Modeling Cognitive-Affective Dynamics with Hidden Markov Models

Sidney K. D'Mello (sdmello@memphis.edu)
Institute for Intelligent Systems, University of Memphis,
Memphis, TN 39152 USA

Art Graesser (a-graesser@memphis.edu)
Department of Psychology, University of Memphis
Madison, TN 38152 USA

Abstract

We present and test a theory of cognitive disequilibrium to explain the dynamics of the cognitive-affective states that emerge during deep learning activities. The theory postulates an important role for cognitive disequilibrium, a state that occurs when learners face obstacles to goals, contradictions, incongruities, anomalies, uncertainty, and salient contrasts. The major hypotheses of the theory were supported in two studies in which participants completed a tutoring session with a computer tutor after which they provide judgments on their cognitive-affective states via a retrospective judgment protocol. Hidden Markov Models constructed from time series of learners' cognitive-affective states confirmed the major predictions as well as suggested refinements for the theory of cognitive disequilibrium during deep learning.

Keywords: affect dynamics, hidden markov model, learning.

Introduction

Deep learning and problem solving are emotionally rich experiences. Students experience boredom when the material does not appeal to them, confusion when they have difficulty comprehending the material and are unsure about how to proceed, frustration when they make mistakes and get stuck, and perhaps even despair and anxiety when their efforts seem to be futile and the big exam is creeping around the corner. This negative picture of the emotional experiences that accompany learning has a complimentary positive side. Students experience curiosity when they encounter topics that interest them, eureka moments when insights are unveiled and major discoveries made, delight when challenges are conquered, and perhaps even flow-like states (Csikszentmihalyi, 1990) when they are so engaged in learning that time and fatigue disappear.

There have been several theories that link cognition and affect very generally (Bower, 1981; Mandler, 1984; Ortony, Clore, & Collins, 1988; Russell, 2003; Stein & Levine, 1991). While these theories convey general links between cognition and emotions, they do not directly explain and predict the sort of emotions that occur during complex learning, such as attempts to master physics, biology, or computer literacy. Researchers in many different fields are familiar with Ekman's work on the detection of emotions from facial expressions (Ekman, 1984). However, the emotions that Ekman intensely investigated (e.g., sadness, happiness, anger, fear, disgust, surprise) have minimal relevance to learning in typical academic settings (D'Mello,

Craig, Sullins, & Graesser, 2006; Kort, Reilly, & Picard, 2001; Lehman, D'Mello, & Person, 2008). Instead, the pervasive cognitive-affective states during complex learning include confusion, frustration, boredom, flow/engagement, and sometimes delight, surprise, anxiety, and curiosity (D'Mello et al., 2006; Lehman, Matthews, D'Mello, & Person, 2008).

The identification of the cognitive-affective states that occur during learning is critical, but it could be argued that merely knowing *what* states occur has limited utility. What is missing is a specification of *how* these states evolve, morph, interact, and influence learning and engagement. What is required is a fine-grained analysis of the rapid dynamics of the cognitive-affective processes that naturally occur during effortful learning activities.

Although affect dynamics has been generally ignored by theories that link affect and cognition during learning, one theory, called the cognitive disequilibrium theory, does address transitions between states. The theory postulates an important role for *cognitive disequilibrium* in comprehension and learning processes, a notion that has a long history in psychology (Berlyne, 1960; Festinger, 1957; Piaget, 1952). Cognitive disequilibrium is a state that occurs when learners face obstacles to goals, contradictions, incongruities, anomalies, uncertainty, and salient contrasts (Graesser, Lu, Olde, Cooper-Pye, & Whitten, 2005; Otero & Graesser, 2001; Piaget, 1952).

