

# Evaluating the Consequences of Affective Feedback in Intelligent Tutoring Systems

Jennifer Robison  
North Carolina State University  
Raleigh, North Carolina, USA  
jlrobiso@ncsu.edu

Scott McQuiggan  
SAS Institute  
Cary, North Carolina, USA  
Scott.McQuiggan@sas.com

James Lester  
North Carolina State University  
Raleigh, North Carolina, USA  
lester@ncsu.edu

## Abstract

*The link between affect and student learning has been the subject of increasing attention in recent years. Because of the possible impacts of affective state on learning, it is a goal of many intelligent tutoring systems to attempt to control student emotional states through affective interventions. While much work has gone into improving the quality of these interventions, we are only beginning to understand the complexities of the relationships between affect, learning, and feedback. This paper investigates the consequences associated with providing affective feedback. It represents a first step toward the long-term objective of designing intelligent tutoring systems that can utilize this information for analysis of the risks and benefits of affective intervention. It reports on the results of two studies that were conducted with students interacting with affect-informed virtual agents. The studies reveal that emotion-specific risk/reward information is associated with particular affective states and suggests that future systems might leverage this information to make determinations about affective interventions.*

## 1. Introduction

Affect has begun to play an increasingly important role in intelligent tutoring systems. The intelligent tutoring community has seen the emergence of work on affective student modeling [6], detecting frustration and stress [4, 18], modeling agents' emotional states [1, 14] devising affectively informed models of social interaction [16, 21], detecting student motivation [11], and diagnosing and adapting to student self-efficacy [3]. 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 [17, 13,

22, 24]. All of these factors have the potential to impact both how a student learns immediately and their learning behaviors in the future. Consequently, developing techniques for keeping students in an affective state that is conducive to learning has been a focus of recent work [5, 10, 12, 20, 23].

However, while much work has targeted the development of optimal techniques for supporting student affect, the nature of the problem introduces a significant degree of uncertainty. In human-human social interaction it is often difficult to determine how best to respond to an individual's affective states. The problem is significantly more challenging for computational systems: they must first be able to correctly recognize student affective states and then decide how best to respond. Systems may frequently encounter situations in which they are uncertain about how to provide affective support and what the effects of a possible intervention may be.

In these inevitable cases of uncertainty, systems should have a means for analyzing the possible impacts of their actions. It may be the case that the system gives a response which is well received and positively impacts the student. However, since the system is uncertain of the quality of its potential response, it may also be the case that the student will respond negatively to the type of feedback provided by the system. The system should therefore have a means of weighing this uncertainty with the potential risks and benefits of intervention in order to decide whether or not to perform the action.

In this paper we empirically investigate the consequences of affective feedback on students interacting with a virtual agent in a narrative-centered learning environment. We examine the negative consequences of poorly selected feedback strategies and contrast these with the positive effects that stand to be gained from appropriate support strategies. This analysis represents a first step towards providing affective support systems the capability of performing risk/benefit analyses when reasoning about potential affective interventions.

## 2. Background

A broad range of techniques have been developed to provide appropriate affective support. For example, Forbes-Riley and Litman [12] performed an analysis of

human-tutor responses to student affect in human-human interactive sessions. It was found that human tutors used different tutorial strategies based on student perception of uncertainty even when controlling for actual incorrectness of solutions. Alternatively, other work has focused on directly modifying student affect. Chaffar and McLaren [5] propose an Emotional Intelligent Agent, which can detect and attempt to modify student affect based on personality characteristics. This agent determines the most positive affective state for any given student and utilizes guided imagery, music and presented images to bring about this state. D’Mello *et al.* [10] have also proposed methods for responding to student affect using empathetic and tutorial dialogue acts accompanied by visual facial expressions and emotionally synthesized speech. In similar work, McQuiggan *et al.* [19] have investigated the impact of emotionally intelligent social agents who utilize varying tactics when responding to student affect. It was found that the likelihood of a student entering a given affective state after an intervention is altered based on the type of response given by the virtual agent. Additionally, studies of affective trajectories by Baker *et al.* [2] and D’Mello *et al.* [9] have shown that, in general, students are more likely to remain in the same affective state if no intervention is provided. This tendency appears to be particularly strong for students in negative affective states, a finding that motivates the need for developing a clearer understanding of when interventions are warranted and what types of interventions are effective.

