Predicting User Physiological Response for Interactive Environments: An Inductive Approach

Scott W. McQuiggan, Sunyoung Lee, and James C. Lester

Department of Computer Science
North Carolina State University
Raleigh, NC 27695, USA
{swmcquig, slee7, lester}@ncsu.edu

Abstract

Affective reasoning holds significant potential for interactive digital entertainment, education, and training. Incorporating affective reasoning into the decision-making capabilities of interactive environments could enable them to create customized experiences that are dynamically tailored to individual users' ever changing levels of engagement, interest, and emotional state. Because physiological responses are directly triggered by changes in affect, biofeedback data such as heart rate and galvanic skin response can be used to infer affective changes. However, biofeedback hardware is intrusive and cumbersome in deployed applications. This paper proposes an inductive framework for automatically learning models of users' physiological response from observations of user behaviors in interactive environments. These models can be used at runtime without biofeedback hardware to continuously predict users' physiological state directly from situational context in the interactive environment. Empirical studies with induced decision tree, naïve Bayes, and Bayesian Network physiological response models suggest that they may be sufficiently accurate for practical use.

Introduction

There is a growing demand for interactive technologies in entertainment, education, and training to create engaging experiences for increasingly sophisticated users. With the appearance of computational models of affective reasoning [6,8,14], we are now well positioned to investigate techniques that leverage the power of emotion to tailor the experience of individual users. Exploratory work in affective reasoning has studied the synthesis of emotion in animated agents [1,2,8,13] and models that inform social interaction [10,15,18]. Complementary efforts have begun to explore the automated recognition of user frustration [4], stress [19], interest [11], and motivation [3,21].

Affective reasoning offers much promise for interactive technologies. One can imagine an "affective barometer" [16] that could be incorporated into interactive environments to create customized experiences that are optimally stimulating and maintain ideal levels of engagement for individual users. Of course, the barometer would need to detect changes in affective state since users' emotions vary widely. For example, a player can be surprised, relaxed, fearful, frustrated, bored, and excited within a single gaming session. Changes in affective state are accompanied by physiological responses such as changes in heart rate, respiration, temperature, and perspiration [7]. Thus, detecting emotions is typically accomplished by "wiring" users, i.e., attaching biofeedback devices that monitor their physiological state changes. Although biofeedback has been demonstrated to accurately predict users' affective state [5,17], it is intrusive and cumbersome to use in practice. The limited feasibility of employing a biofeedback apparatus in a deployable system calls for an alternate approach to modeling users' physiological states that leaves them untethered. Because physiological state is so closely associated with affective state, an accurate model of physiological response could enable interactive environments to effectively reason about users' affective states, their level of stress, and their level of interest to craft customized interactions that are appropriately stimulating and engaging.

This paper presents an inductive approach to modeling users' physiological response in interactive environments. The Physiological Response Predictor (PRP) framework learns empirically informed models of physiological response from observations of user interactions in an interactive virtual environment. During training sessions, users are outfitted with biofeedback sensors (Figure 1). PRP monitors both *situational* data, including locational, intentional, and temporal information, and *physiological* data, including heart rate and galvanic skin response as users direct their characters to perform a sequence of tasks. PRP then induces models of physiological response from situational data. These models can be used without biofeedback hardware to continuously predict a user's physiological state directly from her situational context in the virtual environment. Empirical studies of PRP indicate that it can induce models of physiological response that appear to be sufficiently accurate for practical use.

Figure 1: Training user outfitted with biofeedback.

This paper is structured as follows. First it provides background on affective reasoning. Then it presents the PRP framework and includes a description of Treasure Hunt, the virtual environment testbed developed with the Valve Source™ engine (the game engine for Half-Life 2) in which PRP has been investigated. It next describes an evaluation of three PRP models (naïve Bayes, decision tree, Bayesian network) that were constructed from 20 users' interactions with the Treasure Hunt environment, and finally discusses design implications. Concluding remarks and directions for future work follow.

Affective Reasoning

Affective reasoning has been the subject of increasing attention among cognitive scientists in recent years, and the study of *affective computing* is becoming a field in its own right. Affective computing investigates techniques for enabling computers to recognize, model, understand, express and respond to emotion effectively. Affect influences humans' interactions with one another, their behaviors, and cognitive processes. Because of the central role played by emotions in decision making, perception, and learning, it is widely believed that affect can contribute in important ways to a broad range of computational tasks [16]. In particular, incorporating affective reasoning into digital entertainment, training, and education systems could enable them to create more effective, interesting, and engaging experiences for their users.