The cognitive disequilibrium theory is depicted in Figure 1 as a state transition network. The nodes (circles) in the figure represent the cognitive-affective states (in parentheses) and their presumed causes (in bold). Links represent situations that trigger transitions between the different states.
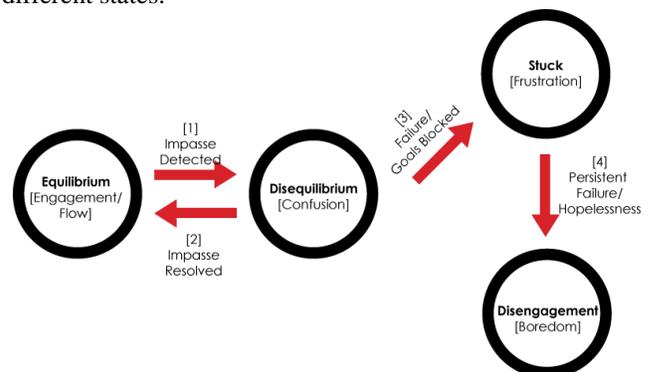

Figure 1. Cognitive Disequilibrium Theory

The theory assumes that learners are in a base state of engagement (perhaps a degree of flow) until they are confronted with a contradiction, anomaly, system breakdown, or error, and when they are uncertain about what to do next (Forbes-Riley & Litman, 2009; Graesser et al., 2005; Siegler & Jenkins, 1989; VanLehn, Siler, Murray, Yamauchi, & Baggett, 2003). Confusion is a key signature of the cognitive disequilibrium that occurs when an impasse is detected (Link 1). Learners must engage in effortful problem solving activities in order to resolve the impasse and restore equilibrium. Equilibrium is restored when the source of the discrepant information is discovered and the impasse is resolved, thereby causing learners to revert back to the engaged state (Link 2).

However, this form of productive confusion associated with impasse resolution can be contrasted with hopeless confusion. This occurs when the impasse cannot be resolved, the student gets stuck, and important goals are blocked. The theory hypothesizes that learners will experience frustration in these situations (Link 3). Furthermore, persistent frustration may transition into boredom, a crucial point at which the learner disengages from the learning process (Link 4).

We have confirmed some of the predictions of the theory in previous publications (D'Mello & Graesser, in review; D'Mello, Taylor, & Graesser, 2007). In particular, we have assessed the presence of oscillations between flow and confusion as well as transitions from confusion to frustration and frustration to boredom. However, verifying the presence of these transitions represents only one important component of the theory. The other crucial component that has not been yet empirically supported pertains to the internal causes that give rise to the observed cognitive-affective patterns. These include an *equilibrium* state that presumably activates the flow/engaged experience, a *disequilibrium* state that causes confusion, a *stuck* state that causes frustration, and a *disengaged* state that emits boredom. Our previous analyses so far have exclusively focused on transitions between the cognitive-affective states but have not explicitly addressed their causes. It is important, however, that both components of the theory be verified before it can be accepted as a useful explanation of the cognitive-affective phenomena that underlies deep learning.

Unfortunately, it is difficult to model the causes that underlie the cognitive-affective expressions. These states can be *observed* via facial expressions, body movements, and contextual cues, but the internal causes are *hidden* (i.e. they cannot be directly observed). This limitation can be alleviated via modeling techniques that permit the simultaneous modeling of both hidden and observed variables. In particular, the present paper describes a study in which Hidden Markov Models (HMMs) were used to model both the observed cognitive-affective states (confusion, frustrations, etc) and their hidden causes (equilibrium, stuck, etc), thereby testing the two components of cognitive disequilibrium theory. The HMMs were parameterized from learners' self reports on their cognitive-affective states via a retrospective judgment protocol after a tutorial session with AutoTutor, an Intelligent Tutoring System with conversational dialogues (Graesser et al., 2004).

Brief Description of HMMs

Hidden Markov Models are valuable tools for modeling system with sequential observable outcomes when the states producing the outcomes cannot be directly observed (i.e. they are hidden). They are widely used to model complex phenomenon with applications in a variety of disparate domains, such as automatic speech recognition, tutorial discourse, computational biology, financial economics, computer vision , and earthquake detection (Jurafsky & Martin, 2008; Rabiner, 1989).