Throughout learning interactions, affect support systems continually encounter situations in which they must determine whether to intervene to provide affective support. Each time a system reaches such a juncture, it should be able to weigh the potential benefits gained from providing a response that may be effective against the potential hazards of providing responses that may have negative consequences. However, without a clear understanding of the impact of these two categories of feedback, a system will be unable to make an informed decision.

### 3. The CRYSTAL ISLAND Environment

The studies were conducted in a narrative-centered inquiry-based learning environment, CRYSTAL ISLAND (Figure 1). This environment is being created in the domains of microbiology and genetics for middle school students. It features a science mystery set on a recently discovered volcanic island where a research station has been established to study the unique flora and fauna.

The user plays the protagonist, Alex, who is attempting to discover the source of an unidentified infectious disease at the 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



Figure 1. The user, Alex, with Jin, the camp nurse, on CRYSTAL ISLAND

specific source of the outbreak. She is free to explore the world and interact with other characters while forming questions, generating hypothesis, collecting data, and testing her hypotheses. She can pick up and manipulate objects, and she can talk with characters to gather clues about the source of the disease. In the course of her adventure she must gather enough evidence to correctly identify the type and source of the disease that has infected the camp members.

## 4. Method

To empirically investigate the preference and effects of feedback strategies, two studies were conducted with subjects interacting with pedagogical agents. In the first of the two studies, agents could respond to student affect with either parallel or reactive empathetic feedback (described below). Preference results between these two methods were used to inform the empathetic capabilities of agents in the second study. However, in this case, agents could also respond with affect-directed task based feedback. The mechanisms by which agents exhibited these behaviors are presented below.

### 4.1. Parallel and Reactive Empathy

Empathy is the expression of emotion based on another’s situation and not merely one’s own [8]. Its expression can demonstrate that the target’s (the recipient of empathetic expression) feelings are understood or shared. In the case of *parallelemathy*, an individual exhibits an emotion similar to that of the target [8]. This is typically based on an understanding of the target’s situation and shows the empathizer’s ability to identify with the target. *Reactive empathy*, in contrast, focuses on the target’s affective state, in addition to her situation [8]. Reactive empathizers will display emotions that are different from the target’s, often in order to alter or enhance the target’s own affective state.

In the studies, empathetic responses were realized with short, text-based responses consisting of one to two sentences. Parallel responses consisted of the character expressing the same emotion as the user through text responses. For example, in response to student frustration, the agent may say, “Yes, I’m very frustrated

as well!” On the other hand, reactive empathetic responses are designed to be more motivating and thus reveal the character’s desire for the user to be in a positive emotional state. In this case an agent may respond to frustration by saying, “I can understand why you are frustrated, but if you keep working, I’m sure you will figure it out.”

## 4.2. Task-Based Feedback

As a supplement to empathetic strategies, task-based feedback was incorporated into agent behavior as an alternative method for supporting student affect. CRYSTAL ISLAND agents used the following strategies to deliver task-based feedback in response to student affect. If students report a positive affective state, it is assumed that they are proceeding well through the environment, or are at least confident that they are able to achieve the goal on their own. Therefore, to avoid interrupting this state, agents provided feedback in the form of a summary of their current progress. This strategy served to ground students’ current states within the environment and provide reinforcement of previously learned material [14].

However, if a student reported a negative state, agents intervened to try to assist them in overcoming the source of this problem. Because of the open-ended nature of the CRYSTAL ISLAND environment, sources of negative emotions are difficult to distinguish. A student may be confused over some aspect of pedagogical content or over the next goal they are supposed to achieve in the narrative. Rather than attempt to guess at the source of negative emotion and risk an inappropriate response, the agent directed students towards information that would help them complete the goal. This may have entailed reading a book to acquire additional domain knowledge or talking to a character who could provide further direction. In addition to these hints, agents also provided a short summary of learned information, similar to the response to positive affect, to provide reinforcement and grounding. All responses were kept brief (2-3 sentences) to avoid overloading the student or leading them to avoid reading the material all together [25].

## 4.3. User Studies

The first study varied responses across parallel and reactive empathy. The subjects of the study consisted of 35 graduate students ranging in age from 21 to 60 including 9 females and 26 males. Participants were given thirty-five minutes to solve the mystery.

When subjects decided to interact with the agents, the following schema was used to direct subject-character interactions and virtual character feedback:

1. The agent queries the subject for a self-reported affective state by asking the question, “Hi Alex, how are you feeling?” The subject may respond by selecting one of the available emotions (*anger, anxiety, boredom, confusion, delight, excitement, flow, frustration, sadness, fear*).

