International Journal of Artificial Intelligence in Education XX (20XX) NN–NN IOS Press

Towards Systems That Care: A Conceptual Framework based on Motivation, Metacognition and Affect

Benedict du Boulay, Human Centred Technology Research Group, University of Sussex, Brighton, UK B.du-Boulay@sussex.ac.uk

Katerina Avramides, Rosemary Luckin, London Knowledge Laboratory, Institute of Education, UK {K.Avramides, R.Luckin}@ioe.ac.uk

Erika Martínez-Mirón, Departamento de Tecnologías de la Información, CCADET -UNAM, Mexico erika.martinez@ccadet.unam.mx

Genaro Rebolledo Méndez, Facultad de Estadística e Informática, Universidad Veracruzana, Mexico GRebolledo@uv.mx

Amanda Carr, Department of Psychology, Roehampton University, London, UK Amanda.Carr@roehampton.ac.uk

Abstract: This paper describes a Conceptual Framework underpinning "Systems that Care" in terms of educational systems that take account of motivation, metacognition and affect, in addition to cognition. The main focus is on *motivation*, as learning requires the student to put in effort and be engaged, in other words to be motivated to learn. But motivation is not the whole story as it is strongly related to metacognition and affect. Traditional intelligent educational systems, whether learner-centred or teacher-centred in their pedagogy, are characterised as having deployed their intelligence to assist in the development of the learner's knowledge or skill in some domain. They have operated largely at the cognitive level and have assumed that the learner is already able to manage her own learning, is already in an appropriate affective state and also is already motivated to learn. This paper starts by outlining theories of motivation and their interactions with affect and with metacognition, as developed in the psychological and educational literatures. It then describes how such theories have been implemented in intelligent educational systems. The first part of the Conceptual Framework develops the notion of a partial hierarchy of systems in terms of their pedagogic focus. These range from traditional, cognitively intelligent systems, essentially concerned with cognition up to "Systems that Care". Intermediate classes of system include Metacognitively Intelligent systems. Affectively Intelligent systems and Motivationally Intelligent systems. The second part of the Conceptual Framework is concerned with the design of systems. This is characterised in terms of (i) the kinds of diagnostic input data (such as the learner's facial expression offering clues as to her demeanour) and (ii) the repertoire of tactical and strategic pedagogic moves (such as offering encouragement), applicable at different levels of the hierarchy. Attention is paid to metacognition, meta-affect and meta-motivation covering the capability of both the learner and the educational system to understand, reason about and regulate cognition, affect and motivation. Finally, research questions and areas of further work are identified in theory development, the role of the meta levels, and design considerations.

1560-4292/08/\$17.00 © 20XX - IOS Press and the authors. All rights reserved.

INTRODUCTION

Various researchers are building educational systems that attempt to take the learner's motivational, metacognitive and/or affective state dynamically into account (see e.g., Paiva, Prada, & Picard, 2007). For example, AutoTutor (S. D'Mello, Graesser, & Picard, 2007; Graesser, et al., 2008) has been instrumented to observe the facial expression and posture of the learner in an attempt to move the learner towards (or maintain the learner in) a positive learning state such as "engaged", and also to move the learner away from negative learning states such as "bored". In a similar vein, Kapoor and others (2007) have developed a system that is able to make a good estimate of when a learner is "frustrated" based on a number of features such as hand pressure by the learner on the mouse. These technologies are also now finding their way into mainstream education (Arroyo, Cooper, et al., 2009; Dragon, et al., 2008) and the more general issue of the affective dimension of human computer interaction is well established (Picard, 2000).

These kinds of system raise a number of theoretical and practical questions: What kinds of data are available on which to make inferences about the motivational, metacognitive and affective states of the learner? What is the nature of the theory that links such data to inferred motivational, metacognitive and affective states? What kinds of motivational states are to be distinguished, one from another? What are the relationships between learning and either relatively stable personality traits or relatively transient motivational states and feelings, or less transient affective states such as moods? What are the predictable trajectories between affective states over the duration of a lesson or of a course? And what is the nature of the theory that determines how a caring system might best assist the learner to move away from trajectories or states that might inhibit learning towards those that might enhance it? In other words, what is the nature of the theory that helps the learner follow a trajectory of states that enhances and opens new possibilities (see e.g., S. D'Mello, Person, & Lehman, 2009; Kort & Reilly, 2002), even if there are individual negative episodes along the way, as opposed to a trajectory that limits possibilities and is dysfunctional or maladaptive, even if there are individual positive states along the way?

Much of the detailed practical work in this area has emerged out of the field of Artificial Intelligence in Education where the tradition of building educational systems that model and react to the state of the knowledge and skill of the learner is central. This paper briefly reviews those intelligent educational systems and intelligent learning environments that go beyond simply modelling knowledge and skill but also attempt to take the learner's motivational, metacognitive and affective states *dynamically* into account in the way they interact. The main contribution of this paper is the development of a conceptual framework with which to view these kinds of system in terms of what they might seek to achieve and how they might seek to achieve it. A number of systems are mentioned as part of this framework to illustrate different aspects of it, but the paper does not provide an exhaustive review of all systems and related work in this area.

According to Lepper et al., expert human teachers include among their goals "first, to sustain and enhance their students' motivation and interest in learning, ... and second, to maintain their pupils' feelings of self-esteem and self-efficacy, even in the face of difficult or impossible problems" (Lepper, Aspinwall, Mumme, & Chabay, 1990, p. 219). Note that two different but related issues are addressed here: one concerned with motivation specifically, the other concerned with feelings. One goal of designing caring systems is to try to understand how exactly to do this: how to interweave the motivational and affective with cognitive and metacognitive tactics so as to try to reproduce this kind of human expert behaviour. As De Rosis (2001) points out, affective issues are linked to learner goals and to their beliefs about the achievability of those goals, are time-dependent, are influenced by context, depend on the internal state of the student, and are mutually interdependent so modelling them is both complex and uncertain. So how can a system take the above motivational and affective issues into account? For example, how could a system distinguish a "clever, confident but lazy student" from a "clever, anxious and hard-working one", and even if it could make this distinction how should its behaviour towards these two kinds of student differ? This distinction takes on further force if we hope that the overall educational goals of a system could be to help the student improve their metacognitive and meta-affective capability as learners and their willingness and ability to engage effectively in further learning activities.

The structure of the paper

We start by examining theories of motivation and their interactions with metacognition and affect from a psychological and pedagogical point of view. In particular we draw on the work of Pintrich (2003) in terms of his analysis of motivation in terms of "values", "expectancies" and "affect"¹. The paper then reviews *operational models* of metacognition, motivation and affect as implemented and adapted to building educational systems, ranging from relatively complex models of the cognitive appraisal of emotion to simple models based round a small number of key motivational variables. These operational models embody the reasoning capability of the systems. For example, a system whose rationale is based on notions of learner-independence and learner-confidence (as well as learner-performance and skill) will have, in principle, pedagogic tactics and strategies that apply to these two motivational variables, but will not be able to reason about anxiety (Pekrun, Goetz, Titz, & Perry, 2002) unless that joins independence and confidence as a "first class" notion within the system.