HMMs are characterized by a set of parameters that can be estimated from available data. If there are m hidden states $(H = h_1, h_2, h_3, \ldots h_m)$ and n observable states $(O = o_1, o_2, o_3, \ldots o_n)$, then the parameters include a $m \times n$ emission probability matrix (E) and a $m \times m$ transition probability matrix (T). The emission probability matrix specifies the conditional probability of emitting an observed state o_t at time t given that the system is a hidden state h_t at the same time point $[\Pr(o_t|h_t)]$. On the other hand, the transition probability matrix specifies the conditional probability of transitioning from the current hidden state h_t to the next (or same) hidden state at the next time interval h_{t+1} $[\Pr(h_{t+1}|h_t)]$.

As an example consider a simplified model of two hidden states for equilibrium (E) and disequilibrium (D) and two observed states for flow (F) and confusion (C). Here, $m = n = 2$ and both matrices are of size 2×2. The emission probability matrix would consist of the following four conditional probabilities: $\Pr(F|E)$, $\Pr(C|E)$, $\Pr(F|D)$, and $\Pr(C|D)$. Since it is assumed that a given hidden state emits one of the observable states, $\Pr(F|E) + \Pr(C|E) = 1$ and $\Pr(F|D) + \Pr(C|D) = 1$.

The transition probability matrix would also consist of four probabilities: $\Pr(E|E)$, $\Pr(F|E)$, $\Pr(D|D)$, and $\Pr(E|D)$. Once again, $\Pr(E|E) + \Pr(D|E) = 1$ and $\Pr(D|D) + \Pr(E|D) = 1$. Hence, given that a learner is in one of the hidden states, we can probabilistically determine which cognitive-affective state is most likely to be observed as well as what the next hidden state is likely to be.

Methods

Study 1
Participants. 28 undergraduate students (5 male and 23 female) from a large mid-south university participated for extra credit in their psychology courses.

Interaction with AutoTutor. Participants interacted with AutoTutor for 32 minutes on one of three randomly assigned topics in computer literacy: hardware, Internet, or

operating systems. AutoTutor is a validated intelligent tutoring system that helps learners construct explanations by interacting with them in natural language with adaptive dialogue moves similar to human tutors (Graesser et al., 2004). AutoTutor's dialogues are organized around difficult questions, such as why, how, what-if, what if not, how is X similar to Y, that require answers involving inferences, explanations, and deep reasoning. Although each question requires 3-7 sentence-like ideas in a correct answer, learners rarely give the complete answer in a single conversational turn. Therefore, the tutor scaffolds the construction of an answer by an adaptive dialogue with pumps for information, hints, prompts, assertions, summaries, and feedback. AutoTutor delivers its dialogue moves via an animated conversational agent that speaks the content of the tutor's turns.

A video of the participant's face and computer screen was recorded during the tutorial session (see Figure 2). Gross body language was tracked using Tekscan's Body Pressure Measurement System (not described here).

Figure 2. Learner interacting with AutoTutor

Judging Cognitive-Affective States. Participants provided self-judgments of their cognitive-affective states immediately after the tutorial session; learning activities during the session were not interrupted. Similar to a cued-recall procedure (Rosenberg & Ekman, 1994), the judgments for a learner's tutoring session proceeded by playing a video of the face along with the screen capture video of interactions with AutoTutor on a dual-monitor computer system (see center and right monitor in Figure 2). The screen capture included the tutor's synthesized speech, printed text, students' responses, dialogue history, and images, thereby providing the context of the tutorial interaction.

Participants were instructed to make judgments on what affective states were present at any moment during the tutoring session by manually pausing the videos (called *spontaneous* judgments). They were also instructed to make judgments at each 20-second interval; the video automatically stopped every 20 seconds (called *fixed* judgments). If the learner was experiencing more than one affective state, the learner was instructed to mark each state

and indicate which was most pronounced. However, only the first choice (more prominent) affective states were included in the subsequent analyses.

