

A Motivationally Supportive Affect-Sensitive AutoTutor

Sidney K. D’Mello, Blair Lehman, and Art Graesser

Introduction

This chapter describes a fully automated affect-sensitive Intelligent Tutoring System (ITS) called the Affective AutoTutor (also called the Supportive AutoTutor). AutoTutor is an ITS that helps students learn topics in Newtonian physics, computer literacy, and critical thinking via natural language dialogues that simulate the dialogue patterns observed in human–human tutoring (Graesser, Chipman, Haynes, & Olney, 2005; Graesser et al., 2004; Storey, Kopp, Wiemer, Chipman, & Graesser, in press). The AutoTutor system uses state-of-the-art natural language understanding mechanisms to model student cognitive states and plan its dialogue moves in a manner that is sensitive to these states. The Affective AutoTutor takes the level of intelligence and interactivity even further by using emerging technologies from the field of Affective Computing (Calvo & D’Mello, 2010; McNeese, 2003; Paiva, Prada, & Picard, 2007; Picard, 1997) to model and respond to students’ affective states in addition to their cognitive states.

The achievement of an affect-sensitive tutorial interaction engages the tutor and learner in an *affective loop* (Conati, Marsella, & Paiva, 2005). This includes the *identification* of the affective states relevant to learning, the real-time *detection* of those states, the *selection* of appropriate tutor actions that maximize learning while influencing learner affect, and the *synthesis* of emotional expressions by the tutor as it attempts to engage learners in a more human-like manner.

Implementing the affective loop in an integrated ITS must incorporate the perspective of both the learner and the tutor. The learner-centric view consists of analyzing the prominent affective states of the learner, assessing their potential impact

S.K. D’Mello(✉)
Department of Psychology, Institute for Intelligent Systems,
University of Memphis, Memphis, TN, USA
e-mail: sdmello@memphis.edu

on learning, identifying how these states are expressed by the learner, and developing automatic systems to detect these states in real time. The tutor-centric view explores how good human tutors or theoretically ideal tutors adapt their instructional agenda to encompass the emotions of the learner. This expert knowledge can then be transferred to computer tutors. Animated pedagogical agents that simulate human tutors can also be programmed to synthesize affective elements through the generation of facial expressions, inflections of speech, and the modulation of posture (Graesser, Jeon, & Dufty, 2008; Johnson, Rickel, & Lester, 2000; Moreno, Mayer, Spiers, & Lester, 2001).

We have implemented aspects of both the learner-centric and tutor-centric perspectives in the Affective AutoTutor, as will be described in this chapter. We begin with a brief description of the AutoTutor system followed by an analysis on five studies that have systematically tested links between emotions and learning with AutoTutor. We then describe how the Affective AutoTutor detects, responds to, and synthesizes affect followed by the results of an experiment that evaluated the efficacy of the system in promoting learning and engagement.

AutoTutor

AutoTutor is a dialogue-based ITS for Newtonian physics, computer literacy, and critical thinking. The impact of AutoTutor on facilitating the learning of deep conceptual knowledge has been validated in over a dozen experiments on college students (Graesser et al., 2004; Storey et al., in press; VanLehn et al., 2007). Tests of AutoTutor have produced learning gains of 0.4–1.5 sigma (a mean of 0.8), depending on the learning measure, the comparison condition, the subject matter, and the version of AutoTutor. So we take it as a given that (the nonaffective) AutoTutor helps learning and now the pertinent question is whether the new Affective AutoTutor can yield further enhancements of learning.

The major components of AutoTutor include an animated conversational agent, dialogue management, speech act classification, a curriculum script, semantic evaluation of student contributions, and electronic documents (e.g., textbook and glossary). AutoTutor communicates through an animated conversational agent utilizing speech, facial expressions, and some rudimentary gestures.

AutoTutor's dialogues are organized around difficult questions or problems (called main questions) that require reasoning and explanations in the answers. When presented with these questions, students typically respond with answers that are only one word to two sentences in length, which is typically not sufficient to answer these main questions. In order to guide students in the construction of an improved answer, AutoTutor actively monitors learners' knowledge states and engages them in a turn-based dialogue.

As with most ITSs, AutoTutor fits within VanLehn's analyses of the outer loop and the inner loop when characterizing the scaffolding of solutions to problems, answers to questions, or completion of complex tasks (VanLehn, 2006).

The outer loop involves the selection of topics and problems to cover, assessments of the student's topic knowledge and general cognitive abilities, and global aspects of the tutorial interaction. The inner loop consists of covering individual steps within a problem at a micro-level.

The outer loop of AutoTutor consists of a series of didactic lessons and challenging problems or questions (such as *why*, *how*, *what-if*). An example main question is "When you turn on the computer, how is the operating system first activated and loaded into RAM?" The order of lessons, problems, and questions can be dynamically selected based on the profile of student abilities, but the order is fixed in most versions of AutoTutor we have developed.