Out of this background we develop a conceptual framework². This has two aspects. The first is concerned with the pedagogic focus of educational systems. For example, we distinguish systems that aim to manage the affective state of the learner from those that aim to manage her motivation, though clearly the two are linked. The second aspect of the conceptual framework is concerned with the design of systems in terms of the categories of diagnostic data they might use and the kinds of pedagogic tactics they might deploy. Each category is illustrated with one or more systems that exemplify work in that category, though no attempt is made to offer a comprehensive review of all such systems. The final section of the paper draws conclusions and sets an agenda for future work.

MOTIVATION, METACOGNITION AND AFFECT

This section outlines theories of motivation from the psychological and educational literature and briefly examines the relationship between motivation, affect and metacognition. The purpose is not to review the motivation literature *per se*, but (a) to show the diversity in the definitions of motivation, (b) to illustrate the difficulty in translating motivational theory into system design, and (c) to highlight that, while motivation is intertwined with affect, the two need to be clearly distinguished in AIED.

¹ This section of the paper draws heavily on (Avramides & du Boulay, 2009)

² This section of the paper draws heavily on and develops (du Boulay, Luckin, Martinez-Miron, Rebolledo-Mendez, & Harris, 2008; du Boulay, Rebolledo-Mendez, Luckin, & Martinez-Miron, 2007)

Theories of motivation

Motivation can be broadly defined as the force behind action that explains why a person acts in a particular way. Beyond this broad definition, a number of theoretical frameworks have tried to define the components of motivation and explain what determines it. The following overview is based on Pintrich's (2003) framework for integrating the literature on motivation, as he tried to incorporate different theoretical frameworks and identify the central components of motivation (in an educational context). (For a more detailed discussion of the literature see e.g., Murphy & Alexander, 2000; Pintrich, 2003; Schunk, Pintrich, & Meece, 2008). Pintrich (2003) identifies three motivational components that are present across motivational theories (though the conceptualisation of each varies): beliefs about one's ability to perform a task (*expectancy* component), beliefs about the value of the task (*value* component), and affective reactions to the task (*affective* component). Only the first two are considered in this section. The relationship between motivation and affect is discussed more generally later.

According to Pintrich's analysis, the expectancy component has been considered in two senses: beliefs about the control one has over the outcome of the task (or one's environment more generally) and beliefs about one's efficacy. Pintrich draws a broad conclusion from this research: believing that one has control over the outcome of a task (e.g. "if I study hard I will get a good grade") leads to higher cognitive engagement and performance. In contrast, having a low belief in one's degree of control leads to a low outcome. The notion of self-efficacy is related to control but is less stable and varies depending on the task and environment. There is strong evidence that self-efficacy beliefs are related to learning and performance (Bandura, 1997; Schunk, et al., 2008) believing that one is able to perform a task is strongly related to high performance and learning.

The value component of motivation in Pintrich's analysis is broken down into two central components: goal orientation and task value. People's goal orientation has typically been defined in terms of two broad orientations, (though this conceptualisation varies, e.g., Ames, 1992; Boekaerts, de Koning, & Vedder, 2006; Dweck & Leggett, 1988): an orientation towards increasing competence (mastery orientation) or an orientation toward increasing performance relative to others (performance orientation). Evidence suggests that the former leads to higher performance and learning, but results regarding the latter are mixed. Specifically, a distinction is made between being orientated towards achieving high performance (approach) in contrast to avoiding low performance (avoidance). There is some evidence to suggest that having an approach performance orientation leads to high achievement and learning, whereas an avoidance performance orientation leads to low learning outcomes (e.g. Harackiewicz, Barron, & Elliot, 1998). The other element of the value component of motivation, task value, has been defined by Eccles (1983) in terms of three components: how important the task is for the individual, their personal interest in the task, and their perception of the utility of the task for future goals. Evidence suggests that the higher the perceived value of a task, the higher the engagement and learning outcome. The notion of value applies to both intrinsic and extrinsic motivation (Ryan & Deci, 2000).

Motivation and affect

Motivation and affect are closely intertwined in a bidirectional relationship. For example, if I perform well on an exam, I am likely to feel positive, which in turn is likely to increase my motivation to study for the next exam, which is likely to lead to a high outcome. However, it is important to distinguish

between the two, as a positive affective state is neither necessary nor sufficient for high motivation and learning. The learning process may involve and even require negative affective states (for example, the frustration associated with problem-solving). It is the learner's motivation (e.g. task value or self-efficacy) that will determine how they react to those states (e.g. whether and how hard they persevere). Moreover, a positive affective state does not necessarily imply that a learner will be motivated to engage in increasing their competence; they may be content with avoiding having their performance compared to others'.

It is particularly important to make this distinction in developing *motivationally* intelligent tutoring systems, as this endeavour has been closely linked to the development of *affectively* intelligent tutoring systems. A system that can detect and react to a learner's affective state is not necessarily motivationally intelligent. For example, a learner who is anxious may be helped by the system reacting to the affective state *per se*, e.g. through reassurance, but if the anxiety is due to low self-efficacy and an avoidance performance orientation, then the impact of the reassurance may be temporary and relatively ineffective. It may calm them and, therefore, help them focus, but it will not push the learner to engage in increasing their understanding. That is not to say that a positive affective state will not have a positive influence on their learning. A motivationally intelligent tutoring system must also be affectively intelligent. But it must go beyond that. It is not enough to assess whether a learner is in a positive state and engaged in interacting with the system. The system must be able to diagnose the nature of the learner's engagement, e.g. their valuation of the task and their expectations, in order to accurately assess their motivational state and to react to it effectively.

Motivation and metacognition

Meta-cognition is normally regarded as knowledge about what we know. In relation to learning this means both our ability to *monitor* how well we understand something as well our ability to *regulate* our learning activities (Flavell, 1979). So, for example, someone who deliberately engages in self-explanation while they learn new material would be showing evidence of well-developed meta-cognitive self-regulatory ability (Conati & Van Lehn, 2000).

Motivation and metacognition are also closely intertwined in a bidirectional relationship. For example, if I perform well on an exam, I am likely to increase my belief in my mastery of the exam material, and the effectiveness of my learning strategies, which will increase my self-efficacy and motivation to study for the next exam, which is likely to lead to a high outcome. In Pintrich's (2003) terms, the "expectancy" component of motivation is closely linked to metacognition through the notion of self-efficacy. Classically metacognition is defined in terms of "knowing what one knows" as well as in the ability to regulate one's learning (Flavell, 1979). That is, it covers both the declarative notion of understanding what we do and do not understand or we can do or cannot do, as well as the procedural notion of being able to apply effective tactics and strategies to improve that understanding or that skill.

OPERATIONAL MODELS OF MOTIVATION

As illustrated earlier, motivation is a multi-faceted construct that is determined by many factors. Moreover, as discussed in (Murphy & Alexander, 2000), there are many terms to denote closely related constructs, which makes it difficult to integrate the literature. More generally, there is no

overarching theory on how these different elements of motivation interact. Motivation research is also still at early stages in terms of our understanding of how motivation impacts the learning *process* and how tutors and school environments can foster a constructive motivational state. In terms of motivational diagnosis, there is the issue of self-report (either direct or indirect) on which the majority of research on motivation is based. How reliable are self-report measures? Another issue is that of generality and context-dependency of self-efficacy and control beliefs. How stable across contexts are these beliefs? Moreover, the scientific understanding that has been developed in this field is not formulated in formal terms that can be easily applied to the design of tutoring systems (Herrington & Herrington, 2005).