Foundational models of affect can be traced back to the pioneering work of Ortony, Clore, and Collins [14]. The OCC model supports twenty-two affective states that arise from valenced reactions (positive and negative) to appraised situations consisting of agents, events, and objects. The outcome of situations in the OCC model is a synthesized affective state. The OCC model was first implemented in the Affective Reasoner [6], which extended the model to include twenty-six affective states. Perhaps the only computational model of affect to equal the breadth of OCC is EMA [8], which is based on Smith and Lazarus' theory of emotion [12]. The Smith and Lazarus model centers on the cycle of event appraisal and coping strategies. EMA's implementation of Smith and Lazarus is used to guide virtual humans' behaviors in the Mission Rehearsal Exercise training environment [8] to increase their believability and realism.

The complementary processes of affect synthesis and affect recognition have been studied extensively in the context of animated agents. Work on *affect synthesis* has investigated the control of expressive models of embodied cognition and behavior in animated agents that support rich interaction in virtual environments [1,2,8,15] and pedagogical agents that support emotive expression in intelligent tutoring systems [10,13,18]. *Affect recognition* is the task of identifying the emotional state of an individual – frequently, this is the user interacting with a system – from a variety of physical cues, which are produced in response to affective changes in the individual. These include visually observable cues such as body and head posture, facial expressions, and posture, and changes in physiological signals such as heart rate, skin conductivity, temperature, and respiration [7]. Affect recognition work has explored emotion classification from self reports [3], post-hoc reports [21], physiological signals [5,19], and from combinations of visual cues and physiological signals [4]. This body of work serves as a springboard for the work described in this paper, which reports on techniques for recognizing users' physiological response without direct access to physiological signals in the runtime environment.

Modeling Physiological Response Prediction

Because users' physiological responses follow directly from their affective states, accurate models of physiological response could be used to enable interactive environments to effectively determine users' level of interest, stress, and emotion in order to guide customized interactions. However, for practical purposes, these determinations should ideally be made without resorting to the use of invasive biofeedback devices. We propose the Physiological Response Prediction (PRP) framework (Figure 2) that first acquires situational and physiological data from gameplay and biofeedback hardware and then learns models of physiological response from this training. PRP's induced models can then be used at runtime to predict users' physiological response directly from situational data without requiring biofeedback monitoring.

The PRP framework operates in two modes, model induction and model usage. During model induction (represented in Figure 2 with dashed lines), PRP acquires training data and learns models of physiological response from training users interacting with the interactive environment. The training user is outfitted with biofeedback equipment which monitors her heart rate and galvanic skin response. Biofeedback signals are recorded in training logs via the interactive environment, which also records an event stream produced by the training users' behaviors in the environment. Together, the biofeedback

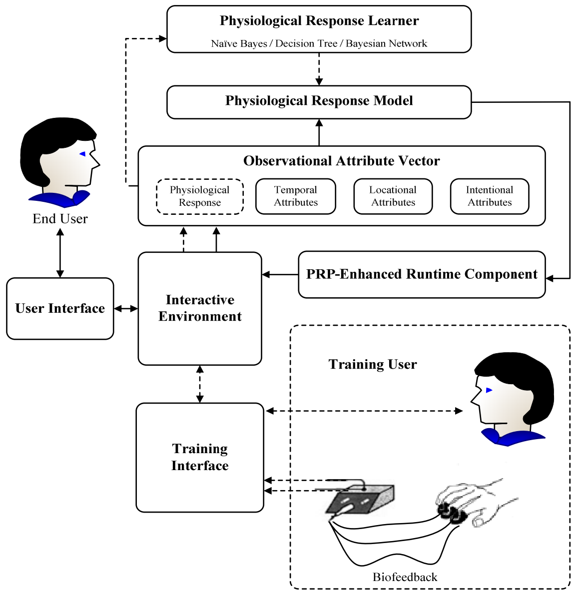

Figure 2: PRP Framework.

signals and the corresponding elements in the event stream are assembled in temporal order into the observational attribute vector. After training sessions (typically involving multiple training users) are complete, the physiological response learner induces models from the observed situational data and physiological data. The physiological data serves as class labels for the training instances. During model usage (represented in Figure 2 with solid lines), which is the mode used in runtime interactive environments to be deployed, the induced models inform the decision making of PRP-enhanced runtime components by predicting end users' physiological responses. Examples of candidate PRP-enhanced runtime components include NPC behavior controllers for games, narrative planners for interactive story worlds, and tutorial planners for intelligent tutoring systems. This paper focuses on PRP model induction, and the automated acquisition of training data for learning such models.

