

4. Towards a Motivationally-Intelligent Pedagogy: How should an intelligent tutor respond to the unmotivated or the demotivated?

Benedict du Boulay

Interactive Systems Research Group, School of Informatics, University of Sussex, Brighton
BN1 9QH, UK

4.1 Introduction

This paper delineates some of the pedagogy needed by a motivationally-intelligent tutoring system. Such a system combines the expertise and knowledge of systems able to reason and react effectively at the cognitive and metacognitive levels with those able to reason and react at the affective and meta-affective levels. Three big problems face the designer of such systems. First is determining the internal motivational states of learners given their behaviours, their demeanours and what they say. Second is figuring out what might have caused that state. Third is choosing how to act or react in a way that is likely to make the situation better (Avramides & du Boulay, 2009; du Boulay, Rebolledo-Mendez, Luckin, & Martinez-Miron, 2007). The main argument of this paper is around the second and third steps: identifying the causation of negative motivational states and remediating those states.

4.1.1 Motivational states

Pintrich (2003) categorised research on motivation as falling into the three overlapping areas of “Values”, “Expectancies” and “Feelings”. Here, Values refer to the personal, social and cultural rationale that underpins the learner’s participation in the educational activity in question. Expectancies refer to the learners’ expectations of their lived experience of doing the learning, for example in terms of success or failure. Feelings refer to the emotions engendered by the learning experience: frustration (say) when a problem is hard, elation when the solution seems to appear from nowhere, or boredom when the material or the interaction is dull. Generalising from this analysis of the literature we characterize the motivational state of a learner as a triple of <Feelings, Expectancies, Values>. As time unfolds, the things that happen to learners, the things that they do and their own

reflections on these change their appraisal of the degree of fit of the Values and Expectancies components of the triple. In its turn, the motivational state helps determine the extent to which, and the method by which, the learner engages (or not) in ongoing activity that may be “constructive” or “unconstructive” (Rosiek, 2003) with respect normal educational goals. Negative motivational states are regarded as those where the causal chain of events has resulted in mismatches or violations of Values or confirmation or disconfirmation of Expectancies, so giving rise to the feeling associated with the negative motivational state and possibly also to unconstructive behaviour such as passivity or gaming the system.

A convenient way to refer to a motivational state is via the main feeling associated with it. So we can talk about the feeling of elation (say), but also of the motivational state within which elation is the main feeling. Two learners may feel equally elated, but be in different motivational states when their Expectancies and Values are different. So for example, an elated learner who rather expected to do well, will be in a different motivational state from one who expected to do badly.

Various researchers have developed ways to detect particular feelings associated with the motivational states that occur in learning. For example, these include frustration (Kapoor, Burleson, & Picard, 2007) as well as more positive feelings such as interest, excitement and confidence (Arroyo, et al., 2009). In broader terms, attempts have been made to detect learners’ overall motivation (see e.g. Johns & Woolf, 2006). In narrower terms others have detected particular symptoms of negative motivational states such as when learners engage in potentially mal-adaptive learning behaviours, e.g. “gaming the system” (for a review see, e.g. Baker, et al., 2008).

Adopting an effective pedagogic response to negative feelings or unconstructive behaviours will depend on the reason why the learner is in that motivational state or exhibiting that behaviour. For example, Baker et al. (2008) examined thirteen hypotheses as to why learners might game the system and found supporting evidence for several of them including dislike of the subject matter, lack of self-drive, and frustration with the level of the material or with the difficulty of reading it. We could add further hypotheses. For example, the learner may never have wanted to be in this class in the first place and was persuaded into it by ambitious parents. By contrast they may find the issue of seeking out the weak points in the system’s tutorial strategy just inherently more interesting than the material they are supposed to learn from the tutor. More mundanely, they might find the material just dull, or indeed too easy. By contrast the learner might have imported feelings from some event prior to logging-on to the system (a row at breakfast with mum, for instance), or may lack confidence in their ability to solve the problems posed by the system. Each of these needs to be dealt with in a different way, and that is what this paper is about.