The design of motivationally intelligent tutoring systems has followed several approaches. Three aspects of these approaches are examined: (i) *how motivation is defined*, (ii) *what information is used in order to diagnose* a learner's motivational state, and (iii) what is *the nature of the motivational pedagogy* that is applied. This is not intended to be a comprehensive overview of all approaches. The purpose is to draw out important design issues.

Motivational ontology

A fundamental issue in operational models is how motivation is conceptualised. Del Soldato and du Boulay (1995) and de Vicente and Pain (2002) specify the components of motivation (such as confidence and effort) on which motivational diagnosis is based. However, much research is based on more open and pragmatic definitions of motivation. For example, Arroyo and Woolf (2005) state that their approach "merges motivation, learning, and misuse of tutoring systems in one single Bayesian model [...] advocating for data-driven models that integrate cognition, motivation and their expression with different behavioral patterns" (p. 33). They do not confine themselves to a narrow definition of motivation, and the questionnaire instrument they used to diagnose learners' motivation and attitudes included generic questions that reflect many components of learners' motivation. For example, learners were asked to state the extent to which the following statement was representative of their use of the system: "I just wanted to get the session over with, so I went as fast as possible without paying much attention". A learner might have wanted to get the session over with because she did not believe she could do any of the activities well and wanted to avoid performing poorly, but she could also have wanted to get the session over with because she found it boring. The system may be able to accurately diagnose when the learner is not paying much attention, but it will fail in terms of diagnosing the learner's *motivational state* and reacting appropriately to it. Moreover, learners' attitudes are likely to vary during their use of the system. Other indicators were also used, such as requests for help and timing of help (e.g. before or after making an attempt to solve the problem). However, as discussed by de Vicente and Pain (2002), the relation between such interactional data and motivation needs to be validated.

There are also implicit assumptions about the connection between affect and motivation. For example, in the work on AutoTutor, D'Mello et al. (2008) discuss how deep learning involves negative affective states, such as frustration, and confusion. Once in these states, the learner is assumed to be unmotivated. The envisioned intelligence in AutoTutor appears to seek to react to negative affective states by trying to change them to positive ones through increasing engagement or challenge. There is no reference to the motivational state of the learner, which will determine how they react to those negative states and to what extent they will be engaged in learning. D'Mello et al. (2007) do consider learning goals in the design of the system's *reaction* in order to engage the learner.

However, there is no *diagnosis* of the nature of the learner's goals. Moreover, the basis of the tutor's reaction is not completely clear. It is proposed that by having the tutor display empathy the learner will be more likely to adopt the learning goals put forth by the tutor.

A similar criticism also applies to the implicit consideration of motivation in OCC theory (Ortony, Clore, & Collins, 1988) which has formed the basis of much research on ITSs. The basis of this theory is that learners' (or people's more generally) affective states arise from their reaction to goals, events, and agents. The learner is assumed to be motivated to achieve a set of goals. However, it does not take into account variables relating to the nature of these goals which are important in the learning process. This may be appropriate in the context of a game, as in Conati and Maclaren's (2005) research. In this context a learner may well have well-defined goals that are achieved or not, such as to win. But in a more general context of learning with tutoring systems, a learner's motivation to engage with the system requires a more complex definition.

Other researchers start from a different theoretical position on emotion to refine down to emotions relevant to learning but without such a clear-cut emotional causality as in OCC. (S. D'Mello, et al., 2007), for example, refine the Basic Emotions of Ekman and his colleagues (1972) such as anger, fear and happiness to a more learning-centric list. Others have worked empirically by questioning students about the emotions they experience in class, in studying and around assessment (Pekrun, et al., 2002). We note that Pekrun and his colleagues found as many examples of positive emotions during learning as of negative, and of the negative, anxiety was the one most commonly experienced.

Motivational diagnosis

De Vicente and Pain (2002) specifically consider the issue of motivational diagnosis. Their definition of the components of motivation is based on (Keller, 1983) and (Malone & Lepper, 1987). A distinction is made between 'trait' variables, which are assumed to remain stable throughout the learning session and 'state' variables, which are assumed to vary. The motivational state of the learner is inferred solely from interaction with the tutoring system (e.g. mouse movements and use of help), which also incorporates a facility for self-reporting one's motivational state. A set of rules was developed for inferring a learner's motivational state based on work with human tutors. The tutors inferred a learner's state based on pre-recorded interactions. The results were integrated into a set of rules that could then be implemented in the system.

A different approach was taken by del Soldato (1995). The theoretical basis was similar in terms of the components of the learner's motivational state (Keller, 1983; Malone & Lepper, 1987). The diagnosis was made along three components (effort, confidence, and independence) and was based on the degree of persistence the learner showed in solving problems (effort), the learner's self-reported degree of confidence in solving a problem before attempting it (confidence), and the learner's use of help (independence). The system reacted with comments, encouragement, provision of help, or choice of activity (e.g. a more challenging problem). These reactions were determined on the basis of a set of production rules that fired in response to the values of the three variables.

In contrast to the above, the approach taken by Arroyo and Woolf is more directly data-driven (Arroyo & Woolf, 2005). They used log data from a tutoring system to explore the relationships between learners' observable interaction with the system and their answers to a retrospective questionnaire relating to their attitudes and motivation (though the definition of motivation on which these questions are based is not always clear). For example, the questionnaire asked learners to specify

how seriously they tried to learn from the system. Arroyo and Woolf built a Bayesian Network in order to diagnose a learner's 'hidden' variables (obtained through the retrospective questionnaire) from their observable interaction with the system.

A similar approach is taken by Conati and her colleagues. They have developed an emotional model that guides the actions of a pedagogical agent in the context of an arithmetic game (Conati & Zhou, 2002; Manske & Conati, 2005). Because the context is a game, they feel that it is inappropriate to break off the interaction and request input from the learner about how she feels. So they are devising a means to determine the learner's affective state from external clues. At present their affective user model has been developed separately from a learning model but is being integrated with it (Conati & Manske, 2009; Manske & Conati, 2005). The researchers have used a subset of 6 of the 22 emotions delineated in the OCC theory (Ortony, et al., 1988). These are divided into three subsets focusing on different issues, and each expressing an emotional dimension: joy/distress about the current state of the game; pride/shame for the learner's own performance so far; and admiration/reproach for the behaviour of the agent³.

Students are characterized in terms of a number of standard personality traits: neuroticism, agreeableness, conscientiousness and extraversion and they are assumed to entertain a range of goals such as: have fun, avoid falling, beat partner, learn math[s] and succeed by myself. When playing the game their actions are logged under a number of headings such as: use of the game tools available, speed of play, requests for help, use of that help, and quality of moves.