The Treasure Hunt Environment

The PRP framework has been studied in Treasure Hunt, a prototype virtual environment testbed (Figure 3) implemented with Valve's Source™ engine. The Treasure Hunt game world features a sub-tropical island once inhabited by pirates, who have left their treasures behind. Players begin by finding their character on the docks of the shores of the island and then explore the island by foot in a first-person point of view. Traveling on the beaches and through the water near the shore, they can venture into an enormous labyrinthine abandoned warehouse. Treasure boxes are scattered throughout the island. Some are visible, but most are concealed; some are hidden behind rocks on the beach, some are under the dock, and some are in various rooms in the warehouse behind oil drums and cargo containers. Each treasure box is labeled with the value of its contents, representing the points obtained by collecting the associated treasure. Players acquire the treasure by opening the boxes with a crowbar they find near the docks. Their quest in this testbed environment, which was designed specifically to study affective modeling, is to collect as much treasure as possible in the fixed amount of time (7 minutes) available to them.

Figure 3: Users' view of Treasure Hunt environment.

Training Data Acquisition

Accurately modeling physiological response requires a representation of the situational context that satisfies two requirements: it must be sufficiently rich to support assessment of physiological changes, and it must be encoded with features that are readily observable at runtime. The PRP framework therefore employs an expressive representation of all activities in the virtual environment, including those controlled by users and the interactive system, by encoding them in an observational attribute vector, which is used in both the model induction and model usage modes of operation. During model induction, the *observational attribute vector* is passed to the physiological response learner for model generation; during runtime operation, the attribute vector is monitored by a PRP-enhanced runtime component that utilizes knowledge of user physiological response to make its decisions. The observable attribute vector represents three interrelated categories of features for making decisions:

- *Temporal features*: The PRP framework continuously tracks the amount of time that has elapsed since the user arrived at the current location, since the user achieved a goal, and since the user was last presented with an opportunity to achieve a goal.
- *Locational features*: The PRP framework continuously tracks the location of the user's character. It monitors locations visited in the past, locations recently visited, locations not visited, and locations being approached.
- *Intentional features*: The PRP framework continuously tracks goals being attempted (as inferred from locational and temporal features, e.g., approaching a location where a goal can be achieved), goals achieved, the rate of goal achievement, and the effort expended to achieve a goal (as inferred from recent exploratory activities and locational features). These features enable models to incorporate knowledge of potential and user-perceived valence (positive and negative perceptions) of a given situation.

In the PRP implementation for Treasure Hunt, the observational attribute vector encodes 203 features. During model induction, a continuous stream of physiological data is collected and logged approximately 30 times per second. In addition, an instance of the observational attribute vector is logged every time a significant event occurs, yielding, on average, hundreds of vector instances each minute. At runtime, the same features are continuously monitored by the interactive environment.

Learning PRP Models

During PRP model induction, the framework learns models of physiological response from the attribute vectors. Many types of models can be learned. Work to date has investigated two families: rule-based models (decision trees) and probabilistic models (naïve Bayes and Bayesian networks). Naïve Bayes and decision tree classifiers are effective machine learning techniques for generating preliminary predictive models. Naïve Bayes classification approaches produce probability tables that can be implemented into runtime systems and used to continually update probabilities for predicting physiological responses. Decision trees provide interpretable rules that support runtime decision making. With both naïve Bayes and decision tree classifiers, PRP-enhanced runtime components can monitor the state of the attributes in the probability tables (for naïve Bayes) or rules (for decision trees) to determine when conditions are met for predicting particular physiological responses. Both naïve Bayes and decision tree classification techniques are useful for preliminary predictive model induction for large multidimensional data, such as the 122-attributes taken from the 203-observed attribute vector used for learning in the Treasure Hunt testbed. Two approaches can be distinguished in learning techniques: those that are completely automated, and those that require the knowledge provided by a domain expert. PRP experiments reported below focus on fully automated learning approaches. PRP model induction proceeds in four phases:

- **Data Construction**: Each training log is first translated into a full observational attribute vector. For example, blood volume pulse (BVP) and galvanic skin response (GSR) readings were taken nearly 30 times every second reflecting changes in both heart rate and skin conductivity. The 122 attributes observed directly in the environment were combined with the selected BVP and GSR class labels compressed from the full 203-element attribute vector for constructing the training sets.
- **Data Cleansing**: First, data are converted into an attribute vector format. Second, a dataset is generated that contains only instances in which the biofeedback equipment was able to successfully monitor BVP and GSR throughout the entire game session. For example, in the evaluation described below, data from one session out of twenty had to be removed for this reason: BVP (used for monitoring heart rate) readings were difficult to obtain from this participant.
- **Naïve Bayes Classifier and Decision Tree Learning**: Once the dataset is prepared, it is passed to the learning systems. The Treasure Hunt data was loaded into the WEKA machine learning tool [20], a naïve Bayes classifier and decision tree were learned, and tenfold cross-validation analyses were run on the resulting models. The entire dataset was used to generate several types of physiological response models. These included models of changes to both heart rate and galvanic skin response.
- **Bayesian Network Learning**: Decision tree results can then inform Bayesian network structure learning. In the Treasure Hunt PRP implementation, the top thirty-three attributes providing the most information were used to learn the structure of a Bayesian network.

Evaluation

Participants

In a formal evaluation, data was gathered from 20 subjects in an Institutional Review Board (IRB) of NCSU approved user study. There were 11 female and 9 male participants. Participants average age was 21.6 (SD = 2.96).

Procedure

First participants completed and reviewed pre-experiment materials consisting of a demographic survey, Half-Life 2 controls reference sheet, and a controlled backstory in preparation for interacting within the environment. The pre-experiment phase also contained a practice task from the game Half-Life 2 presenting an opportunity for training users to become familiar with the controls. Participants were then outfitted with a biofeedback apparatus for the experiment. The experiment consisted of two 3D Treasure Hunt virtual environments, each of varying degrees of difficultly. The easiest version of Treasure Hunt offered many opportunities to find treasures and meet the expectations that were set in the backstory. The most challenging version of Treasure Hunt made it difficult to find treasures; there were fewer treasures worth less value and more occluded treasure boxes making it difficult to meet backstory expectations. Participants collected as many treasures as possible in the allotted 7 minutes. The post-experiment materials consisted of a survey about the training user's experience and opinions on affect in applications such as games.

Results

All models were evaluated using a tenfold cross-validation scheme for producing training and testing datasets. In this scheme, data is decomposed into ten equal partitions, nine of which are used for training and one used for testing. The equal parts are swapped between training and testing sets until each partition has been used for both training and testing. Tenfold cross-validation is widely used for obtaining a sufficient estimate of error [20].

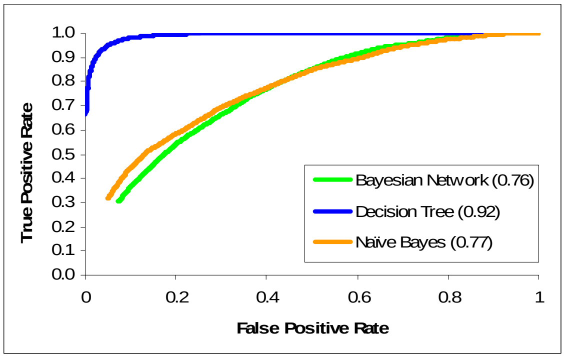

Figure 4: ROC curves of PRP induced models for heart rate change prediction. Areas under the curves are in parentheses.

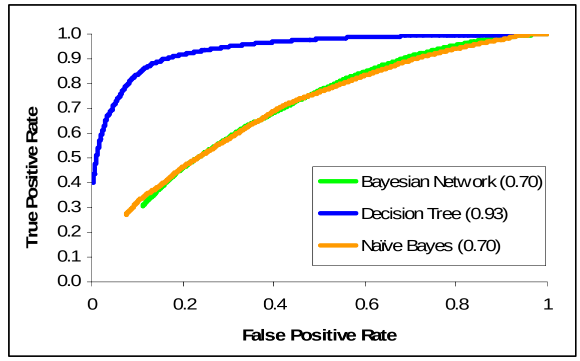

Figure 5: ROC curves of PRP induced models for galvanic skin response prediction. Areas under the curves are in parentheses.