4.1.2 Motivational intelligence for computer tutors

Several researchers argue that cognition and emotion are interwoven in learning and hard to disentangle, both for the learners themselves and for their tutor's understanding of their learning (see, e.g. Bickhard, 2003). A consequence of this is that the tutor needs to reason about both the likely emotional and the likely cognitive consequences of a tutorial intervention if it is to succeed in acting in a motivationally-intelligent way. For example, just preventing the learner from engaging in gaming behaviour by adjusting the way the help mechanism works may only succeed in encouraging the learner's (possible) frustration to emerge in other ways (Baker, et al., 2008).

The multifaceted aspects of motivation and the interwoven nature of cognition and emotion make the design and development of motivationally-intelligent tutors especially complex. The aim of this chapter is to try to tease apart some of the factors that might assist in the design of the diagnostic and remedial components of motivationally-intelligent tutors. The work described is at an early stage, without empirical support as yet. So the diagnostic part corresponds in part to the hypothesis generation stage of the work of Baker et al. (2008). The remedial part corresponds in part to the to strategy generation proposals on managing mood, attitudes, and interpersonal stances of Blanchard et al. (2009). For example, they suggest that a tutor might improve learners' attitudes to learning by considering their self-efficacy and their personal goals.

This chapter concentrates on negative motivational states as these need to be dealt with if the learner is to make good progress. Positive motivational states are also important to the tutor, not just as a goal to achieve in their right, but also as states to be recorded as potential sources of encouragement and reflective advice to the learner should things not go so well later. For example, as we see later in this chapter, one way to counter certain kinds of anxiety (say) is to remind the learner about past learning episodes where the anxiety turned out to be unfounded. Detecting positive states poses similar difficulties to detecting negative states, but ongoing good performance and effort on the learning task at hand are a good guide. Recording both positive and indeed negative motivational states opens up the possibility for the tutor to be more proactive, possibly heading off a shift towards an unwanted negative motivational state before the feelings or behaviour that would accompany it manifest themselves.

The chapter is organised into 5 sections. The next section looks briefly at the kinds of data available about motivational states and at the kinds of learner state that are normally distinguished in such systems. The main section of the paper takes three motivational states whose associated feelings are frustration, anxiety and boredom. For each of these motivational states it suggests a set of pedagogic tactics to remediate that state. These tactics respond differently to the different causal chains arising from Values and Expectancies issues. Finally there is a conclusions section with some indications of future directions.

4.2 Diagnosing Motivational States

This section looks briefly at the kinds of data potentially available to the tutor, and the temporal or pedagogical granularity with which that data is observed and considered. It also considers the kinds of motivational state into which learners are typically categorized.

4.2.1 *Kinds of data*

Much of the contemporary work is focused on broadening the bandwidth of data available to the tutor beyond what can be gleaned from the learners' responses, either in terms of dealing with the domain itself or in terms of self-reports about their cognitive, metacognitive and affective reactions to what is going on. As technology becomes more sophisticated and cheaper we find cameras being used to record focus of attention and facial expression, sensors to record skin conductance, heart rate and brain waves, pressure sensors to record posture and wriggling in the chair, linguistic analysis to infer affect (see e.g. Dong, Chapter 18 in this volume) as well as the force exerted on the mouse (see e.g. S D'Mello, et al., 2008). Each of these channels provides clues of differing quality with respect to the emotional state of the learner, with some being good for confirming some states and others being good for disconfirming other states (Arroyo, et al., 2009).

In addition to looking at the emotional state there are also clues to be found in the learner's unconstructive behaviour. In diagnosing negative motivational states, we may classify this crudely into (i) the presence of unconstructive activities that are mal-adaptive, or into (ii) the absence of constructive activity that should occur. So in terms of mal-adaptive activity we list gaming the system and other misuses of the help facility, engaging in off-topic activities such as surfing the web or using email instead of studying, or making poor choices as to the difficulty of the problems tackled (whether too easy or too hard). In terms of mal-adaptive inactivity we might list a general lack of activity at all, listlessness, passivity and failure to engage.