The above factors (the traits and goals) are linked together in a Dynamic Bayesian Network (essentially a more general form of a Hidden Markov Model). Establishing exactly which nodes in the network should be linked to which, as well as the initial probability values between nodes was derived from an analysis of logfiles of users with the system as well as by triangulating this data via offline evaluations. One of the outcomes of this analysis was that students vary their goals (e.g. have fun, learn maths etc) during the course of the interaction, depending on how things turn out. The analysis also lead to the introduction of a new goal, "want help" that better explained some of the user data than the existing set of goals. The Bayesian Dynamic Network therefore models both the students' emotional states as well as the motivational and affective pedagogic theory of what actions the system should choose in order to optimize a learner's state.

D'Mello et al. (2008) are exploring how a system can be designed to detect the learner's affective state through the detection of conversational cues, facial features, and posture. The authors select the following affective states as those most relevant to learning, based on observations of learners using the system: boredom, engagement, confusion, and frustration. The system focuses on the learner's affective state, not their motivational state, but there appears to be an implicit definition of motivation in terms of engagement (without a clear separation of affect and motivation). For example, they use the term motivation quite loosely, as illustrated in the following quote in which they explain the goals of the system: "the other essential component is to build mechanisms that empower AutoTutor to intelligently respond to these emotions, as well as to their state of cognition, motivation, social

³ The game involves "climbing a mountain" of numbered positions organized around the notion of factorisation. The game can be played with a partner, in which case the pedagogical agent helps the player about to make a move, or played in practice mode just against the pedagogical agent. Making a good move involves climbing higher. Making a bad move involves falling back. So one could imagine that when the student makes a good move, following help from the pedagogical agent, her degrees of joy, self-pride and admiration of the pedagogical agent would all increase.

sensitivity, and so on. In essence, how can an affect-sensitive AutoTutor respond to the learner in a fashion that optimizes learning and engagement? Therefore, the next phase of our research focused on fortifying AutoTutor with the necessary pedagogical and motivational strategies to address the cognitive and the affective states of the learner" (p.37).

A further issue is what information is used in order to make a diagnosis about the learner's motivational state. The above approaches have included interactional data and physiological and behavioural indicators of the learner's affective state. The approach by Arroyo and Woolf (2005) raises questions about the reliability of the diagnosis of the learner's attitudes based on self-report at a single point in time. Many of the components of motivation are not stable, but are likely to vary within a session (Keller, 1983; Malone & Lepper, 1987). The question is then, to what extent are students' self-reported motivations at the end of a session representative of their motivations during the session? Therefore, in what sense can it be related back to their interactions with the system? AutoTutor has been designed to take into account physiological and behavioural data to diagnose the learner's affective state. However, as discussed previously, these need to be coupled with data about the learner's motivational state. Finally, approaches based on OCC theory assume that a learner's affective state can be diagnosed based on whether they have achieved their goals or not. As discussed above, this may be appropriate in the narrow context of a game, but the definition of a learner's goals and whether or not they have been achieved is more complex in other learning contexts (see e.g., Pekrun, et al., 2002). Learner self-report was used as part of the diagnostic toolkit by del Soldato (del Soldato, 1994) and more recently has been used as a way of calibrating other methods of gathering information about learner state (see e.g., Robison, McQuiggan, & Lester, 2009). The main worry is not so much the reliability and accuracy of self-report (Balaam, Harris, & Fitzpatrick, 2009) as the intrusion of making self-reports during learning

Motivational pedagogy

The motivational, metacognitive and affective theories briefly outlined earlier provide some background to understand the design of systems, though the literatures on motivation, metacognition and affect are large and with many conflicting views. We should not lose sight of the fact that most systems are designed, at base, to improve the cognitive state of the student⁴. They employ motivational, metacognitive and affective reasoning as a means to this end. So their pedagogy focuses on diagnosing the student's state, and if that state is sub-optimal with respect to learning, helping the student move into a state more conducive to learning. Once the student is in a good state for learning the pedagogy aims to maintain that state.

The underlying premise is that the student's Cognitive, Motivational, Metacognitive and Affective (CMMA) state can be modelled, albeit with imprecision and uncertainty. From this standpoint *operational* motivational and affective pedagogy comprises three kinds of reasoning.

Let us assume that the student is in cognitive, motivational, metacognitive and affective state $CMMA_{Current}$, then the three kinds of reasoning are as follows:

⁴ Of course, some systems are designed to improve the student's metacognitive state *per se* (see e.g., Aleven, McLaren, Roll, & Koedinger, 2006; Rose Luckin & Hammerton, 2002).

- (i) Reasoning about causes: How did the student get into the state CMMA_{current}? Some states, for example, frustration, boredom and anxiety have a number of possible causes. It is not enough simply to ascertain that the student is bored (say) for the system to make right pedagogical move. It is important to trace the cause of that state. The student might be bored because the work is too easy. But they may be bored because the work is too hard and is "going over their heads". What needs to be done in these two cases is very different. The issue of causation is developed more fully in du Boulay (in press)
- (ii) Reasoning about consequences: How will the student's state CMMA_{Current} be affected by an event such as completing a problem successfully, or being praised for her effort, or taking time to think about her goals. In other words what will CMMA_{Nexr} be like? For example, in many tutors completing a problem successfully produces positive increments in the model of the student's ability to solve similar problems in the future, and also positive increments in whatever variables are used to represent the student's feeling of well-being. This kind of reasoning can be used by the system to predict the consequences of events including those events not provoked by the system itself.
- (iii) Reasoning about means: the inverse of the above is reasoning by the system as to what action it should take that would be most likely to help move the student from her current state CMMA_{Current} to a desired state CMMA_{Desired}: for example, from a student state of relative ignorance about how she might feel about a future learning outcome to one of considered anticipation, or from a state where the student tends to overuse the help facility in solving problems to one where her use is more careful. This kind of reasoning can be used to help the system select from its repertoire of actions what might be best to do next: offer advice, ask the student to explain she has just done, ask her to consider how much help she is using and so on.

Reasoning about consequences

Exactly how cognition, motivation, affect and learning interact is still largely understood in the qualitative terms of psychological and educational theory rather than in the detail needed for system design (see e.g., Herrington & Herrington, 2005; Wentzel, 2002), though to some extent Conati and her colleagues' methodology (see earlier) avoids the need to grapple with this tricky issue.

Most models centre around a "node and link model" of CMMA, with varying complexity in terms of the number of nodes and the kinds of interactions between them. Node and link models themselves vary along a continuum of internal complexity as well as along a continuum of motivational and affective richness. At one end of the continuum there are complex models of emotional processing in general (cognitive) terms that provide a way to understand how emotions emerge, develop and change: OCC being a well-known instance (Ortony, et al., 1988). At the other end of the continuum are much simpler models of motivation only, based on the interactions between a small number of variables, as in MORE (del Soldato & du Boulay, 1995). Finally there are models somewhere between the two that are complex but cover only those emotions that are relevant in educational situations. Conati and her colleagues' work falls into this latter class (Conati & Maclaren, 2005; Conati & Zhou, 2002).

Reasoning about means

The theory of the relation between motivational, metacognitive and affective states and learning is relatively undeveloped at the kind of detailed level required by intelligent educational systems. For recent progress in the area of affect see Blanchard and his colleagues for a useful short guide to "affect management" in AIED systems, including distinctive strategies for managing emotions, moods, attitude and interpersonal stance (E. G. Blanchard, Volfson, Hong, & Lajoie, 2009); also see Lehman and his colleagues for detailed analyses of how human tutors deal with student Emotions (Lehman, Matthews, D'Mello, & Person, 2008). In the general most approaches centre around getting the student into a good state for learning and then keeping her there.