The interactive dialogue occurs during the construction of a response to the problems and questions. The answer to a question (or solution to a problem) requires several sentences of information in an ideal answer. AutoTutor assists the learner in constructing their answer after the student enters their initial response. The inner loop of AutoTutor consists of this collaborative interaction while answering a question (or solving a problem). It is this inner loop that is the distinctive hallmark of AutoTutor. AutoTutor adaptively manages the tutorial dialogue by providing feedback on the learner's answers (e.g., "good job," "not quite"), pumps the learner for more information (e.g., "What else?"), gives hints (e.g., "What about X?"), prompts for specific words (e.g., "X is a type of what?"), corrects misconceptions, answers questions, and summarizes topics. The inner loop dialogue between AutoTutor and the student takes approximately 100 dialogue turns to answer a single challenging question, approximately the length of a conversation with a human tutor (Graesser, Person, & Magliano, 1995).

AutoTutor can keep the dialogue on track because it is always comparing what the student says to anticipated input from a curriculum script. This constitutes AutoTutor's model of the student's knowledge and cognitive states. Pattern matching operations and pattern completion mechanisms drive the comparison. These matching and completion operations are based on symbolic interpretation algorithms (Rus & Graesser, 2007) and statistical semantic matching algorithms (Graesser, Penumatsa, Ventura, Cai, & Hu, 2007).

In summary, AutoTutor uses natural language processing techniques, recent advances in agent technologies, insights from discourse processing, the dialogue moves and tactics of human tutors, and strategies from constructivist theories of pedagogy (Chi, Roy, & Hausmann, 2008; Jonassen, Peck, & Wilson, 1999; Moshman, 1982) to allow students to chart their own course through the tutorial dialogue and to construct their own answers to difficult questions.

Identifying Affective States

What are the affective states that learners experience during interactions with AutoTutor and other learning environments? Do the "basic emotions" (anger, sadness, fear, disgust, happiness, and surprise) (Ekman, 1992), constitute learners' primary emotional reactions. Or are the "academic emotions" (e.g., anxiety, boredom) more

relevant in learning contexts (see Pekrun, 2011)? We addressed this fundamental question by conducting a number of studies that aimed at identifying the affective states that learners typically experience while interacting with AutoTutor, with the expectation that these findings will generalize to other learning environments (Baker, D'Mello, Rodrigo, & Graesser, 2010).

In the *observational* study, five trained judges observed the affective states (boredom, confusion, frustration, eureka, flow/engagement, vs. neutral) of 34 students who were learning introductory computer literacy with AutoTutor (Craig, Graesser, Sullins, & Gholson, 2004b). In the *emote-aloud* study, seven college students verbalized their affective states while interacting with AutoTutor (D'Mello, Craig, Sullins, & Graesser, 2006). The *multiple-judge* study consisted of 28 learners completing a 32-min session with AutoTutor, after which their affective states were judged by the learners themselves, untrained peers, and two trained judges. Judgments were based on videos of learners' faces and computer screens which were recorded during the tutorial session (Graesser et al., 2006). The *speech recognition* study was similar to the multiple-judge study with the exception that learners spoke their responses to the AutoTutor system instead of typing them. Retrospective self-reports by the learners constituted the primary affect measure in this study (Graesser, Chipman, King, McDaniel, & D'Mello, 2007). The *physiological* study also implemented the retrospective affect judgment procedure, however, the learners were 27 engineering students from an Australian University instead of the undergraduate psychology students from the USA, who comprised the samples in the previous four studies (Pour, Hussein, AlZoubi, D'Mello, & Calvo, 2010).

When averaged across studies, flow/engagement was the most frequent state, comprising 24% of the observations. Boredom and confusion were the second most frequent states (18 and 17%, respectively) followed by frustration (13%). Neutral was reported for 19% of the observations, while delight (6%) and surprise (3%) were rare.

Although the present set of studies did not directly compare the occurrence of these learning-centered affective states with the basic emotions, other studies have demonstrated that the basic emotions are comparatively rare in learning sessions (one exception is happiness which does occur in some contexts) (Lehman, D'Mello, & Person, 2008; Lehman, Matthews, D'Mello, & Person, 2008). The basic emotions have claimed center-stage of most emotion research in the last four decades, but our results suggest that they might not be relevant to learning, at least for the short learning sessions of these studies. In contrast, confusion, frustration, and boredom, were the prevalent negative emotions, indicating that it is critically important for the Affective AutoTutor to respond to these states.