It is completely understandable as to why researchers are keen to find ways to ascertain learner's motivational states with as little intrusion as possible. However there are limits as to how accurately even a human teacher can gauge the learner's state, not least when the learner may wish to mask it (Balaam, Luckin, & Good, 2009).

The chapter assumes that it is going to be difficult to distinguish between some motivational states just from the sensor data and the external behaviour on their own, and just as difficult to distinguish Values issues from Expectancies issues. This suggests that the motivationally-intelligent tutoring system needs to engage in some kind of dialogue with the learner, just as a human teacher would need to

do in similar circumstances. In order to sidestep the obvious problems of NLP-based interactions, we currently favour some kind of menu-based interaction, similar to those employed elsewhere to gauge motivational states (Arroyo, et al., 2009; del Soldato & du Boulay, 1995). This might be triggered in two stages. The first stage might simply ascertain the learner's perceived motivational valence (positive or negative). Should this be negative and should the demeanour of the learner give other causes for concern, the system might then ask the learner to choose as many items from a set of menu items that might apply, somewhat along the lines of:

- I don't see why are we learning this
- This is too hard
- This is dull
- This is too easy
- Something happened outside the lesson which upset me
- I would rather be doing something else
- There is no problem, I am feeling fine
- There is a problem but it's private
- Other . . .

In addition there could be a text box for the learner to type in whatever they want. While this would not be reacted to directly by the system, it could be used by the system-designer to improve the menu items over the longer term.

4.2.2 Temporal and pedagogical granularity

Many tutors focus most of their attention on the problem at hand or at least on the current session. A few (e.g. Weber & Brusilovsky, 2001) look back to the detail of interactions in previous sessions, and none (as far as we know) anticipate the sessions yet to come and reason about how best to plan for the future *from a motivational point of view*. As tutors cover greater amounts of material and the interaction data logged by the tutor becomes more extensive, issues around the how best to exploit what may be many hours of motivational experience with a particular learner come to the fore.

4.2.3 Distinctions made amongst learner states

There are differing views as to how best to categorise the possible types of motivational state that a learner may be in. Given the tri-faceted view of motivation espoused, some of the categories are based around Expectancies, some around Values and some around Feelings. While the occurrence of pure strong emotions

in a learning situation is rare (disgust, anger, surprise), more nuanced emotional states (feelings) are common. Some researchers focus on emotions *per se* (e.g. Conati and Maclaren, see below) and some on motivational states that are designated in terms of their emotional component (see e.g. Graesser et al., below).

A simple but effective method to distinguish affective states is simply between the positive and the negative (the valence) and react both to absolute values of valence and to changes of valence (Zakharov, Mitrovic, & Johnston, 2008). In their work on understanding the phenomenon of gaming the system, Baker et al. (2008) distinguish between learner characteristics (such goals, attitudes, beliefs, general approaches and emotions) rather than motivational states as such, though there is overlap with Expectancies, Values and Feelings. In trying to calibrate the utility of different sensors, Arroyo et al. (2009) distinguish between learners who are Confident, Frustrated, Excited and Interested. Graesser and his colleagues distinguish the states of Confusion, Frustration, Boredom, Flow/Engagement, Eureka and Neutral (Graesser, et al., 2008).

Others adopt a subset of the emotion states developed by OCC theory (Ortony, Clore, & Collins, 1988), or variations on this, to reason about the causality in learning situations. So, for example, Conati and Maclaren (2005) distinguish the emotions of Joy, Distress, Admiration and Reproach as part of their approach to modeling the causes of emotion in the classroom. Following classroom-based empirical work with adolescents, Balaam (2010) asked her participants to distinguish between Happy, Tired, Proud, Bored, Nervous, Angry and Frustrated.

The purpose of the chapter is to show how a system might react effectively to a range of negative motivational states, where the reactions would attempt to deal with the underlying causes. So we examine Frustration, Boredom and Anxiety.