Participants were provided with a checklist of seven states (boredom, flow/engagement, confusion, frustration, delight, surprise, and neutral) for them to mark along with definitions of the states. Hence, judgments were made on the basis of the participants' facial expressions, contextual cues via the screen capture, and the definitions of the cognitive-affective states.

Study 2
The participants were 30 undergraduate students (13 male and 17 female) from a mid-south university in the U.S. who participated for extra course credit.

Study 2, was similar to Study 1, but with two important differences. While participants in Study 1 interacted with the traditional typed-input version of AutoTutor, Study 2 participants spoke their responses to a new spoken-input AutoTutor. In addition to changing the input modality, there were a number of technical improvements in the new version of AutoTutor (version 3.1). These include improvements in conversational smoothness via a contextually-sensitive dialogue management module, state-of-the-art semantic and statistical natural language understanding mechanisms (Jurafsky & Martin, 2008), and an updated domain knowledge base for computer literacy.

The second difference between the two studies pertains to the retrospective affect judgment protocol. While participants in Study 1 provided affect judgments every 20 seconds and in-between each 20 second block, participants in Study 2 provided judgments at three pre-selected points plus some random points in the tutorial session. These included: (1) a few seconds after AutoTutor completed a dialogue move, (2) immediately before the learner started expressing his or her spoken response to the tutor, and (3) other randomly selected points in the dialogue. Participants provided approximately 30-35 cognitive-affective ratings at each of these three judgment points. These constituted the fixed judgment points. Similar to Study 1, the participants could stop the video at any time and make spontaneous judgments.

Results and Discussion

The retrospective affect judgment procedure yielded 2967 and 3099 self reported cognitive-affect judgments for Studies 1 and 2, respectively. A time series that preserved the temporal ordering of the cognitive-affective states was constructed for each participant. On average, there were 106 states ($SD = 9$) per time series for Study 1 and 103 states ($SD = 14$) for Study 2.

Since the goal of this paper is to investigate transitions between different states, and not persistence in the same state, the data was recoded to eliminate repetitions between states. For example, the sequence $X \rightarrow Y \rightarrow Y \rightarrow Z$ was converted to $X \rightarrow Y \rightarrow Z$. This process reduced the length of the time series to a mean of 64 states for both studies ($SD_1 =$

19, $SD_2 = 15.24$). On average, there was a state transition every 32.38 and 32.77 seconds for Studies 1 and 2, respectively ($SD_1 = 11.17$, $SD_2 = 9.58$). The recoding process did not alter the distribution of the cognitive-affective states.

Estimating Parameters of HMMs

The current analyses focused on discovering the parameters of HMMs that best explain the relationship between observable cognitive-affective states and the hidden variables that presumably govern their behavior. In particular, we estimated the parameters of an HMM with six observable states and four hidden states. The hidden states were equilibrium, disequilibrium, stuck, and disengaged, whereas the observable states were boredom, flow/engagement, confusion, frustration, delight, and surprise. Although the theory does not explicitly address the presence of delight and surprise, the states were included in the present analyses because they occasionally occur during learning sessions with AutoTutor (Graesser et al., 2006).

The present analyses constructed separate HMMs for each study from the time series of the cognitive-affective states. Parameters of the two matrices of each HMM were estimated with the Baum-Welch algorithm, which is the standard procedure used to train HMMs (Jurafsky & Martin, 2008; Rabiner, 1989). The algorithm begins with a set of initial parameters and then iteratively improves the estimates of these parameters by comparing how well the model constructed at each iteration fits the data. The algorithm converges when the discrepancy between the predictions made by the model and the training data minimally vary (i.e. within a preset threshold).