Detecting Affective States

The affect detection system monitors conversational cues, gross body language, and facial features to detect boredom, confusion, frustration, and neutral (no affect) (see Fig. 1). Automated systems that detect these emotions have been integrated into




Fig. 1 Monitoring affective states during interactions with AutoTutor

AutoTutor and have been extensively described and evaluated in previous publications (D’Mello, Craig, Witherspoon, McDaniel, & Graesser, 2008; D’Mello & Graesser, 2009). The system is capable of correctly identifying learner affect with approximately 50% accuracy (base-rate=25%) (see D’Mello & Graesser, 2009 for details).

It is beyond the scope of this chapter to describe the individual components of the system, however, it is useful to get a grasp of the multimodal affect detection system as a whole. The system uses a decision-level fusion algorithm where each channel independently provides its own diagnosis of the student’s affective state. These individual diagnoses are combined with an algorithm that selects a single affective state and a confidence value of the detection. The algorithm relies on a voting rule enhanced with a few simple heuristics.

A spreading activation network (Rumelhart, McClelland, & PDP Research Group, 1986) with *projecting* and *lateral* links is used to model decision-level fusion. A sample network is presented in Fig. 2. This hypothetical network has two sensor nodes, *C1* and *C2*, and three emotion nodes, *E1*, *E2*, and *E3*. Each sensor is connected to each emotion by a projecting link (solid lines). The degree to which a particular sensor activates a particular emotion is based on the accuracy by which the sensor has detected the emotion in past offline evaluations (see weights in Fig. 2). So if one sensor is more accurate at detecting boredom than confusion, it will excite the boredom node more than the confusion node, even if its current estimates on the probability of both emotions are approximately equivalent.

Each emotion is also connected to every other emotion with a *lateral* link (dotted lines). These links are weighted and can be excitatory or inhibitory (see weights in Fig. 2). Related emotions excite each other while unrelated emotions inhibit each other. For example, confusion would excite frustration but boredom would inhibit engagement.

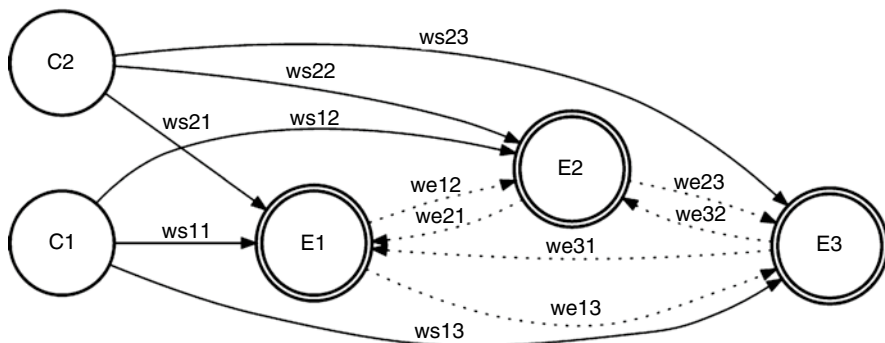


Fig. 2 Sample activation spreading network for decision-level fusion

Each emotion node receives activation from both link types and maintains an *activation value*. At any time, the emotion node with the highest activation value is considered to be the emotion that the learner is currently experiencing. The decision-level fusion algorithm operates in four phases.

1. *Detection by Sensors*. Each sensor provides an independent estimate of the likelihood that the learner is experiencing an emotion. The likelihood can be represented as a probability value for each emotion (e.g., sensor C1 expresses a .53 probability that the current emotion is E1).
2. *Activation from Sensors*. Sensors spread activation and emotion nodes aggregate this activation.
3. *Activation from Emotions*. Each emotion spreads the activation received from the sensors to the other emotions, so that some emotions are excited while others are inhibited.
4. *Decision*. The emotion with the highest activation is selected to be the emotion that the learner is currently experiencing.

Responding to Affective States and Synthesizing Affect

Despite the complexity associated with real-time affect detection, detection is only one piece of the puzzle. The next challenge is to help students regulate their affective states so that positive states such as flow/engagement and curiosity persevere, while negative states such as frustration and boredom are rapidly eradicated.

The Affective AutoTutor addresses this challenge by adapting its dialogue moves in a manner that is dynamically responsive to students' affective and cognitive states. In particular, at any given turn the Affective tutor keeps track of five informational parameters that provide the foundations for affect-sensitivity (three affective parameters and two cognitive parameters). The three affective parameters include

the current affective state detected, the confidence level of that affect classification, and the previous affective state detected. The cognitive parameters include a global measure of student ability (dynamically updated throughout the session) and the conceptual quality of the student's immediate response. These cognitive measures are accessed via the use of natural language understanding techniques that monitor students knowledge trajectories by constantly comparing their responses to information in the curriculum script (Graesser, Penumatsa et al., 2007).