4.3 Motivational Pedagogy

This section looks at the three different negative motivational states whose main feeling is experienced by the learner as: frustration, anxiety or boredom. In each case we examine that state from the point of view of Values and Expectancies and for each of these we sketch possible causes of that state together with possible remedial actions that might be taken by the tutor.

It is worth stressing that learning difficult material can be hard work and that solving tricky problems can be frustrating. Being frustrated or anxious about outcomes (say) is a natural aspect of learning. Indeed as Pekrun points out (see Chapter 3 in this volume) the same emotion can have either an “activating” or a “deactivating” effect, depending on how the learner appraises the situation. The motivationally-intelligent tutor will recognize this by putting more effort into helping the learners become more aware of these issues themselves and assist them to manage these feelings, than into changing the flow of activity so as to avoid situations that lead to these feelings (Avramides & du Boulay, 2009).

We distinguish two kinds of causes of a transition towards a negative (or more negative) motivational state:

1. **Values-based:** underpinning most formal educational situations there will be a set of values around desirable learner behaviours and learner outcomes. To the extent that the learner “goes along” with these all may be fine, but there may be times where the learner cannot understand the value of a particular activity or outcome, or if it can be understood, its personal value cannot be appreciated. For example, a learner who really cannot see the point of learning about Pythagoras’ Theorem will feel at odds with a situation where that is the goal. So we concentrate on mismatched or violated Values between those of the learner and those of the educational situation.
2. **Expectancies-based:** learners have expectations about how well or badly they are going to succeed in an educational activity and whether it will conclude with a successful or unsuccessful outcome. They will also have views about how much agency and control they are in a position to exert and also about the nature of learning and skill acquisition. For example they may limit their effort and misinterpret errors as evidence not only that they cannot exercise some skill, *but cannot imagine ever coming to be able to acquire it*. So we concentrate on confirmed and disconfirmed Expectancies.

It is important to note that feelings may also be imported from external or past events. First are feelings not directly linked to the learning situation *per se*. For example, a learner may be angry about events that happened outside the classroom, or may be anxious about some future event unconnected with the learning in hand: for example, an ongoing feud with another learner in a different class. Rather than seeking causes in terms of Values and Expectancies within the learning, the tutor would need to help the learners distance themselves from such external causes of the negative state, if an optimal state for learning is to be maintained. This might even involve abandoning the learning activity for the moment to give time and space for this distancing.

By contrast feelings may be imported from the recollection of previously experienced similar learning situations. A learner who has experienced anger, anxiety or frustration about mathematics (say) in the past may well re-experience these feelings in a new, but apparently similar, learning situation. In this case the tutor should reason about the Values and Expectancies of the learner to anticipate this kind of possibility and attempt to forestall the development of the feelings afresh, e.g. by reassurance, or by reference to positive motivational states previously experienced.

The system should be able to gain some diagnostic leverage by examining the time signature of the onset of a motivational state. In the tables below we look back simply at the motivational valence (positive/negative) of the learner at the beginning of the session, and on average in the previous session. This helps to distinguish those motivational states that can become negative quite quickly (frustration) from those where the build-up is likely to be slower (anxiety, boredom).

Knowing something of the personality of the learner will also help identify the state in terms of the nature of the behavioural consequences (symptoms) of that state. In the tables that follow, motivational valence is designated as “-ve” for negative, “neutral” for neutral, and “+ve” for positive. “Any” means that the motivational valence can take any value.

4.3.1 Anxiety

Pekrun (see Chapter 3, this volume) provides a detailed account of the different antecedents and different consequences of anxiety. He also emphasizes, as we do, that the antecedents of any particular emotion can vary from one individual to another, and the consequences for the learner’s behaviour also vary. Here we distinguish anxiety arising from issues around Values from that arising from issues around Expectancies (see Table 1).