2. The agent then responds to the subject’s reported affective state with a randomized feedback response. Responses varied between parallel and reactive empathetic statements.

3. A follow-up dialog box is then presented to the subject asking her to respond with the prompt, “... and you respond.” The subject is able to choose from four Likert-scaled responses designed to evaluate the appropriateness and effectiveness of the virtual character’s response. Subjects can issue responses ranging from (1) “That does not help me at all.” to (4) “Thanks, that helped a lot!”

4. The agent responds with a one-word quip (e.g., “Thanks,” or “Great!”) directed towards the subject’s evaluation response (Step 3).

5. At the conclusion of the interaction, the agent again asks the subject how she feels. The subject is presented a dialog box similar to the one described in Step 1 without the character greeting. Here, the character prompts the subject with, “How are you feeling now?”

The second study built on the first study by supplementing empathetic feedback with task-based feedback. The second study also utilized empathetic strategies informed by the first study’s results. The subjects of the second study consisted of 41 college students ranging in age from 19 to 30 and included 12 females and 29 males.

The methodology of the second study was nearly identical to that of the first with the following exceptions. First, agent behavior was randomized between empathetic and task-based feedback. Additionally, the possible set of emotions that students could report was slightly altered. Due to extremely infrequent reports of *sadness* and *fear* by subjects in the first study, they were removed from the set of possible emotions in the second study, and *curiosity* was included as an alternative positive emotion to *flow*.

## 5. Results

After discussing students’ ratings of the quality of agent responses, we will then turn to the overall impact of affective intervention.

### 5.1. Quality of Responses

Because the focus of this investigation is on the impact of low and high quality responses, we will first examine the frequency of these ratings by emotion (Figure 2). In both studies, agents’ responses to negative emotions were rated significantly worse than agents’ responses to positive emotions. While it may be the case that students in a negative state are less likely to give positive ratings to attempts at intervention, it may also be likely that we simply have not yet captured a method for responding appropriately to these states. This appears to be a reasonable explanation given that,

with the exception of *anger*, all negative affective states are being met with responses that are deemed to be “poor” on average.

Additionally, these particularly poor ratings apply only to negative learning emotions (*frustration, anxiety, boredom, confusion*) and not the negative “narrative” emotions (*anger, sadness, fear*), suggesting that negativity alone is not causing the discrepancy in ratings. The worst ratings are those in response to *confusion* and *frustration*. The average quality rating for these emotions is significantly worse than ratings for positive emotions such as flow and excitement ( $p < 0.05$ ). This suggests that since agents are more likely to give inappropriate responses to these states, these emotions should be the focus of analysis on the impacts of inappropriate feedback. However, it is also true that poor responses to positive affective states may have particularly severe consequences since dropping from a positive state to any of the negative states is a highly undesirable event.

## 5.2. Effects of Affective Intervention

In order to examine the impact of affective intervention, the data were divided into groups based on the student-rated quality of responses. Because the difference between responses rated 2 or 3 is not very great, analysis focused on ratings of responses for 1 – “very bad” and 4 – “very good”. The effects are quantified by examining the likelihood of affective transitions across interventions of both types. To compute transition likelihoods we adopt D’Mello *et al.*’s  $L$  [9], which is based on Cohen’s Kappa [7], and has been used by Baker *et al.*[2] and McQuiggan *et al.* [19] for affective transition analysis.  $L$  computes the probability that a transition between two affective states (CURRENT  $\rightarrow$  NEXT) will occur, where CURRENT refers to a reported emotion at time  $t$ , while NEXT refers to the next reported emotion at time  $t+1$ .  $L$ ’s values range between  $-\infty$  and 1 where a result of  $L$  equal to 1 translates to emotion NEXT always following the CURRENT emotion; an  $L$  value equal to 0 means the likelihood of transitioning to emotion NEXT is equal to chance, i.e., the probability of experiencing NEXT (the base rate) regardless of the CURRENT emotion. An  $L$  value less than 0 translates to the likelihood of

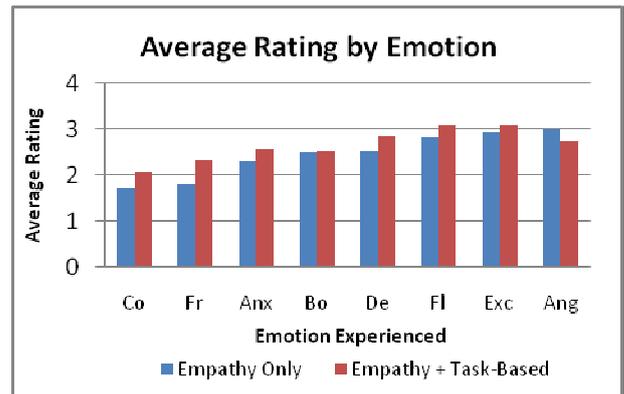


Figure 2 - Quality of response by emotion and study

transitioning to emotion NEXT being less than chance (the probability of experiencing NEXT regardless of the CURRENT emotion).