One aspect of this approach was explored nearly a century ago and encapsulated in the Yerkes-Dobson law which suggests that the most productive state of learning is neither when the student is emotionally under-aroused or over-aroused, but in some intermediate state between the two (For an application of this law in an educational context and for a more general discussion, see Bregman & McAllister, 1982; Picard, 2000).

In terms of motivational pedagogy, the rules for the behaviour of systems, the work of Keller and his ARCS model provides some qualitative guidance. The model is a qualitative theory of motivational pedagogy based around the key factors of Attention, Relevance, Confidence, Satisfaction, curiosity, and independence (Keller, 1983). It sets out a qualitative theory of what the teacher needs to do to ensure that students learn effectively. This was expressed in terms of making sure that the students were attending to what was to be learned, that it was relevant to their needs, that they remained confident in their ability to tackle the material, that they were satisfied by their interaction, that their curiosity was stimulated and that they enjoyed a sense of independence.

Various researchers have built motivational models, including Keller himself in his Genetics program based on ARCS (Song & Keller, 2001). For example, the authors of this paper have developed three systems that are based on ARCS (del Soldato & du Boulay, 1995; Martinez-Miron, Harris, du Boulay, Luckin, & Yuill, 2005; Rebolledo-Mendez, du Boulay, & Luckin, 2006). Each of these systems has modelled motivation by choosing numeric variables to represent some of the ARCS factors, typically including *confidence* and *independence*. Both student actions and system reactions are presumed to have effects on these variables, either incrementing or decrementing their values. When a value reaches a threshold it will either trigger a system reaction or will cause the system to react in a different way to an event, such as solving a problem. We describe these as "thermostat" models in that they effectively attempt to maintain the variables representing the student's motivational state between upper and lower bounds of normality – so acting like a motivational thermostat.

Another approach to reasoning about means divides motivational and affective states between the positive (e.g. delight) and the negative (e.g. confusion) and systems are developed and evaluated in terms of maximizing the frequencies of the student transitioning into a positive state (including from a positive state) and minimizing frequencies of transitioning into a negative state (see e.g., Baker, Rodrigo, & Xolocotzin, 2007).

Given the complexity of operationalising motivational and educational theory, new useful approaches are now being adopted by various researchers. One such involves trying to relate motivational states to learning empirically. This can be done "live", as it were, by observing students undertaking some learning task, categorising their demeanours into such states as "engagement" or "frustration" and then trying to either link this to learning outcomes or back to particular events or interaction in the learning (Baker, et al., 2007). In a similar vein D'Mello and his colleagues are

exploring the interactions between problem-solving states and affective states and have found certain cyclic phenomena such as the virtuous cycle of curiosity, expected positive feedback, happiness and further curiosity (S. D'Mello, et al., 2009).

A further approach is to mine the large amounts of user interaction data now available. These methods are used to find relationships in that data to measurable variables such as post-test performance e.g. via Hidden Markov Models (Soller & Lesgold, 2003); or via other machine learning techniques (Mavrikis, Maciocia, & Lee, 2007). For example, Arroyo & Woolf (2005) use Dynamic Bayesian Networks to infer (probabilistic) relationships between externally measurable values such as the number of times the student sought help with hidden, internal variables such as the student's attitude to challenge. By correlating these values with self-report post-tests of student attitudes and performance (such as "I just wanted to get the session over with, so I went as fast as possible without paying much attention"), the researchers are moving towards a theory of what kinds of actions and reactions of the system are likely to best induce good attitudes as well as good learning in future students.

A PARTIAL HIERARCHY OF INTELLIGENT EDUCATIONAL SYSTEMS

In order to pull some of the preceding theory and implementation together we present a conceptual framework. This is in two parts. First we identify a partial hierarchy of increasingly intelligent educational systems in terms of where their educational priorities and focus lie, and the kinds of resources and reasoning that they are able to deploy (see Table I). In the subsequent section the conceptual framework looks at the design of systems from the point of view of the kinds of diagnostic data available to them (to determine the student CMMA state) and the repertoire of actions available to it to help change that state (see Table II).

The partial hierarchy in Table I is organised as follows. At the base are cognitively and metacognitively intelligent systems together with affectively intelligent and meta-affectively intelligent systems. Cognitively intelligent system have a focus on increasing the learner's knowledge and skill, whereas metacognitively intelligent systems have a focus on helping the learner to better understand and manage her own learning.

At the same base level, affectively intelligent systems have a focus on increasing the student's feelings of well-being as a learner, whereas meta-affectively intelligent systems have a focus on helping to increase the learner's insight into her own feelings as a learner and also her ability to manage and regulate those feelings. While affect and cognition are themselves strongly interlinked, as we have already argued, they are sufficiently distinct to merit different names and both feed "upwards" into motivation. At the next level up we have motivationally intelligent and meta-motivationally intelligent systems. The former focus on improving the learner's motivation and her willingness to expend effort in learning, while the latter focus on helping the learning gain insight into her own motivation and into ways of managing that are effective for her.

At the apex of the partial hierarchy are Caring Systems that have a regard for the growth of the learner as a person.

	Kind of system	Pedagogic Focus	
	Caring systems	The growth of the learner as a person	
	Meta-motivationally intelligent Motivationally intelligent	Increasing the learner's meta-motivational capability e.g. her insight and regulation of her motivation Increasing the learner's desire to learn, e.g. her willingness to expend effort on the learning process.	
Kind of system	D. J E		
King of system	Pedagogic Focus	Kind of system	Pedagogic Focus
Metacognitively intelligent	Increasing the learner's metacognitive capability, e.g. insight into what she understands and can do, and her ability to regulate her learning process effectively.	Meta-affectively intelligent	Pedagogic Focus Increasing the learner's meta- affective capability, e.g. her insight and regulation of her feelings as a learner.
Metacognitively intelligent Cognitively	Increasing the learner's metacognitive capability, e.g. insight into what she understands and can do, and her ability to regulate her learning process effectively.	Meta-affectively intelligent Affectively intelligent	Pedagogic Focus Increasing the learner's meta- affective capability, e.g. her insight and regulation of her feelings as a learner.

Table I: Partial hierarchy of intelligent educational systems

Cognitively intelligent educational systems

At the base of the partial hierarchy on the left lie traditional, cognitively intelligent educational systems whose priority is to improve the knowledge and skill of the learner and which do not take metacognitive, motivational or affective issues dynamically into account. A typical example of this class of system would be the PAT Algebra Tutor (Koedinger, Anderson, Hadley, & Mark, 1997). This

is not to say that this system is not motivating to use, just that the system itself cannot reason about motivational issues.

We propose the following definition of a traditional intelligent educational system:

A **cognitively intelligent educational system** is a system that is able to deploy its resources and tactics dynamically and effectively to support learning but without modelling the metacognitive, affective or motivational states of the learner.