Table 1: Anxiety

Valence start of session	Valence previous session	Motivational Facet	Possible Cause	Remedial Possibility
Neutral or -ve	Neutral or -ve	VALUES	Ongoing mismatch of values.	Try to re-orientate values.
Neutral or -ve	Any	EXPECTANCIES	Work has started to look too difficult and there is fear of failure.	Reassure by finding evidence to the contrary <i>or</i> make the task easier <i>or</i> offer more domain level support.

Where there is anxiety arising from a mismatch of values, this may focus on the alternative activities that the learner might have been engaged in rather than the ones that she does not value. A way to try to deal with this is to help the learner either to value the current activity or the overall goals of the learning within which the current activity is situated. By contrast, anxiety arising from lack of confidence or fear of failure requires a different approach in terms of reassurance and support.

4.3.2 Frustration

Frustration (see Table 2) is often associated with a greater degree of arousal than anxiety considered earlier and so may pose different kinds of remedial pres-

sure (Russell, 1980). In considering the Values aspect, the learner may have something else in mind to do, rather than simply not seeing the point of the current activity. Frustration is also likely to have few precursors from earlier in the session or from the previous session, but arise out of a specific activity. At the Expectancies level the learner can get frustrated if the work is too hard or indeed too easy, so these need to be distinguished in order to take sensible steps.

Table 2: Frustration

Valence start of session	Valence previous session	Motivational Facet	Possible Cause	Remedial Possibility
Any	Any	VALUES	Would rather be doing something else	Discuss comparative value of two activities. Discuss re-scheduling two activities.
Any	Any	EXPECTANCIES	Work has started to look too difficult or too easy	<i>Too hard:</i> Reassure by finding evidence to the contrary <i>or</i> make the task easier <i>or</i> offer more domain level support <i>Too easy:</i> make the work more challenging
			Unable to exercise personal choice in the current activity	Offer more control

Frustration about the work being too easy is perhaps more properly regarded as a Values issue, arising from the sense that the learner's time is not being well used.

4.3.3 Boredom

Like frustration, boredom in educational settings has a higher degree of arousal than anxiety (Sidney D'Mello & Graesser, 2010), though not all agree (Russell, 1980). So one may observe various mal-adaptive activities like gaming the system, chatting to other learners, or being a nuisance. In terms of behavioural cues

these may be just the opposite of listlessness and lack of effort – though the effort may well be misdirected (see Table 3).

Table 3: Boredom

State start of session	State previous session	Motivational Facet	Possible Cause	Remedial Possibility
Neutral or -ve	Any	VALUES	Would rather be doing something else	Discuss comparative value of two activities. Discuss re-scheduling two activities.
			Cannot see the point of the current activity	Try to re-orientate values.
Any	Any	EXPECTANCIES	Work has started to look too easy	Make the work more challenging or add interest and excitement.
			Unable to exercise personal choice in the current activity	Offer more control.

4.4 Discussion

In terms of Values the main remedial method is to attempt to align (or realign) the learner's values with those inherent in the course being taken. In terms of Expectancies, it is important to distinguish the learner's realistic expectancies from unrealistic ones. Unrealistically negative expectancies may be countered by evidence of success in similar circumstances in the past. Unrealistically positive expectancies can either be ignored for the moment or a note of caution suggested, depending on the likely impact of failure on that learner. Realistically negative expectancies may be dealt with by negotiating over whether the task difficulty should be adjusted. Realistic positive expectancies can be affirmed.

In terms of feelings, it is important to distinguish feelings that have been imported into the learning situation from outside from those emanating directly from the learning situation itself. In the former case it may be possible to acknowledge the feelings arising from outside while trying to minimise their effects within the lesson. For feelings arising directly from the learning situation it will be important to decide whether the feelings are well-founded (similar to the realistic/unrealistic distinction for Expectancies). Of course, the feelings need to be acknowledged

whether or not they are well-founded as we assume that the learner is not misrepresenting how she feels. For example, a learner who feels (well-foundedly) ashamed over a poor performance may be consoled and a strategy put in place to improve performance. However in dealing with learners who feels (ill-foundedly) ashamed over a perfectly adequate performance effort may be devoted to helping them to take a more realistic view of their own and others' performances.