Taking these five parameters as input, the Affective AutoTutor is equipped with a set of production rules to map the input parameters with appropriate tutor actions. In particular, the Affective tutor responds with (a) feedback for the current answer with an affective facial expression, (b) an affective statement accompanied by a matching emotional facial and vocal expression by the tutor, and (c) the next dialogue move. Each of these components is described below.

Feedback with Affective Facial Expression: AutoTutor provides short feedback to each student response. The feedback is based on the semantic match between the response and the anticipated answer. There are five levels of feedback: positive, neutral-positive, neutral, neutral-negative, and negative. Each feedback category has a set of predefined expressions that the tutor randomly selects from. "Good job" and "Well done" are examples of positive feedback, while "That is not right" and "You are on the wrong track" are examples of negative feedback. In addition to articulating the textual content of the feedback, the affective AutoTutor also modulates its facial expressions and speech prosody. Positive feedback is delivered with an *approval* expression (big smile and big nod). Neutral-positive feedback receives a *mild approval* expression (small smile and slight nod). Negative feedback is delivered with a *disapproval* expression (slight frown and head shake), while the tutor makes a *skeptical* face when delivering neutral-negative feedback (see Fig. 3). No facial expression accompanies the delivery of neutral feedback.

Affective Response with Affective Facial Expression and Affectively Modulated Speech: After delivering the feedback, the affective AutoTutor delivers an emotional statement if it senses that the student is bored, confused, or frustrated. A non-emotional discourse marker (e.g., "Moving on," "Try this one") is selected if the




Fig. 3 Synthesized facial expressions by the AutoTutor agent

student is neutral. AutoTutor's strategies to respond to boredom, confusion, and frustration are motivated by attribution theory (Batson, Turk, Shaw, & Klein, 1995; Heider, 1958; Weiner, 1986), cognitive disequilibrium during learning (Craig, Graesser, Sullins, & Gholson, 2004a; Festinger, 1957; Graesser & Olde, 2003; Piaget, 1952), and recommendations by pedagogical experts. In general, the Affective AutoTutor responds to the learners' affective states via empathetic and motivational responses. These responses always attribute the source of the learners' emotion to the material instead of the learners themselves. So the supportive AutoTutor might respond to mild boredom with "This stuff can be kind of dull sometimes, so I'm gonna try and help you get through it. Let's go." A response to confusion would include attributing the source of confusion to the material ("Some of this *material* can be confusing. Just keep going and I am sure you will get it") or the tutor itself ("I know I do not always convey things clearly. I am always happy to repeat myself if you need it. Try this one").

As a complete example, consider a student who has been performing well overall (high global ability), but the most recent contribution was not very good (low current contribution quality). If the current emotion was classified as boredom, with a high probability, and the previous emotion was classified as frustration, then AutoTutor might say the following: "Maybe this *topic* is getting old. I'll help you finish so we can try something new." This is a randomly chosen phrase from a list that was designed to indirectly address the student's boredom and to try to shift the topic a bit before the student becomes disengaged from the learning experience. This rule fires on several different occasions, and each time it is activated AutoTutor will select a dialogue move from a list of associated moves. In this fashion, the rules are context sensitive and are dynamically adaptive to each individual learner.

The affective response is accompanied by an emotional facial expression and emotionally modulated speech. These affective expressions include empathy, mild enthusiasm, high enthusiasm, and neutral in some cases. The facial expressions in each display were informed by Ekman's work on the facial correlates of emotion expression (Ekman & Friesen, 1978). The facial expressions of emotion displayed by AutoTutor are augmented with emotionally expressive speech synthesized by the agent. The emotional expressivity is obtained by variations in pitch, speech rate, and other prosodic features. Previous research has led us to conceptualize AutoTutor's affective speech on the indices of pitch range, pitch level, and speech rate (Johnstone & Scherer, 2000). The current quality of the emotionally modulated speech is acceptable, although there is the potential for improvement.

A screenshot of the Affective AutoTutor is shown in Fig. 4. Here the tutor is displaying a skeptical face while delivering neutral-negative feedback (e.g., "kind of," "sort of").

Next Dialogue Move: Finally, AutoTutor responds with a move to advance the dialogue. In the current version of the tutor, this dialogue move is sensitive to the learner's cognitive state but not to his or her affective state (see "AutoTutor" section). That is, affect-sensitivity is currently applied to AutoTutor's motivational (feedback and affective response) but not its pedagogical moves (i.e., hints, prompts, assertions).