Cross-validated ROC (Receiver Operating Characteristic) curves are useful for presenting the performance of classification algorithms for two reasons. First, they represent the positive classifications (true positives), included in a sample, as a percentage of the total number of positives along the vertical axis, against the negative classifications (false positives) as a percentage of the total number of negatives [20]. Second, the area under ROC curves is widely been accepted as a generalization of the measure of the probability of correctly classifying an instance [9].

The ROC curves (Figure 4 and 5) show the results of Bayesian networks, decision tree and naïve Bayes model results for predicting physiological response. The smoothness of the curves indicates that sufficient quantities of data were used to induce PRP models and that there was adequate coverage of possible instances. The highest performing models induced from Treasure Hunt training data was the decision tree models in both classification of heart rate and galvanic skin response, accurately predicting more than 90% of changes in physiological response.

Discussion

The results of the experiment suggest that the PRP framework can support the automated induction of accurate physiological response prediction models. It is interesting that the accuracy levels were surprisingly high, both for heart rate and GSR (changes in skin conductivity) prediction. The naïve Bayes and Bayesian network models both performed reasonably well, and the decision tree model performed particularly well. It seems that the high performance of the decision tree classifier was perhaps influenced by the fact that the data available for learning was voluminous. Decision trees seem to perform well on tasks that can furnish very large data sets, and physiological response modeling is such a task. Drawing strong conclusions such as, "Decision trees are superior to naïve Bayes or Bayesian networks for physiological response prediction," is not supported by a single experiment such as the one reported here. In the Treasure Hunt experiment, both the conditional probabilities and the structure of the Bayesian network were learned automatically from the data. The results were similar to those for naïve Bayes, but it seems possible with the incorporation of additional domain knowledge, a Bayesian network enhanced by a domain expert might achieve stronger results.[1]

In the study, binary labels (*down* and *up*) were induced. Thus, given an environmental context, the model predicts whether the user's heart rate will increase or decrease and whether the user's skin conductivity will increase or decrease within a predefined temporal window. Binary labels were chosen because of their simplicity and because of the potential increases in prediction accuracy that might arise from reduced granularity. One can imagine fine-grained labels, e.g., (*down*, *stable*, *up*) or (*large decrease*, *small decrease*, *no change*, *small increase*, *large increase*).[2] It could be argued that fine-grained physiological reaction models could be beneficial. However, it appears that a PRP-enhanced runtime component can function effectively with coarse grained class labels because predictions will be made continuously and it is the trends in physiological change rather than an isolated prediction that seem to be the most informative for decision making. For example, a sequence of 35 *up* values in heart rate clearly suggests that the user is experiencing an increase in stimulation.

Inspection of the data reveals that there were many repeated instances in the data, i.e., multiple users reacted similarly in identical situations. In the experiment reported here, although the interaction data were gathered from 20 different training users, and they were permitted flexibility in exploring the environment, there is significant regularity in the training data. This training data regularity stems from regularity in the environment itself, regularity in how users explored and interacted with the environment, and regularity with how users reacted physiologically to events in the environment.

[1] The authors also investigated Bayesian networks enhanced with hidden variables suggested by a domain expert. Because of technical limitations associated with learning conditional probabilities for Bayes nets with real-valued variables and hidden variables, success with "authored" Bayesian networks for physiological response prediction has to date been limited.

[2] In fact, another set of experiments were run in which ternary rather than binary class labels were used. Accuracy rates were only slightly less than those for binary class labels, e.g., 88% for the decision tree GSR model.

One possible explanation for the results is that physiological response patterns are common across users, i.e., there are inherent physiological similarities between people. If this is true, then physiological response learning approaches such as that embodied by the PRP framework are promising. An alternative (and complementary) explanation is that lower levels of predictability might hold for a more diverse set of training users with virtual environments that are more expansive, complex and dynamic. This could well be the case and is an interesting direction for future work. Nonetheless, it seems likely that to a great extent, people react in predictable patterns that can be automatically learned from observation.

Conclusion

Dynamically crafting interactive experiences that are highly customized for individual users is a long-term goal of digital entertainment, education, and training. If interactive environments can be given the ability to accurately reason about users' affective characteristics, they can tailor user's experiences to moment-by-moment changes in their levels of engagement, interest, stress, motivation, and emotional state. Physiological responses follow directly from changes in affect and thus can be used as key predictors of affective state. Although biofeedback devices can be used to obtain actual physiological signals, it may be impractical to require users to don biofeedback equipment and deploy additional hardware with applications. The PRP framework can automatically induce models of users' physiological responses that can predict physiological changes from observable events in interactive environments. It appears that the accuracy of PRP induced models may be sufficiently high to facilitate the control of users' experience in the virtual worlds that increasingly dominate gaming, education, and training.

