate into computational models. This approach may help to validate theoretical models of learner emotions. Finally, a more comprehensive set of learning-focused cognitive and affective states could provide additional power to affect-sensitive systems. For example, the current predictive models do not consider a “neutral” affective state, nor do they offer a clear distinction between traditionally cognitive states (e.g. *focused*) and affective states (e.g. *excited*). Representations that distinguish these states may improve predictive and responsive capabilities.

The findings suggest that theoretical models of learner emotions can provide valuable guidance in designing cognitively focused affect detection models. By focusing on modeling the cognitive appraisal process without the support of physical biofeedback sensors, we are able to avoid the costs associated with distributing these sensors to future students and may achieve a greater audience for our educational systems. Future work is needed to investigate how these models may generalize to other learning systems and other populations. Additionally, it will be important to determine what level of accuracy is needed for predictive models to inform affective feedback and ultimately lead to improved learning and motivation.

Acknowledgments. The authors wish to thank members of the IntelliMedia Group for their assistance, Omer Sturlovich and Pavel Turzo for use of their 3D model libraries, and Valve Software for access to the Source™ engine and SDK. This research was supported by the National Science Foundation under Grants REC-0632450, DRL-0822200, IIS-0812291, and CNS-0739216. This material is based upon work supported under a National Science Foundation Graduate Research Fellowship.

References

1. Conati, C., Maclaren, H.: Empirically Building and Evaluating a Probabilistic Model of User Affect. *User Modeling and User-Adapted Interaction*, 19(3) 267-303 (2010)
2. Burleson, W.: Affective Learning Companions: Strategies for Empathetic Agents with Real-Time Multimodal Affective Sensing to Foster Meta-Cognitive and Meta-Affective Approaches to Learning, Motivation and Perseverance. PhD thesis, Massachusetts Institute of Technology (2006)
3. McQuiggan, S., Lee, S., Lester, J.: Early Prediction of Student Frustration. In: *Proc. of the 2nd Intl. Conf. on Affective Computing and Intelligent Interaction* (2007)

4. Marsella, S., Gratch, J.: EMA: A Process Model of Appraisal Dynamics. In: *Cognitive Systems Research* 10(1) 70-90 (2009)
5. Paiva, A., Dias, J., Sobral, D., Aylett, R., Sobrepez, P., Woods, S., Zoll, C., Hall, L.: Caring for Agents and Agents that Care: Building Empathetic Relations with Synthetic Agents. In: *Proc. of the 3rd Intl. Joint Conf. on Autonomous Agents and Multiagent Systems*, pp. 194-201 (2004)
6. de Vicente, A., Pain, H.: Informing the Detection of Students' Motivational State: An Empirical Study. In: *Proc. of the 6th Intl. Conf. on Intelligent Tutoring Systems*, pp. 933-943 (2002)
7. Beal, C., Lee, H.: Creating a Pedagogical Model That Uses Student Self Reports of Motivation and Mood to Adapt ITS Instruction. In: *AIED'05 Workshop on Motivation and Affect in Educational Software* (2005)
8. Kort, B., Reilly, R., Picard, R.: An Affective Model of Interplay Between Emotions and Learning: Reengineering Educational Pedagogy—Building a Learning Companion. In: *Proc. IEEE Intl. Conf. on Advanced Learning Technology: Issues, Achievements and Challenges*. Madison, WI: IEEE Computer Society (2001)
9. Picard, R., Papert, S., Bender, W., Blumberg, B., Breazeal, C., Cavallo, D., Machover, T., Resnick, M., Roy, D., Strohecker, C.: *Affective Learning – A Manifesto*. *BT Technology Journal*, 22(4) (2004)
10. Arroyo, I., Cooper, D., Bursleson, W., Woolf, B., Muldner, K., Christopherson, R.: Emotion Sensors Go to School. In: *Proc. of the 14th Intl. Conf. on Artificial Intelligence in Education*, pp. 17-24 (2009)
11. D'Mello, S., Graesser, A.: Multimodal Semi-Automated Affect Detection from Conversational Cues, Gross Body Language, and Facial Features. *User Modeling and User-Adapted Interaction*, 20(2) 147-187 (2010)
12. Calvo, R., D'Mello, S.: Affect Detection: An Interdisciplinary Review of Models, Methods and Their Applications. In: *IEEE Transactions on Affective Computing* 1, 18-37 (2010)
13. Ortony, A., Clore, G., Collins, A.: *The Cognitive Structure of Emotions*. Cambridge University Press. (1990)
14. Craig, S., Graesser, A., Sullins, J., Gholson, B.: Affect and Learning: An Exploratory Look Into the Role of Affect in Learning with AutoTutor. In: *Journal of Educational Media* 29, 241-250 (2004)
15. Csikszentmihalyi, M. *Finding Flow: The Psychology of Engagement with Everyday Life*. Basic Books (1997)
16. Elliot, A., McGregor, H., A 2x2 Achievement Goal Framework. In: *Journal of Personality and Social Psychology* 80(3), 501-519 (2001)
17. Elliot, A., Pekrun, R. (2007) Emotion in the Hierarchical Model of Approach-Avoidance Achievement Motivation. In: Schutz, P., Pekrun, R. (eds.) *Emotion in Education*, pp. 57-74. London: Elsevier (2007)
18. McCrae, R., Costa, P.: *Personality in Adulthood: A Five-Factor Theory Perspective* (2nd ed.) New York: Guilford Press (1993)
19. Gernefski, N, Kraaij, V.: Cognitive Emotion Regulation Questionnaire: Development of a Short 18-Item Version (CERQ-Short) In: *Personality and Individual Differences* 41, 1045-1053 (2006)
20. Russell, S., Norvig, P.: *Artificial Intelligence: A Modern Approach* (2nd ed.) Pearson (2003)
21. Robison, J., McQuiggan, S., Lester, J. Evaluating the Consequences of Affective Feedback in Intelligent Tutoring Systems. In: *Proc. of the Intl. Conf. on Affective Computing and Intelligent Interaction*, pp 37-42 (2009)
22. Baker, R.S., D'Mello, S.K., Rodrigo, M.M.T., Graesser, A.C.: Better to Be Frustrated than Bored: The Incidence, Persistence, and Impact of Learners' Cognitive-Affective States during Interactions with Three Different Computer-Based Learning Environments. *Intl. Journal of Human-Computer Studies*. 68(4) 223-241, (2010)