The choice of initial parameters plays an important role in the estimation process (Jurafsky & Martin, 2008). The initial parameters can be randomly seeded if there is no prior theory guiding their selection. In our case, the cognitive disequilibrium theory provides some important guidelines for initial parameter selection. For example, the theory hypothesizes that flow/engagement is expected to accompany the equilibrium state. Hence, the initial emission matrix was seeded such that the Flow|Equilibrium probability was slightly higher (.18) than the other emissions stemming from the equilibrium state. In particular emissions for Boredom|Equilibrium, Confusion| Equilibrium, etc, were set to .164 [$(1 - .180)/5 = .164$]. In this fashion, a small increase in emission probabilities was provided to confusion in the disequilibrium state, frustration in the stuck state, and boredom in the disengaged state.

The initialization process for the transition probability matrix was quite different. Here, transitions into the same hidden states were set to zero (because we are interested in modeling transitions to other states), while transitions to other hidden states were set to .333. Hence, each hidden state had an equal probability of transitioning to any other hidden state. The HMMs were seeded in this fashion to test whether hidden state transitions in the converged HMMs

aligned with predictions of the cognitive disequilibrium theory. For example, equilibrium should transition into disequilibrium more frequently then stuck and disengaged.

It should be noted that the initial distribution of hidden states were also set to .25. The initial parameters of the HMM's are listed in Table 1 (see *Init* band). HMMs initialized on the basis of these parameters converged in 30 and 29 iterations for Study 1 and Study 2, respectively.

Exploring the Structure of the Converged HMMs

Before delving into the structure of the HMMs, we first evaluated how well the HMMs captured the dynamics of the state transitions in the two sets of analyses. In the first analysis, we compared each HMM to its random surrogate, which was an HMM that was seeded with the same initial parameters but was trained on randomly shuffled time series. Random surrogate comparisons provide a convenient face-validity test for time series analyses, because random shuffling eliminates all temporal dependencies between events while preserving the priori probabilities of individual events. The results indicated that the log-likelihood (LL) for HMM's constructed on the basis of a randomly shuffled time series was significantly ($p < .05$) lower than the LL for HMMs constructed from the original time series ($d = 1.36$ and 1.33 for Study 1 and Study 2, respectively).

The second analysis focused on the generalizability of the HMMs. Here we compared HMMs constructed and validated on the entire training set to HMMs constructed on partial data sets using a leave-one-out cross validation procedure (LVOCV). LVOCV involves constructing N HMMs, where each HMM is trained on time series from $N - 1$ participants and tested on the time series of the remaining one participant. Correlations between the LL of LVOCV HMMs and HMMs trained on the entire data set were almost perfect ($r = .99$ for both Studies).

Table 1 lists the parameters of the HMMs for Study 1 and Study 2. As could be expected, the parameters of the emission matrix indicate that the flow state is emitted during equilibrium, confusion during disequilibrium, frustration when stuck, and boredom when disengaged. Hence, the converged emission matrix accurately models the hypotheses of the cognitive disequilibrium theory.

Although the transition matrix was seeded such that transitions between the hidden states were equivalent (.333), a different distribution of transitions emerged after training. In particular, consistent with the theory, the equilibrium state is more likely to transition into disequilibrium than the other states. As predicted, the disequilibrium state is more likely to transition back into equilibrium and the stuck state than the disengaged state.

The patterns were somewhat more murky for the stuck state. Although we hypothesized that stuck should transition into disengagement more frequently than equilibrium or disequilibrium, this pattern was not observed in the HMM for Study 1. The results were more in line with the theory for the HMM for Study 2, where stuck was equally likely to transition into disengagement and disequilibrium, but not

equilibrium. Finally, the theory does not explicitly address transitions from the disengaged state, and the HMMs did not reveal any clear transition pattern for this state.

It is also important to indicate that we constructed two additional HMMs for Studies 1 and 2. These HMMs were identical to the HMMs listed in Table 1 but were seeded with randomly initialized parameters instead of the theoretically derived initial parameters. The structure of these randomly-seeded HMMs were quite similar to the HMMs listed in Table 1, indicating that our theoretically derived initial parameters did not bias the models.