### 5.2.1 Study 1 Results

The results from the study varying parallel and reactive empathy show interesting differences in student response to feedback of varying quality. The largest of these effects are found in students exhibiting positive affective states. For example, when students experiencing *flow* (Figure 3) are presented with an appropriate form of intervention, they are very likely ( $L=0.40$ ) to remain in this optimal state for learning. However, when presented with an inappropriate response they are most likely to transition into a state of confusion ( $L=0.35$ ). The positive affective state of *delight* shows similar results. Inappropriate feedback leads very strongly to a state of confusion ( $L=1.0$ ) while appropriate feedback maintains the positive state ( $L=0.08$ ).

The results for negative affective states are less clear. Responses to *boredom*, for example, follow similar trends to those of positive states (Figure 4). When presented with poorly selected feedback students are extremely likely to transition into frustration ( $L=0.68$ ) or remain bored ( $L=0.16$ ). Positive affective intervention has very beneficial results in this case by redirecting student attention to the learning task and helping learners to experience the positive state of *flow* ( $L=1$ ).

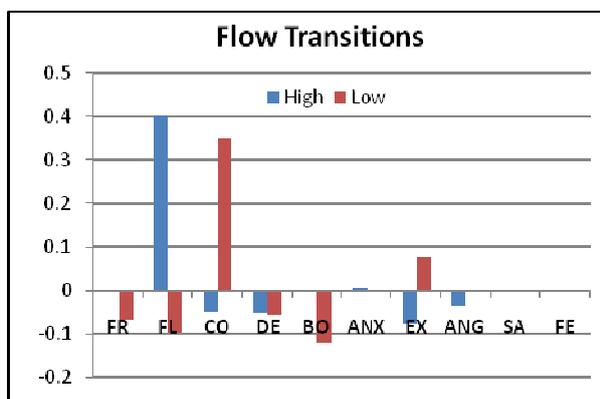


Figure 3 - Transitions from Boredom in Study 1

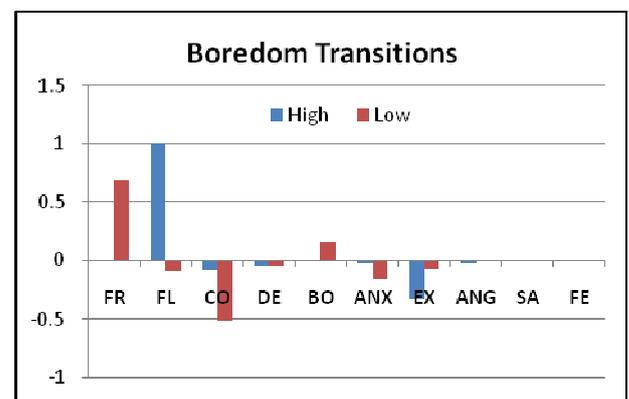


Figure 4 - Transitions from Flow in Study 1

While the impacts of intervention are very clear in the negative state of boredom, the case is less clear in emotions such as *frustration* and *confusion*. For these states the differences in transitions across intervention quality is less notable. Regardless of the rating of intervention quality, these students tend to remain in their negative affective state. Consequently, the risks of causing additional harm to student affective state in these cases is extremely low, suggesting that any attempt to bring about a positive change in affect would be worthwhile. The opposite is true for positive affective states and some negative states such as *boredom* and *anxiety*. In these states, students experience significant negative reactions to inappropriate responses. Therefore, when a student is in one of these affective states, to ensure that unintentional harmful consequences are avoided, caution should be used before attempting an intervention.

### 5.2.2 Study 2 Results

Results from the second user study including task-based feedback mirrored many of the results from the study on parallel and reactive empathy. While the overall effects were typically the same, the level of distinction between effects of high and low quality responses was diminished. It is interesting to note that the similarity in the two sets of results suggest that the interactions of affective state and intervention seem to persist across intervention type.