4.5 Conclusion

This chapter has argued that dealing with poorly motivated learners requires knowledge of the causes of the negative feelings and mal-adaptive behaviour associated with their motivational state. So, noting that the learner is gaming the system, and determining that they are frustrated (say) is just the first step. One needs to work back to the causes of that frustration in terms of Expectancies or Values in order to have some hope of deploying a remedial action that may make the situation better. In working back to the causes one is very likely to have to go beyond simply observing learners' behaviours and demeanours to find out directly from them why they believe that they feel and act as they do.

Making the causation behind learner demotivation explicit is potentially beneficial to the learner as well as to the motivationally-intelligent tutor. From the learner's point of view, attempting to be explicit about the causation of the poor motivation and then experiencing the remedial tactic suggested by the tutor should build up the learner's own understanding of motivation, in other words, improve his or her meta-motivational insight. For the tutor, each of these episodes of diagnosis and remediation can itself become an example that can be used with the learner later: "remember when you said that you were feeling and you tried doing well I think you may be in a similar situation again". Referring to a past incident like this should further increase the learner's meta-motivation.

The pedagogic tactics in the tables above have not yet been implemented in a working system. Once deployed, they could be evaluated in terms of process measures such as decreased frequency or severity of negative motivational states, improved persistence in problems-solving or following set-backs, decreased off-topic or mal-adaptive learning behaviour, more sensible choice of problem difficulty, and so on. In terms of outcome measures one might expect benefits such as increased learning gain, improved willingness to engage with future learning and increased meta-motivational insight.

Acknowledgements: I thank Katerina Avramides, Madeline Balaam, Martin van Zijl and the Intelligent Computer Tutor Group at University of Canterbury, New Zealand for many helpful comments on a draft of this paper.

References

- Arroyo, I., Cooper, D. G., Bursleson, W., Woolf, B. P., Muldner, K., & Christopherson, R. (2009). Emotion Sensors Go to School. In V. Dimitrova, R. Mizoguchi, B. du Boulay & A. Graesser (Eds.), *Artificial Intelligence in Education. Building Learning Systems that Care: from Knowledge Representation to Affective Modelling* (Vol. Frontiers in Artificial Intelligence and Applications 200, pp. 17-24): IOS Press.
- Avramides, K., & du Boulay, B. (2009). Motivational Diagnosis in ITSs: Collaborative, Reflective Self-Report. In V. Dimitrova, R. Nizoguchi, B. du Boulay & A. Graesser (Eds.), *Artificial Intelligence in Education. Building Learning Systems that Care: From Knowledge Representation to Affective Modelling* (Vol. Frontiers in Artificial Intelligence and Applications 200, pp. 587-589): IOS Press.
- Baker, R., Walonoski, J., Heffernan, N., Roll, I., Corbett, A., & Koedinger, K. (2008). Why Students Engage in "Gaming the System" Behaviours in Interactive Learning Environments. *Journal of Interactive Learning Research*, 19(2), 185-224.
- Balaam, M., Fitzpatrick, G., Good, J., & Luckin, R. (2010). *Exploring Affective Technologies for the Classroom with the Subtle Stone*. Paper presented at the Proceedings of the 28th international conference on Human factors in computing systems (CHI 2010).
- Balaam, M., Luckin, R., & Good, J. (2009). Supporting affective communication in the classroom with the Subtle Stone. *International Journal of Learning Technology* 4(3-4), 188-215.
- Bickhard, M. H. (2003). An integration of motivation and cognition *BJEP Monograph Series II, Number 2 - Development and Motivation* (Vol. 1, pp. 41-56): British Psychological Society.
- Blanchard, E. G., Volfson, B., Hong, Y.-J., & Lajoie, S. P. (2009). Affective Artificial Intelligence in Education: From Detection to Adaptation. In V. Dimitrova, R. Mizoguchi, B. du Boulay & A. Graesser (Eds.), *Artificial Intelligence in Education. Building Learning Systems that Care: from Knowledge Representation to Affective Modelling* (Vol. Frontiers in Artificial Intelligence and Applications 200, pp. 81-88): IOS Press.
- Conati, C., & Maclaren, H. (2005). Data-driven refinement of a probabilistic model of user affect. In L. Ardissono, P. Brna & A. Mitrovic (Eds.), *User Modeling 2005, 10th International Conference, UM 2005, Edinburgh, Scotland, UK, Proceedings* (Vol. Lecture Notes in Artificial Intelligence 3538, pp. 40-49): Springer.
- D'Mello, S., & Graesser, A. (2010). Mining Bodily Patterns of Affective Experience during Learning. In R. S. J. d. Baker, A. Merceron & P. I. J. Pavlik (Eds.), *Proceedings of the 3rd International Conference on Educational Data Mining* (pp. 31-40). Pittsburgh.