This works represents a first step toward deployable affective reasoning for interactive environments. It will be interesting to explore the predictive capabilities of induced models in more complex, dynamic environments. Will different types of models be required? Can complementary affective constructs such as immersion, confidence, and fantasy be induced from observable elements in these environments? How can models such as these be incorporated into next-generation entertainment and learning environments? These questions suggest important directions for future work.

References

[1] André, E., and Mueller, M. 2003. Learning affective behavior. In *10th International Conference on Human-Computer Interaction*, 512-516.

[2] Bates, J. 1994. The role of emotion in believable agents. Technical Report CMU-CS-94-136, Carnegie Mellon University, Pittsburgh, PA.

[3] Beal, C., and Lee, H. 2005. Creating a pedagogical model that uses student self reports of motivation and mood to adapt ITS instruction. *Workshop on Motivation and Affect in Educational Software*. *12th International Conference on Artificial Intelligence in Education*.

[4] Burleson, W., and Picard, R. 2004. Affective agents: Sustaining motivation to learn through failure and a state of stuck. *Workshop of Social and Emotional Intelligence in Learning Environments,* In *7th International Conference on Intelligent Tutoring System*.

[5] Conati, C. 2002. Probabilistic assessment of user's emotions in educational games. *Applied Artificial Intelligence*, 16:555-575.

[6] Elliott, C. 1992. The Affective Reasoner: A Process Model of Emotions in a Multi-agent System. Ph.D. Thesis, Northwestern University, Chicago, IL.

[7] Frijda, N. 1986. *The Emotions*. Cambridge University Press.

[8] Gratch, J., and Marsella, S. 2004. A domain-independent framework for modeling emotion. *Journal of Cognitive Systems Research*, 5(4):269-306.

[9] Hanley, J., and McNeil, B. 1982. The meaning and use of the area under the Receiver Operating Characteristic (ROC) curve. *Radiology* 143:29-36.

[10] Johnson, L., and Rizzo, P. 2004. Politeness in tutoring dialogs: "Run the factory, that's what I'd do". In *7th International Conference on Intelligent Tutoring Systems,* 67-76.

[11] Kapoor, A., Picard, R., and Ivanov, Y. 2004. Probabilistic combination of multiple modalities to detect interest. *Proc. of International Conference on Pattern Recognition*.

[12] Lazarus, R. 1991. *Emotion and Adaptation*. Oxford University Press, UK.

[13] Lester, J., Towns, S., and FitzGerald, P. 1999. Achieving affective impact: Visual emotive communication in lifelike pedagogical agents. *The International Journal of Artificial Intelligence in Education*, 10(3-4):278-291.

[14] Ortony, A., Clore, G., and Collins, A. 1988. *The Cognitive Structure of Emotions*. Cambridge University Press, Cambridge, MA.

[15] Paiva, A., Dias, J., Sobral, D., Aylett, R., Woods, S., Hall, L., and Zoll, C. 2005. Learning by feeling: Evoking empathy with synthetic characters. *Applied Artificial Intelligence*, 19:235-266.

[16] Picard, R. 1997. *Affective Computing*. MIT Press, Cambridge, MA.

[17] Picard, R., Vyzas, E., and Healey, J. 2001. Toward machine emotional intelligence: analysis of affective physiological state. *IEEE Transactions Pattern Analysis and Machine Intelligence* 23(10): 1185-1191.

[18] Porayska-Pomsta, K. and Pain, H. 2004. Providing cognitive and affective scaffolding through teaching strategies. In *7th International Conference on Intelligent Tutoring System*, 77-86.

[19] Prendinger, H., Mayer, S., Mori, J., and Ishizuka, M. 2003. Persona effect revisited: Using bio-signals to measure and reflect the impact of character-based interfaces. In *4th International Conference on Intelligent Virtual Agents*. 283-291.

[20] Witten, I., and Frank, E. 2005. *Data Mining: Practical Machine Learning Tools and Techniques*. 2nd Edition, Morgan Kaufman, San Francisco, CA.

[21] de Vicente, A. and Pain, H. 2002. Informing the detection of the students' motivational state: An empirical study. In *6th International Conference on Intelligent Tutoring Systems.*