Table 1. Parameters of HMMs

		Emission Matrix						Transition Matrix			
	Current Hidden	Current Observed State						Next Hidden State			
HMM	State	Bor	Con	Del	Flo	Fru	Sur	Eq.	Dq.	St.	Dg.
Init	Equilibrium	.16	.16	.16	.18	.16	.16	.00	.33	.33	.33
	Disequilibrium	.16	.18	.16	.16	.16	.16	.33	.00	.33	.33
	Stuck	.16	.16	.16	.16	.18	.16	.33	.33	.00	.33
	Disengaged	.18	.16	.16	.16	.16	.16	.33	.33	.33	.00
S1	Equilibrium	.00	.00	.02	**.96**	.00	.01	.00	**.42**	.28	.30
	Disequilibrium	.00	**.94**	.02	.00	.00	.04	**.43**	.00	**.33**	.24
	Stuck	.00	.00	.15	.00	**.79**	.06	.35	.33	.00	.33
	Disengaged	**.80**	.00	.06	.00	.00	.15	.38	.31	.31	.00
S2	Equilibrium	.00	.00	.06	**.89**	.00	.05	.00	**.46**	.27	.27
	Disequilibrium	.00	**.90**	.07	.00	.00	.03	**.37**	.00	**.35**	.27
	Stuck	.00	.00	.13	.00	**.83**	.03	.29	.36	.00	.36
	Disengaged	**.90**	.00	.04	.00	.00	.06	.32	.33	.35	.00

Notes. Eq. = equilibrium*, Dq* = disequilibrium*, St* = stuck*, Dg.* = disengagement

Discussion

The present paper used HMMs to test a theory of cognitive disequilibrium that is applicable to the dynamics of cognitive-affective states in deep learning environments. The major predictions of the theory were verified via the emission and transition matrices of the HMM which aligned with different aspects of the theory. In particular, the results supported an equilibrium state that emitted flow/engagement, a disequilibrium state that emitted confusion, and transitions between the equilibrium and disequilibrium states. These results support the assertion that students in the state of engagement/flow are continuously being challenged within their zones of optimal learning (Brown, Ellery, & Campione, 1998; Vygotsky, 1978) and are experiencing two-step episodes alternating between confusion and insight.

The HMMs confirmed the presence of a transition from disequilibrium to the stuck state that emitted frustration. However, the prediction of a transition from the stuck state to the disengaged state was only partially supported. The converged HMMs suggest that in addition to the predicted transition from the stuck to disengaged states, transitions from stuck to the disequilibrium and even the equilibrium states are permissible.

These transitions from frustration suggest that it is important to differentiate between different exemplars of frustration. Similar to the discrimination between productive and hopeless episodes of confusion, there might also be different manifestations of frustration. For example, being stuck for a short period of time and then obtaining an insight might trigger delight and cause a transition into the equilibrium state. Some evidence for this assertion can be obtained from the emission matrix which indicates that delight is sometimes emitted from the stuck state. Alternatively, the stuck state can transition into the disequilibrium state when an additional impasse is detected. The third manifestation of frustration is one that is predicted by the theory. Here, persistent failure and hopelessness from being stuck will eventually trigger disengagement, where the learner detaches from the learning session.

In summary, there appear to be three alternatives for transitions from frustration and the stuck state: (a) frustration is alleviated when a resolution is reached, (b) frustration oscillates with confusion when a stuck student detects an additional impasse, and (c) frustration transitions into boredom when a hopelessly stuck learner disengages from the learning session. Testing the fidelity of these transitions will require further empirical research.

Acknowledgments

References

Berlyne, D. (1960). *Conflict, Arousal, and Curiosity*. New York: McGraw-Hill Inc.

Bower, G. (1981). Mood and memory. *American Psychologist, 36*, 129-148.