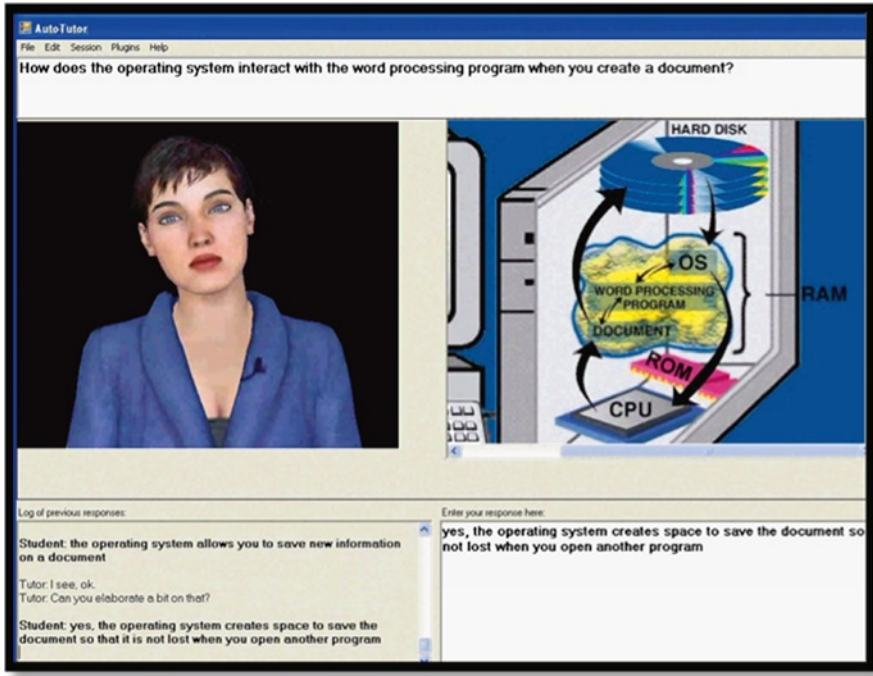


Fig. 4 Screenshot of the Affective AutoTutor

Future affect-sensitive interventions will focus on the tutor’s pedagogical moves as well. This adaptation would increase the bandwidth of communication and allow the Affective AutoTutor to respond at a more sophisticated metacognitive level.

There could be many possible responses to the different affective states of the learner and the context of the interaction. If the affective state of frustration is detected, then the Affective AutoTutor could respond by changing its dialogue strategies to include more direct feedback, assertions, and corrections of detected misconceptions. If the learner is bored then the tutor could engage the learner in a task that increases interest and cognitive arousal, such as a simulation, options of choice, a challenge, or a seductive embedded game.

Confusion presents a key opportunity for the tutor to encourage deep learning. The Affective AutoTutor system could manage confusion in at least two ways. Successful learners might be allowed to work out their own confusion in a discovery learning environment (Bruner, 1961; D’Mello et al., 2010; Vavik, 1993) that requires self-regulated cognitive activities (see Azevedo & Chauncey, 2011). A second method would systematically scaffold the student out of the confused state. This method might work better for learners with lower domain knowledge and lower ability to self-regulate their learning activities.

Evaluating the Affect-Sensitive AutoTutor

We have recently conducted an experiment that evaluated the pedagogical effectiveness of the Affective AutoTutor when compared to the original tutor (D’Mello et al., 2010). Both tutors utilize identical pedagogical strategies, however, the Affective AutoTutor has enhanced motivational moves. The obvious prediction is that learning gains should be superior for the Affective AutoTutor.

The experiment utilized a between-subjects design where 84 learners (a) completed a pretest on topics in computer literacy, (b) were tutored on two computer literacy topics with either the affective or the regular AutoTutor, and (c) completed a posttest. The tests and tutorial sessions were pitched at deeper levels of comprehension with questions that required reasoning and inference instead of the recall of shallow facts and definitions. The tutorial session consisted of two 30-min sessions on different computer literacy topics but with the same version of AutoTutor (i.e., either Affective or Regular). The key dependent variable was *proportional learning gains*, computed as: $(\text{posttest scores} - \text{pretest scores}) / (1 - \text{pretest scores})$. Proportional learning gains represent the degree of improvement at posttest above and beyond pretest performance.

The results of this experiment indicated that the Affective AutoTutor was more effective than the regular tutor for low-domain knowledge students (identified via a median split on pretest scores) in the second session ($d=0.713$), but not the first session (see Fig. 5). This suggests that it is inappropriate for the tutor to be supportive to low-domain knowledge students before there has been enough context to show there are problems. Simply put, do not be supportive until the students need support.

The low-domain knowledge learners also demonstrated more knowledge transfer by scoring higher on related topics that were not covered in the tutorial session. Transfer scores were higher with the Affective AutoTutor when compared to the regular tutor, thereby signaling a unique advantage for this type of motivational support.