While there were few differences between the two sets of results, one exception stands out. In the case of a student experiencing *confusion*, it was found that students receiving helpful responses tended to remain in a state of *confusion* ( $L = 0.20$ ) whereas they were otherwise most likely to transition to *frustration* ( $L = 0.09$ ). In general, *confusion* is considered more positive than *frustration* as it may be that the student is still willing to work towards understanding the source of his or her confusion, whereas in frustration the student is likely to become disengaged. Therefore, this likelihood of transitioning into *frustration* from *confusion* given an inappropriate feedback response indicates a need to provide increased attention to response quality for *confusion*, which was not revealed by the first study.

### 5.3. Limitations

While the results of these studies are encouraging, it is important to note an important limitation. The studies rely heavily on user feedback, i.e., users are asked to report both the emotion they are feeling at any given time and their perception of agent behavior. Self-reports such as these are problematic, in that users may not be able to accurately distinguish their emotional state or may choose not to report it. Consequently, replicating this analysis with expert-judged or physiologically-based assessments of emotion could yield a more accurate view of students' true emotions. Additionally, giving users the opportunity to explain their ratings could provide qualitative evidence about the source of

their judgments, and provide insight about the cause of the affective transitions that occur.

## 6. Conclusion

Designing affective support systems to facilitate learning poses a significant challenge: we can never be entirely certain that they are delivering useful affective feedback. There will always be some risk of unintentional negative consequences when attempting to intervene to modify student affect. Understanding the consequences of both appropriate and inappropriate feedback attempts could enable affect-support systems to make informed decisions as to the potential hazards or benefits which stand to be gained by attempting interventions.

It was found that emotions such as *flow*, *delight* and *boredom* appear to be particularly susceptible to the quality of feedback given. In these cases, caution should be exercised in the face of uncertainty to avoid negative consequences. However, for particularly negative emotions such as *frustration*, the risks of inappropriate delivery are not large enough to warrant extreme caution when providing responses. Instead, it appears that any attempt to alleviate the frustration can only help the student, and consequently, strategies should be used regardless of certainty about their outcomes. In designing future affect-support systems, it is hoped that incorporating results such as these into decision-making about whether or not to provide affective feedback will minimize the occurrence of unintentional negative effects on the student while simultaneously maximizing the likelihood of positive outcomes.

Though preliminary, these results point to several important directions for future work. While the findings give us a clearer understanding of the impact of well-chosen and inappropriate feedback on student affect, it is not sufficient for an analysis of the potential risks and benefits of specific actions. In order to develop a fully functional mechanism for making such determinations, a utility must be assigned to the affective states that the student is likely to experience after an intervention. This is likely not a straightforward task. For example, the state of *confusion* is sometimes considered positive for learning because students will be motivated to overcome the source of their misunderstanding. However, it is also likely that if a student remains confused for an extended period of time, they may disengage and no longer be interested in learning. Consequently, a significant research effort is needed to determine the utility of these affective states within the context of the learning environment.

Once this information is obtained, it can be incorporated along with the results presented here into a model for analyzing the risk and benefits of affective intervention. These findings could be used to determine threshold values of uncertainty at which the expected utility of performing an action is more than that of doing nothing at all. Affect-support systems would then act

only in these instances to minimize the risk of inducing harmful effects on the student's emotional state.

## Acknowledgements

The authors would like to thank the other members of The IntelliMedia Group at North Carolina State University for useful discussions and support. We are grateful to 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, IIS-0757535, DRL-0822200 and IIS-0812291. This material is based upon work supported under a National Science Foundation Graduate Research Fellowship. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