- D'Mello, S., Jackson, T., Craig, S., Morgan, B., Chipman, P., White, H., et al. (2008). *AutoTutor Detects and Responds to Learners Affective and Cognitive States*. Paper presented at the Workshop on Emotional and Cognitive Issues at the International Conference on Intelligent Tutoring Systems.
- del Soldato, T., & du Boulay, B. (1995). Implementation of Motivational Tactics in Tutoring Systems. *International Journal of Artificial Intelligence in Education*, 6(4), 337-378.
- du Boulay, B., Rebolledo-Mendez, G. R., Luckin, R., & Martinez-Miron, E. (2007). Motivationally Intelligent Systems: diagnosis and feedback. In R. Luckin, K. R. Koedinger & J. Greer (Eds.), *Artificial Intelligence in Education: Building Technology Rich Learning Contexts that Work* (Vol. Frontiers in Artificial Intelligence and Applications 158, pp. 563-565): IOS Press.
- Graesser, A. C., D'Mello, S. K., Craig, S. D., Witherspoon, A., Sullins, J., McDaniel, B., et al. (2008). The Relationship Between Affective States and Dialog Patterns During Interactions With Auto Tutor. *Journal of Interactive Learning Research*, 19(2), 293-312.
- Johns, J., & Woolf, B. P. (2006). *A Dynamic Mixture Model to Detect Student Motivation and Proficiency*. Paper presented at the Proceedings of the 21st national Conference on Artificial Intelligence (AAAI-06).
- Kapoor, A., Burleson, W., & Picard, R. W. (2007). Automatic prediction of frustration. *International Journal of Human-Computer Studies*, 65(8), 724-736.
- Ortony, A., Clore, G. L., & Collins, A. (1988). *The Cognitive Structure of Emotions*. : Cambridge University Press.
- Pintrich, P. (2003). Motivation and Classroom Learning. *Handbook of Psychology: Educational Psychology*, 7, 103-122.
- Rosiek, J. (2003). Emotional Scaffolding: An Exploration of the Teacher Knowledge at the Intersection of Student Emotion and the Subject Matter. *Journal of Teacher Education*, 54(4), 399-412.
- Russell, J. A. (1980). A Circumplex Model of Affect. *Journal of Personality and Social Psychology*, 39(6), 1161-1178.
- Weber, G., & Brusilovsky, P. (2001). ELM-ART: An Adaptive Versatile System for Web-based Instruction. *International Journal of Artificial Intelligence in Education*, 12(4), 351-384.
- Zakharov, K., Mitrovic, A., & Johnston, L. (2008). Towards Emotionally-Intelligent Pedagogical Agents. In B. P. Woolf, E. Aïmeur, R. Nkambou & S. L. Lajoie (Eds.), *Intelligent Tutoring Systems, 9th International Conference, ITS 2008, Montreal, Canada, Proceedings* (Vol. Lecture Notes in Computer Science 5091, pp. 19-28): Springer.