This definition is agnostic as to whether the system is teacher-centred or student-centred, or oneto-one or supporting collaborative work. The main issue is that it has explicit representations of the domain, of the learner and of its pedagogy and uses these to reason about how to make dynamic decisions at run-time so as to maximize the learning of the learner(s). Of course the designers of such systems will have designed-in such motivating elements as they can, but the basic assumption is that the system will be used by motivated learners, who are in a happy frame of mind, and able to manage their own learning so that if a particular learner is, or becomes, demotivated (say) the system has no means of noticing this *in these terms* or of reasoning about how to deal with it.

Metacognitively intelligent educational systems

An extension to traditional systems includes those which also reason about the learner's metacognitive capability, either simply in support of developing the learner's knowledge or skill, or more ambitiously, with the focus of developing their metacognitive capabilities *per se*.

A **metacognitively intelligent educational system** is a system that is able to deploy its resources and tactics dynamically and effectively to *improve the metacognitive capability* of the learner.

Examples of such system include Gama's system to develop student's ability to reason about mathematical problem solving (Gama, 2004), Conati and Van Lehn's system to develop the student's ability to self-explain (Conati & Van Lehn, 2000). There are also those tutors aiming to develop students' use of on-line help (Aleven, et al., 2006) and propensity to embrace learning challenges (Rose Luckin & Hammerton, 2002). Betty's Brain is a system designed to confront the learner with the adequacy or otherwise of what she understands in a particular domain by inviting her to teach "Betty" and then have Betty take a test on that material. If the test indicates that Betty has not been able to answer a question correctly, the learner has the opportunity to refine what she taught Betty. The system is based around the metaphor of learning by teaching and a recent version contains specific feedback to the learner about self-regulated learning strategies such as suggesting that Betty's understanding of the domain would need to be updated if she is to score better on the test (Leelawong & Biswas, 2008). A similar concern for the motivational and reflective learning aspects of teaching a (non-embodied) pedagogical agent was explored by Uresti & du Boulay (2004). Their agent (unlike Betty) rather realistically occasionally forgot what it had been taught. This feature was disliked by the students – another plausibility problem.

Affectively intelligent educational systems

At the base of the hierarchy on the right are affectively intelligent educational systems. These are a class of systems which are able to reason about the learner's affective state.

An **affectively intelligent** educational system is a system that is able to deploy resources and tactics dynamically to provide an educational experience that *improves the student's state of well-being* during learning.

Note that the educational focus of such a system may still be largely on the performance and skill of the learner. It may take affective issues into account but only in support of performance goals. Typical examples here would be AutoTutor (S D'Mello, et al., 2008) and Wayang Outpost (Arroyo, Cooper, et al., 2009).

Meta-affectively intelligent educational systems

Also at the base of the hierarchy on the right are meta-affectively intelligent educational systems. These are a class of systems whose focus is on helping the learner to understand and manage her own affective states more effectively.

A **meta-affectively intelligent** educational system is a system that is able to deploy resources and tactics dynamically to provide an educational experience that *improves the student's understanding of, and her ability to manage, the affective dimension of her learning.*

As far as we are aware no-one has attempted to build a meta-affectively intelligent system but Yussof has taken some steps in this direction (Yussof & du Boulay, 2009) and Zhang et al. have built a system in which learners are prompted to reflect on the emotional inner life of the characters in an animated drama (Zhang, Gillies, Dhaliwal, & Gower, 2009).

Motivationally intelligent educational systems

In the second tier up of the hierarchy are systems that reason about the motivational state of the learner. They are placed at this level, as understanding and reasoning about motivation involves understanding and reasoning about both cognition and affect in the first instance as well as metacognition and meta-affect in the second instance. We define a *motivationally intelligent educational system* as follows:

A **motivationally intelligent** educational system is an intelligent system that is able to deploy resources and tactics dynamically to *maintain or increase* the student's *desire to learn* and her *willingness to expend effort in so doing*.

The special goal of such systems is to maintain or even increase the learner's *desire to learn* and *her willingness to expend effort* in undertaking the, sometimes hard, activities that lead to learning. Such a goal is in addition to, but intertwined with, the more traditional educational system goals of offering information, activities and support for learning new knowledge and skills (see e.g., Song &

Keller, 2001). A reason for placing this kind of system higher in the hierarchy than either metacognitively or affectively intelligent systems is that the theories of motivation taken as a whole (see earlier) have both a metacognitive and an affective dimension. The motivationally expert (human) teacher exploits what she understands of both the learner's individual affective disposition and her particular metacognitive insight to find and operate the keys to motivate her. To that extent such a teacher taps into the deeper values of the learner that underpin her willingness to engage *wholeheartedly* in learning at that particular moment.

Meta-motivationally intelligent systems

A meta-motivationally intelligent system aims to improve the student's ability to understand and manage her own motivational processes. Thus we define a meta-motivationally intelligent system as follows:

A **meta-motivationally intelligent** educational system is an intelligent system that is able to deploy resources and tactics dynamically to *increase* the student's *ability to understand and regulate her own motivation*.

As far as we know there is no intelligent system that attempts meta-motivational teaching but the goal is similar to that undertaken by human sports coaches who seek help their coachees understand and regulate their ability to drive themselves forward through training and, in competitive sports, deal with the motivational aspects of starting to lose a game (say), or indeed starting to win.

Caring Intelligent Systems

Finally at the top of the hierarchy we have, to use the term from Self (1999), Caring Intelligent Systems. These are meta-motivationally intelligent systems that also can reason about overall context and environment within which the student is learning. The main difference between a caring system and a meta-motivationally intelligent system is in terms of the degree of regard that the system has for the "whole learner" as a person operating within an educational, physical and social context and also as someone sensitive to, and able to, regulate cognitive and affective aspects of their learning.

Artificial Intelligence is still some way from producing Caring Systems. To gain some insight into what this would mean in practice, see the work of Rosiek (2003). She describes how insightful human teachers can subtly reconceptualise material to be taught in order to make it more accessible to the learner, or less negatively "loaded" emotionally, thus maximising the chance that they will engage positively and minimising the chance that they dismiss the issues out of hand.

A special issue of the International Journal of Artificial Intelligence in Education was devoted to honouring John Self and his notion of systems that care. Various authors delineated aspects of caring, for example, Cooper (2003) noted that care in an educational setting involves "profound empathy in one to one empathic relationships" with concern for both the personal and academic development of the learner. Kay & McCalla (2003) developed the issue of care in terms of "computational mathetics" (a basic science of learning). Bull, Greer & McCalla (2003) described the I-Help system that involves the learner having a personal agent who helps them, sometimes by seeking out someone else who might provide assistance. Dimitrova (2003) saw care in terms of ensuring that the system has as good a model of the learner as possible and argues for interactive open learner modelling. Katz and her

colleagues (2003) saw an aspect of care in careful reflective discussions with the learner after problem-solving. Finally Tedesco (2003) saw care in terms of the role played by a system in managing conflict in group interactions.