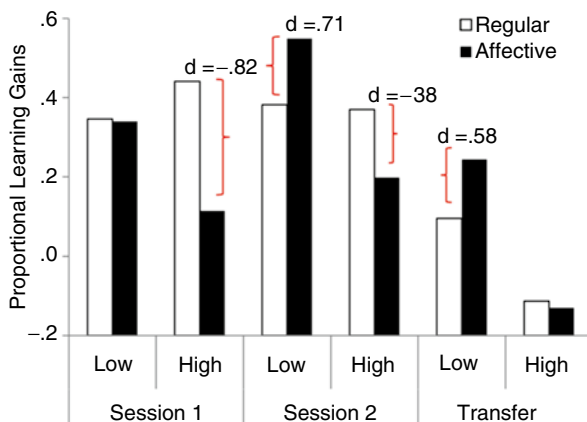


Fig. 5 Proportional learning gains separated by learning session and tutor

The students with more knowledge never benefited from the Affective AutoTutor. These students do not need the emotional support, but rather they need to go directly to the content. There are also conditions when affective support was detrimental to these high domain knowledge students. There appears to be a liability to quick support and empathy compared to no affect-sensitivity.

The central message is that there is an appropriate time for affect-sensitivity in the form of supportive dialogues. Just as there is a “time for telling”; there is a “time for emoting.” We could imagine a strategy where low-knowledge students start out with a nonemotional regular tutor until they see there are problems. Then after that they need support, as manifested in the second tutorial session. Regarding high-knowledge students, they are perfectly fine working on content for an hour or more and may get irritated with an AutoTutor showing compassion, empathy, and care. But later on there may be a time when they want an Affective AutoTutor. These are all questions to explore in future research.

Conclusions

Once no more than a mere seductive vision, the idea of having a tutoring system detect, respond to, and synthesize emotions is now a reality (Picard, 1997). The fact that the Affective AutoTutor is the first affect-sensitive ITS to facilitate deep learning gains of difficult technical materials over and above nonaffective controls signals an important advancement in the field of ITSs and the more general areas of Affective Computing and Human–Computer Interaction. However, there is still much room for further research and technological development. In addition to providing affective motivational support, there is the key challenge of providing affect-sensitive pedagogical support as well. Another challenge is to provide more contextually sensitive affective responses that take into account both the causes and effects of learner emotions (see Conati, 2011). It is also important to consider the “afterglow of affect-sensitivity” which would involve monitoring the learner after an affective-sensitive intervention. One might even consider inducing certain emotions that are considered to be beneficial to learning.

Individual differences play their usual important role so more research is needed into their effects. In addition to prior knowledge, individual differences in motivation, attribution styles, academic risk taking, cognitive complexity, affective traits, and baseline mood states are likely to interact with student affect (Clifford, 1988; Clore & Huntsinger, 2007; Fletcher, Danilovics, Fernandez, Peterson, & Reeder, 1986; Hidi, 2006; Isen, 2008; McAuley, Duncan, & Russell, 1992; Meyer & Turner, 2006). Identifying how these and other individual differences moderate the experience and impact of student emotions and developing affective interventions that capitalizes on these relationships represents a significant challenge for next-generation affect-sensitive ITSs.

Acknowledgments We thank our research colleagues in the Emotive Computing Group and the Tutoring Research Group (TRG) at the University of Memphis (<http://emotion.autotutor.org>). We gratefully acknowledge our partners in the Affective Computing group at the MIT Media Lab.

This research was supported by the National Science Foundation (REC 0106965, ITR 0325428, HCC 0834847). Any opinions, findings and conclusions, or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of NSF.