## References

- [1] E. André and M. Mueller. Learning affective behavior. In *Proceedings of the 10th Intl. Conf. on Human-Computer Interaction*. Lawrence Erlbaum, Mahwah, NJ, 512-516, 2003.
- [2] R. Baker, M. Rodrigo, and U. Xolocotzin. The dynamics of affective transitions in simulation problem-solving environments. *Proceedings the 2nd International Conference on Affective Computing and Intelligent Interactions*, Lisbon, Portugal, 666-677, 2007.
- [3] C. Beal and H. Lee. 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, in conjunction with the 12th Intl. Conf. on Artificial Intelligence in Education*, 2005.
- [4] W. Burleson. 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.
- [5] S. Chaffar and C. Frasson. Using an emotional intelligent agent to improve the learner's performance. *Proceedings of the Workshop on Social and Emotional Intelligence in Learning Environments in conjunction with Intelligent Tutoring Systems*, Maceio, Brazil, 2004.
- [6] C. Conati and H. McLaren. Data-driven refinement of a probabilistic model of user affect. *Proceedings of the 10th Intl. Conf. on User Modeling*, Springer-Verlag, New York, NY, 40-49, 2005.
- [7] J. Cohen. A coefficient of agreement for nominal scales. *Educational and Psychological Measurement*, 20:37-46, 1960.
- [8] M. Davis. *Empathy: A Social Psychological Approach*, 1994. Brown and Benchmark Publishers, Madison, WI.
- [9] S. D'Mello, R.S. Taylor, and A. Graesser. Monitoring Affective Trajectories during Complex Learning. *Proceedings of the 29th Annual Meeting of the Cognitive Science Society*, Austin, TX, 203-208, 2007.
- [10] S. D'Mello, T. Jackson, S. Craig, B. Morgan, P. Chipman, H. White, N. Person, B. Kort, R. el Kaliouby, R. Picard, and A. Graesser. AutoTutor detects and responds to learners affective and cognitive states. *Proceedings of the Workshop on Emotional and Cognitive issues in ITS in conjunction with the 9th International Conference on Intelligent Tutoring Systems*, 31-43, 2008.
- [11] A. de Vicente and H. Pain. Informing the detection of the students' motivational state: an empirical study. *Proceedings of the 6th International Conference on Intelligent Tutoring Systems*, Springer-Verlag, New York, NY, 933-943, 2002.
- [12] K. Forbes-Riley and D. Litman. Investigating human tutor response to student uncertainty for adaptive system development. *Proceedings the 2nd International Conference on Affective Computing and Intelligent Interactions*, Lisbon, Portugal, 2007.
- [13] D. Goleman. *Emotional Intelligence*, 1995. Bantam Books, New York
- [14] A. Graesser, N. Person, and J. Magliano. Collaborative dialogue patterns in naturalistic one-to-one tutoring. *Journal of Applied Cognitive Psychology* 9:495-522, 1995.
- [15] J. Gratch and S. Marsella. A domain-independent framework for modeling emotion. *Journal of Cognitive Systems Research*, 5(4):269-306, 2004.
- [16] L. Johnson and P. Rizzo. Politeness in tutoring dialogs: "Run the factory, that's what I'd do". *Proceedings of the 7th Intl. Conf. on Intelligent Tutoring Systems*. Springer-Verlag, New York, NY, 2004.
- [17] B. Kort, R. Reilly, and R. Picard. An affective model of interplay between emotions and learning: Reengineering educational pedagogy—building a learning companion. *Proceedings IEEE Intl. Conf. on Advanced Learning Technology: Issues, Achievements and Challenges*. Madison, WI: IEEE Computer Society. 43-48, 2001.
- [18] S. McQuiggan, S. Lee, and J. Lester. Early prediction of student frustration. In *Proc. of the 2nd Intl. Conf. on Affective Computing and Intelligent Interaction*, Portugal, 2007.
- [19] S. McQuiggan, J. Robison, and J. Lester. Affective transitions in narrative-centered learning environments. *Proceedings of the 9th International Conference on Intelligent Tutoring Systems*, Montreal, Canada, 2008.
- [20] S. McQuiggan, J. Robison, R. Phillips, and J. Lester. Modeling parallel and reactive empathy in virtual agents: An inductive approach. *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multi-Agent Systems*, Portugal, 2008.
- [21] A. Paiva, J. Dias, D. Sobral, R. Aylett, S. Woods, L. Hall, and C. Zoll. Learning by feeling: Evoking empathy with synthetic characters. *Applied Artificial Intelligence*, 19:235-266, 2005.
- [22] R. Picard, S. Papert, W. Bender, B. Blumberg, C. Breazeal, D. Cavallo, T. Machover, M. Resnick, D. Roy, and C. Strohecker. Affective learning – a manifesto. *BT Technology Journal*, 22(4), 2004.
- [23] J. Robison, S. McQuiggan, and J. Lester. Modeling task-based vs. affect-based feedback behavior in pedagogical agents: an inductive approach. To appear in *Proceedings of the 14th International Conference on Artificial Intelligence in Education (AIED)* Brighton, England, 2009.
- [24] N. Schwarz. Emotion, cognition, and decision making. *Journal of Cognition and Emotion*, 14(4):443-440, 2000.
- [25] V. Shute. *Focus on Formative Feedback*. ETS. Princeton, NJ, 2007.