Brown, A., Ellery, S., & Campione, J. (1998). Creating Zones of Proximal Development Electronically in Thinking Practices in Mathematics and Science Learning. In J. Greeno & S. Goldman (Eds.), (pp. 341-368). Mahawah, NJ: Lawrence Erlbaum.

Csikszentmihalyi, M. (1990). *Flow: The Psychology of Optimal Experience*. New York: Harper and Row.

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., & Graesser, A. (in review). Dynamics of Cognitive-Affective States during Deep Learning.

D'Mello, S., Taylor, R., & Graesser, A. (2007). Monitoring affective trajectories during complex learning. In D. McNamara & G. Trafton (Eds.), *Proceedings of the 29th Annual Cognitive Science Society* (pp. 203-208). Austin, TX: Cognitive Science Society.

Ekman, P. (1984). Expression and the nature of emotion. In K. Scherer & P. Ekman (Eds.), *Approaches to emotion* (pp. 319-344). Hillsdale, NJ: Erlbaum.

Festinger, L. (1957). *A theory of cognitive dissonance*. Stanford, CA: Stanford University Press.

Forbes-Riley, K., & Litman, D. (2009). Adapting to student uncertainty improves tutoring dialogues. In V. Dimitrova, R. Mizoguchi & B. Du Boulay (Eds.), *Proceedings of the 14th International Conference on Artificial Intelligence in Education* (pp. 33-40). Amsterdam: IOS Press.

Graesser, A., Lu, S., Olde, B., Cooper-Pye, E., & Whitten, S. (2005). Question asking and eye tracking during cognitive disequilibrium: Comprehending illustrated texts on devices when the devices break down. *Memory and Cognition, 33*, 1235-1247.

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.

Jurafsky, D., & Martin, J. (2008). *Speech and language processing: An introduction to natural language processing, computational linguistics, and speech recognition*. Upper Saddle River, NJ: Prentice-Hall.

Kort, B., Reilly, R., & Picard, R. (2001). *An affective model of the interplay between emotions and learning*. Paper presented at the Proceedings of the International Conference of Advanced Learning Technologies, Madison, WI.

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, A. E., N. R. & L. S. (Eds.), *Proceedings of the 9th International Conference on Intelligent Tutoring Systems* (pp. 50-59).

Mandler, G. (1984). *Mind and Body: Psychology of Emotion and Stress*. New York: W.W. Norton & Company.

Ortony, A., Clore, G., & Collins, A. (1988). *The cognitive structure of emotions*. New York: Cambridge University Press.

Otero, J., & Graesser, A. (2001). PREG: Elements of a model of question asking. *Cognition and Instruction, 19*(2), 143-175.

Piaget, J. (1952). *The origins of intelligence*. New York: International University Press.

Rabiner, L. (1989). A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE, 77*(2), 257-286.

Rosenberg, E., & Ekman, P. (1994). Coherence between expressive and experiential systems in emotion. *Cognition & Emotion, 8*(3), 201-229.

Russell, J. (2003). Core affect and the psychological construction of emotion. *Psychological Review, 110*, 145-172.

Siegler, R., & Jenkins, E. (Eds.). (1989). *Strategy Discovery and Strategy Generalization*. Hillsdale, NJ,: Lawrence Erlbaum Associates.

Stein, N., & Levine, L. (1991). Making sense out of emotion. In A. O. W. Kessen, & F, Kraik (Eds.) (Ed.), *Memories, thoughts, and emotions: Essays in honor of George Mandler* (pp. 295-322). Hillsdale, NJ: Erlbaum.

VanLehn, K., Siler, S., Murray, C., Yamauchi, T., & Baggett, W. (2003). Why do only some events cause learning during human tutoring? *Cognition and Instruction, 21*(3), 209-249.

Vygotsky, L. (1978). *Mind in society: The development of higher psychological processes*. Cambridge, MA: Harvard University Press.