References

- Azevedo, R., & Chauncey, A. (2011). Integrating cognitive, metacognitive, and affective regulatory processes with MetaTutor. In R. Calvo & S. D'Mello (Eds.), *Affective prospecting (Explorations in the learning sciences, instructional systems and performance)*. New York: Springer.
- Baker, R., D'Mello, S., Rodrigo, M., & Graesser, A. (2010). Better to be frustrated than bored: The incidence and persistence of affect during interactions with three different computer-based learning environments. *International Journal of Human-Computer Studies*, 68(4), 223–241.
- Batson, C., Turk, C., Shaw, L., & Klein, T. (1995). Information function of empathic emotion – learning that we value the others welfare. *Journal of Personality and Social Psychology*, 68(2), 300–313.
- Bruner, J. (1961). The act of discovery. *Harvard Educational Review*, 31(1), 21–32.
- Calvo, R. A., & D'Mello, S. K. (2010). Affect detection: An interdisciplinary review of models, methods, and their applications. *IEEE Transactions on Affective Computing*, 1(1), 18–37.
- Chi, M., Roy, M., & Hausmann, R. (2008). Observing tutorial dialogues collaboratively: Insights about human tutoring effectiveness from vicarious learning. *Cognitive Science*, 32(2), 301–341.
- Clifford, M. (1988). Failure tolerance and academic risk-taking in ten- to twelve-year-old students. *British Journal of Educational Psychology*, 58, 15–27.
- Clore, G. L., & Huntsinger, J. R. (2007). How emotions inform judgment and regulate thought. *Trends in Cognitive Sciences*, 11(9), 393–399.
- Conati, C. (2011). Combining cognitive appraisal and sensors for affect detection in a framework for modeling user affect. In R. Calvo & S. D'Mello (Eds.), *Affective prospecting (Explorations in the learning sciences, instructional systems and performance)*. New York: Springer.
- Conati, C., Marsella, S., & Paiva, A. (2005). Affective interactions: The computer in the affective loop. In J. Riedl & A. Jameson (Eds.), *Proceedings of the 10th International Conference on Intelligent User Interfaces* (p. 7). New York: ACM.
- Craig, S., Graesser, A., Sullins, J., & Gholson, B. (2004). Affect and learning: An exploratory look into the role of affect in learning. *Journal of Educational Media*, 29, 241–250.
- D'Mello, S., Craig, S., Sullins, J., & Graesser, A. (2006). Predicting affective states expressed through an emote-aloud procedure from AutoTutor's mixed-initiative dialogue. *International Journal of Artificial Intelligence in Education*, 16(1), 3–28.
- D'Mello, S., Craig, S., Witherspoon, A., McDaniel, B., & Graesser, A. (2008). Automatic detection of learner's affect from conversational cues. *User Modeling and User-Adapted Interaction*, 18(1–2), 45–80.
- D'Mello, S., & Graesser, A. (2009). Automatic detection of learners' affect from gross body language. *Applied Artificial Intelligence*, 23(2), 123–150.
- D'Mello, S., Lehman, B., Sullins, J., Daigle, R., Combs, R., Vogt, K., et al. (2010). A time for emoting: When affect-sensitivity is and isn't effective at promoting deep learning. In J. Kay & V. Alevan (Eds.), *Proceedings of 10th International Conference on Intelligent Tutoring Systems* (pp. 245–254). Berlin: Springer.
- Ekman, P. (1992). An argument for basic emotions. *Cognition & Emotion*, 6(3–4), 169–200.

- Ekman, P., & Friesen, W. (1978). *The facial action coding system: A technique for the measurement of facial movement*. Palo Alto: Consulting Psychologists Press.
- Festinger, L. (1957). *A theory of cognitive dissonance*. Stanford, CA: Stanford University Press.
- Fletcher, G. J. O., Danilovics, P., Fernandez, G., Peterson, D., & Reeder, G. D. (1986). Attributional complexity – An individual-differences measure. *Journal of Personality and Social Psychology*, 51(4), 875–884.
- Graesser, A., Chipman, P., Haynes, B., & Olney, A. (2005). AutoTutor: An intelligent tutoring system with mixed-initiative dialogue. *IEEE Transactions on Education*, 48(4), 612–618.
- Graesser, A., Chipman, P., King, B., McDaniel, B., & D’Mello, S. (2007). Emotions and learning with AutoTutor. In R. Luckin, K. Koedinger, & J. Greer (Eds.), *13th International Conference on Artificial Intelligence in Education* (pp. 569–571). Amsterdam: Ios Press.
- Graesser, A., Jeon, M., & Dufty, D. (2008). Agent technologies designed to facilitate interactive knowledge construction. *Discourse Processes*, 45(4–5), 298–322.
- Graesser, A., Lu, S. L., Jackson, G., Mitchell, H., Ventura, M., Olney, A., et al. (2004). AutoTutor: A tutor with dialogue in natural language. *Behavioral Research Methods, Instruments, and Computers*, 36, 180–193.
- Graesser, A., McDaniel, B., Chipman, P., Witherspoon, A., D’Mello, S., & Gholson, B. (2006). *Detection of emotions during learning with AutoTutor*. Paper presented at the 28th Annual Conference of the Cognitive Science Society, Vancouver, Canada.
- Graesser, A., & Olde, B. (2003). How does one know whether a person understands a device? The quality of the questions the person asks when the device breaks down. *Journal of Educational Psychology*, 95(3), 524–536.
- Graesser, A., Penumatsa, P., Ventura, M., Cai, Z., & Hu, X. (2007). Using LSA in AutoTutor: Learning through mixed-initiative dialogue in natural language. In T. Landauer, D. McNamara, S. Dennis, & W. Kintsch (Eds.), *Handbook of latent semantic analysis* (pp. 243–262). Mahwah, NJ: Erlbaum.
- Graesser, A., Person, N., & Magliano, J. (1995). Collaborative dialogue patterns in naturalistic one-to-one tutoring. *Applied Cognitive Psychology*, 9(6), 495–522.
- Heider, F. (1958). *The psychology of interpersonal relations*. New York: Wiley.
- Hidi, S. (2006). Interest: A unique motivational variable. *Educational Research Review*, 1, 69–82.
- Isen, A. (2008). Some ways in which positive affect influences decision making and problem solving. In M. Lewis, J. Haviland-Jones, & L. Barrett (Eds.), *Handbook of emotions* (3rd ed., pp. 548–573). New York, NY: Guilford.
- Johnson, W., Rickel, J., & Lester, J. (2000). Animated Pedagogical Agents: Face-to-face interaction in interactive learning environments. *International Journal of Artificial Intelligence in Education*, 11, 47–78.
- Johnstone, T., & Scherer, K. (2000). Vocal communication of emotion. In M. Lewis & J. Haviland-Jones (Eds.), *Handbook of emotions* (2nd ed., pp. 220–235). New York: Guilford Press.
- Jonassen, D., Peck, K., & Wilson, B. (1999). *Learning with technology: A constructivist perspective*. Upper Saddle River, NJ: Prentice Hall.
- Lehman, B., D’Mello, S., & Person, N. (2008). *All alone with your emotions: An analysis of student emotions during effortful problem solving activities*. Paper presented at the Workshop on Emotional and Cognitive issues in ITS at the Ninth International Conference on Intelligent Tutoring Systems
- Lehman, B., Matthews, M., D’Mello, S., & Person, N. (2008). What are you feeling? Investigating student affective states during expert human tutoring sessions. In B. Woolf, E. Aimeur, R. Nkambou, & S. Lajoie (Eds.), *Proceedings of the 9th International Conference on Intelligent Tutoring Systems* (pp. 50–59). Berlin: Springer.
- McAuley, E., Duncan, T. E., & Russell, D. W. (1992). Measuring causal attributions – the Revised Causal Dimension Scale (Cdsii). *Personality and Social Psychology Bulletin*, 18(5), 566–573.
- McNeese, M. (2003). New visions of human-computer interaction: making affect compute. *International Journal of Human-Computer Studies*, 59(1–2), 33–53.