DIAGNOSTIC INPUT TO AND FEEDBACK REACTIONS FROM SYSTEMS

We now turn to the second part of the conceptual framework, namely identifying issues around diagnosis and reaction. So we consider the nature of the data that the different kinds of intelligent system make use of and the kinds of adaptive pedagogic tactics and strategies that they are able to deploy. There is clearly a relationship between the position of the system in the partial hierarchy explained earlier (see Table I) and the data that the system uses and the pedagogic moves it can deploy: broadly systems nearer the top of the hierarchy are able to deploy a more expert set of moves and usually need a broader range of data in order to determine those moves. However the relationship needs to be viewed with care, as we argue below.

We define four broad categories⁵ of diagnostic input and feedback reaction, see Table II. These categories mirror, to some extent, the hierarchy of systems developed in Table I. By "inputs" here we mean the kind of event or measurement that provides input data to the system, such as the student asking for help, completing a problem, dominating a discussion with a peer, or changing their posture. By "reactions" we mean actions, reactions or outputs by the system, such as setting a harder problem, putting two students in touch with each other, changing the facial expression of an online pedagogical agent, or providing a deeper level of help and so on.

These four categories largely mirror the earlier part of the paper and are: (i) the cognitive and metacognitive; (ii) the affective and the meta-affective; and (iii) the motivational and meta-motivational. To these three we add a fourth category, namely the context and meta-context: in what kind of location and milieu is the learning taking place? What is the physical ambience? And so on. At the meta-level there are issues around the degree that the learner can articulate and possibly control the effects of these factors on her learning e.g. does she need to listen to music as a background to studying.

Student actions and their own or the system's reactions do not necessarily operate within the same category. A system may take diagnostic input in one category but react in others. For example, imagine that a system is able to detect *physical* symptoms of nervousness on the part of a student, e.g. via sensors connected to the student's hands or via wriggle detectors in her chair. On the basis of this input data the system might decide that one way to assist that particular student might be to react in terms of the educational *context* by changing the nature of the interaction from one-to-one, to many-to-one by inviting some of the student's peers to also take part in the same activity.

⁵ Earlier papers by the authors have described these categories in different ways.

CATEGORY		DIAGNOSTIC INPUTS	FEEDBACK REACTIONS
CONTEXT	The spatial, physical, social and temporal milieu within which the student is learning.	Location e.g. classroom, home, library. Peer group. Physical ambience.	Change of location, use of available peers and others, Change of ambience e.g. light levels.
META-CONTEXT	What the learner knows and can articulate and regulate about the context in which she is learning.	Comments from the learner about the context.	Conversations about the nature of the context and how different contexts affect learning.
MOTIVATION	What drives the learner to learn or not to learn, what they think they are going to achieve, why they are learning at all and the social and temporal milieu within which the learner is learning.	Goals, expectations, values, needs, rationale as well as contextual features such as educational milieu.	Flagging of goals and outcomes, intrinsic and extrinsic rewards. Use of available peers and others, change of educational milieu.
META- MOTIVATION	What the learner knows and can articulate and regulate about her motivation.	Comments from the learner about the motivation.	Conversations about the nature of the development of motivation.
AFFECTIVE	How the learner feels about the learning activity.	Demeanour of the learner e.g. happy, engaged.	Praise, encouragement, criticism, politeness, teacher's demeanour.
META- AFFECTIVE	What the learner knows, can articulate and regulate about her actual and expected feelings.	Comments from learner about expectations of feelings, motivation.	Conversations about expectations of feelings, state of motivation, engagement.
PHYSIOLOGICAL	Bodily aspects such as heart and breathing rate, skin conductance, facial expression, body language and posture.	Sensors: skin, body movements, Cameras: facial expression, posture.	Breathing exercises, mantras, pauses. Changes in the physical context e.g. warmth, light, ambient noise.
META- PHYSIOLOGICAL	What the learner knows and can articulate and regulate about her physiological responses.	Comments from learner about her body.	Conversations about physiological response.
COGNITIVE	Knowledge and skills of the learner.	Performance, latencies, effort, focus of attention.	Activity choice, pace or order of work, provision of help.
META- COGNITIVE	What the learner knows, can articulate and regulate about her knowledge and skills.	Difficulty of work chosen, use of available help (including gaming), goal orientation.	Conversation about performance, degree of challenge, use of help, narrative framework.

Table II: Categories of input and reaction

The paper now takes each block of Table II in turn and describes one or more systems that have operated primarily at that level, allowing for the fact that many systems have operated at more than a single level, and sometimes have operated *at one level for diagnostic input and at another for feedback reaction*.

Cognitive and metacognitive

Systems operating exclusively at the cognitive level for diagnostic input and feedback reaction are what we have earlier described as traditional intelligent educational systems and we need say no more here.

Diagnostic input

Various systems have been built to deal with different metacognitive issues. These have made use of data both at the cognitive level and at the metacognitive level. Ecolab II, for example, was designed to teach to 10-11 year olds, at the domain level, the concepts of food webs and chains, and at the metacognitive level, effective choice of problem and effective use of help (Rose Luckin & Hammerton, 2002). The system offers a virtual laboratory which can be viewed from different perspectives, and into which children can introduce various creatures and populations of creatures to see what eats what and how the sizes of population vary over time. Diagnostic input included the accuracy of the student's answers to problems, the difficulty of the problems chosen and the complexity or amount of help sought from the system. Aleven et al. have built a tutor based on both a model of desired help-seeking behaviour as well as rule-based analysis of inappropriate uses of help (such as "clicking through hints"). Their system is built on top of, and uses the same technology as, earlier Cognitive Tutors. The input to the system is the learner's use and misuse of the help system including timings (Aleven, et al., 2006).

Feedback reaction

In terms of reactive feedback most systems attempt to guide the learner to behave in ways similar to a metacognitively more aware learner, but do not engage the learner in a dialogue about their metacognitive insight as such. This means that the metacognitive advice is focused on the specific application domain without explicitly attempting to help the student reflect and generalise this to other domains. So the Ecolab II system commented to the student about her strategy, perhaps suggesting that she tackle more demanding problems, or perhaps try to solve the problems using less help. The help-seeking tutor (Aleven, et al., 2006) tries to guide the learner to behave more like an ideal student, e.g. by making use of hints or an online glossary as initial tactics when she believes that she needs help.

Affective and meta-affective

Diagnostic input

The affective level has received a lot of attention both in educational systems and in intelligent interaction in general (Paiva, et al., 2007). For example, on the diagnostic input side, D'Mello et al. (2006) found ways to predict the learners' "affective states (e.g. confusion, eureka, frustration)" by looking at the conversational patterns between the learner and AutoTutor. McQuiggan and his colleagues (2007) have developed a method for detecting the onset of frustration based on an inductive approach working from the learners "actions, locations [within the learning environment, Crystal Island], goals and temporal information" (p. 699). Arroyo & Woolf (2005) detected the learner's hidden affective state from external task performance data using a Bayesian Network. Chaffar & Frasson (2004) determined the learner's optimum emotional state for learning using the learner's choice of a sequence of colours that had been calibrated earlier via a decision tree. De Vincente & Pain (2002) helped teachers articulate rules that infer learners' affective state from interaction data between students and an online mathematics teaching system. Hernandez & Noguez (2005) predicted the learner's affective state in terms of the OCC model, based on "personality traits, student knowledge state, mood, goals, and tutorial situation (i.e. outcome of the student's actions)". Zhang et al. (2005; 2009) have explored methods of detecting learners' affective state from their textual contributions to what characters should say in an e-drama, and Beal & Lee (2005) have used learner's self-reports to determine their affective state for the Wayang-West ITS. Finally, we note the work of Balaam et al. (2009). They have designed a classroom technology for use by learners called a Subtle Stone that has both communicative (to the teacher) and reflective (for the learner herself) capability. The technology enables the leaner to communicate her current affective state to the teacher privately without other learners in the classroom being aware.