- Meyer, D., & Turner, J. (2006). Re-conceptualizing emotion and motivation to learn in classroom contexts. *Educational Psychology Review*, 18(4), 377–390.
- Moreno, R., Mayer, R., Spiers, H., & Lester, J. (2001). The case for social agency in computer-based teaching: Do students learn more deeply when they interact with animated pedagogical agents? *Cognition and Instruction*, 19(2), 177–213.
- Moshman, D. (1982). Exogenous, endogenous, and dialectical constructivism. *Developmental Review*, 2(4), 371–384.
- Paiva, A., Prada, R., & Picard, R. (Eds.). (2007). *Affective computing and intelligent interaction*. Heidelberg: Springer.
- Pekrun, R. (2011). Emotions as drivers of learning and cognitive development. In R. Calvo & S. D'Mello (Eds.), *Affective prospecting (Explorations in the learning sciences, instructional systems and performance)*. New York: Springer.
- Piaget, J. (1952). *The origins of intelligence*. New York: International University Press.
- Picard, R. (1997). *Affective computing*. Cambridge: MIT Press.
- Pour, P. A., Hussein, S., AlZoubi, O., D'Mello, S. K., & Calvo, R. (2010). The impact of system feedback on learners' affective and physiological states. In J. Kay & V. Alevan (Eds.), *Proceedings of 10th International Conference on Intelligent Tutoring Systems* (pp. 264–273). Berlin, Heidelberg: Springer.
- Rumelhart, D., McClelland, J., & PDP Research Group. (1986). *Parallel distributed processing: explorations in the microstructure of cognition. Volume 1: Foundations*. Cambridge, MA: MIT Press.
- Rus, V., & Graesser, A. (2007). Lexico-syntactic subsumption for textual entailment. In N. Nicolov, K. Bontcheva, G. Angelova, & R. Mitkov (Eds.), *Recent advances in natural language processing IV: Selected papers from RANLP 2005* (pp. 187–196). Amsterdam: John Benjamins.
- Storey, J., Kopp, K., Wiemer, K., Chipman, P., & Graesser, A. (in press). Critical thinking tutor: Using AutoTutor to teach scientific critical thinking skills. *Behavioral Research Methods*.
- VanLehn, K. (2006). The behavior of tutoring systems. *International Journal of Artificial Intelligence in Education*, 16(3), 227–265.
- VanLehn, K., Graesser, A., Jackson, G., Jordan, P., Olney, A., & Rose, C. P. (2007). When are tutorial dialogues more effective than reading? *Cognitive Science*, 31(1), 3–62.
- Vavik, L. (1993). Facilitating discovery learning in computer-based simulation learning environments. In R. Tennyson & A. Baron (Eds.), *Automating instructional design: Computer-based development and delivery tools* (pp. 403–449). Berlin: Springer.
- Weiner, B. (1986). *An attributional theory of motivation and emotion*. New York: Springer.