Lehman and his colleagues have built on the work of Lepper and others to investigate the range of learner affective states as they occur working with expert human tutors in a classroom and found that the most common emotions were confusion, anxiousness and happiness. They also looked at the ways that expert tutors reacted to different student states e.g. "they were more likely to provide positive feedback when the students were confused than when frustrated", or that they were "more likely to utilize off-topic conversation when students were happy or anxious" (Lehman, et al., 2008, pp. 55-56).

Feedback Reaction

On the reactive side different researchers have opted to try to deal with different aspects of affect. Some have been concerned to try to map out the territory of student affect and of teachers' ability to recognise and codify affective states (see e.g., Alexander, 2005; de Vicente & Pain, 2002; Porayska-Pomsta, Mavrikis, & Pain, 2008). Despite the evidence that teachers can adapt at the macro level to some changes in the affective states of students or predict their affective states based on experience, there are serious questions to ask as to how far human teachers are able to identify and classify accurately the changing affective states of students at the micro level over the course of a lesson (see e.g. Balaam, et al., 2009; S. D'Mello, Taylor, Davidson, & Graesser, 2008). Whether this is to do with insufficient specific teacher training or to do with the inherent ambiguity of the available affective

clues or to do with the affective impact of factors outside the educational interaction is a moot point (S. D'Mello, et al., 2009).

Researchers have been concerned to explore potential strategies for managing affect (see e.g., E. G. Blanchard, et al., 2009; Rosiek, 2003). In terms of implemented systems, Baylor et al. have explored ways of dealing with learner *frustration* either by having an interface agent emit apologetic messages (after a system malfunction) or by having it emit empathetic messages such as, "It must have been very frustrating to try to finish the survey with the problem you were experiencing. I sympathize with how you feel. I wish that I could have helped you overcome this problem. Please take a few minutes to describe your experiences from the previous screens. Thank you" (Baylor, Warren, Park, Shen, & Perez, 2005, p. 75). The results suggested that *plausibility* (du Boulay & Luckin, 2001) was an issue, more so for apologetic messages than for the empathetic. Forbes-Riley & Litman (2009) have developed a tutor that takes account of the learner's self-expressed degree of *uncertainty* in answering physics problems. Using a Wizard of Oz methodology, they found that adapting feedback specifically to learner uncertainty via empathetic responses improves both learning efficiency and learner satisfaction.

At the meta-affective layer there has not yet been much progress in the sense of the system engaging the learner directly in discussion about affective issues. So for example in the systems described above empathy is used as part of the feedback to self-reported, observed or inferred learner states, but not in *anticipation* of such states. Indeed the systems that have been implemented tend to concentrate on the meta-affective aspects of scenarios other than the current learning situation. An example is the work of Marsella et al. (2003) who have developed a system to conduct an interactive pedagogical drama for health interventions with a view to assisting the learner gain insights into her own affective reactions to situations portrayed in the drama. In a similar vein, but working more indirectly on meta-affective issues is the work of Zhang and colleagues (Zhang, et al., 2005; Zhang, et al., 2009).

Yussof has developed a tutor for programming that invites the learner to pause and undertake relaxation exercises between problems. This is aimed to improve their sense of well-being, especially if they have not done so well on a problem (Yussof & du Boulay, 2009).

Pedagogic agents

One of the areas in which motivational and affective modelling comes into sharp, contemporary focus involves the use of animated pedagogical agents. Systems employ such agents in a wide variety of roles: as embodiments of one or more teachers, one or more peers, as team-mates, as pets to be managed and so on. In terms of their reactions, such agents can operate at all the levels identified. So they can offer advice and guidance at the cognitive level or at the metacognitive level, can offer affective feedback by smiling or scowling, and indeed adjust the problem context by offering teammates to help someone learn a complex, collaborative, procedural task (Johnson, Rickel, & Lester, 2000). A particular recent focus has been on agents which provide affective feedback in the form of encouragement, empathy, smiley faces and so on. Much of the research has concentrated on mapping out the consequences of employing different kinds of pedagogical agents for different kinds of learner. So for example, Arroyo and her colleagues have compared the effects of agents of different gender which presented themselves in varying affective states in order to demonstrate empathy with the learner's affective state, e.g. confident, excited, bored, focused, frustrated, anxious (Arroyo, Woolf, Royer, & Tai, 2009). In a similar vein Baylor & Plant (2005) compared how agents which varied in

terms of gender, age, attractiveness and coolness affected female students' views about engineering. Haake & Gulz have compared the learner preferences along a number of dimensions including their "visual static appearance, pedagogic role, and communicative style" (Haake & Gulz, 2009). Kim (2007) explored interactions between learner characteristics and agent characteristics e.g. competency and control, and found differences in preference between academically strong and academically weak learners. For example strong learners liked strong agents and valued their high degree of control.

Baylor and Kim (2005) compared agents playing three different roles: an expert, a motivator and a mentor. The motivator was designed to encourage by using an effusive and enthusiastic tone of voice and emotional animation thereby operating on the affective aspects of acknowledgement, confusion, disapproval, excitement, pleasure and surprise. The mentor was similar to the motivator though it also provided information and adopted a confident and calm voice as opposed to an effusive and enthusiastic one. The expert provided information in an authoritative manner and did not explicitly attend to the affective dimension (though of course, a teacher failing to attend to the affective dimension may well have affective consequences for some learners). They found that "the Expert agent led to increased information acquisition, the Motivator led to increased self-efficacy, and the Mentor led to overall improved learning and motivation" (Baylor & Kim, 2005, p. 95).

Clearly animated pedagogical agents now bring to the fore a number of interesting "interpersonal" and social issues that have a bearing on motivation. Should the agent face the learner? How much of the agent should be visible? What facial expression should be adopted? What body language and posture should be shown? What degree of politeness offered? And so on. Many of these issues raise design questions that were not so problematic when educational systems could offer only typed messages on the screen back to learners. Clearly the issue of plausibility is central too.

So for example, we have Johnson & Rizzo (2004), Porayska-Pomsta and colleagues (2008; 2004) working on the relation between the *politeness* and social sensitivity of the tutor and the expected motivational state of the learner including maintaining the learner's sense of "face"; there are Chen and his colleagues (Chen, Chou, Deng, & Chan, 2007; 2005) working on the effects of pedagogic pets of different degrees of cuteness and their effects on learners; and Gulz & Haake (2005) and Kim (2005) investigated how pedagogical agents should represent themselves to users.

Physiological and meta-physiological

There is an obvious overlap between dealing with diagnostic input at the affective level and with its physiological manifestations, such as sweaty hands or engaged posture.

Diagnostic Input

We already mentioned the work of D'Mello et al. (2008) and Kapoor et al. (2007) in the Introduction. Heraz & Frasson (2009) are exploring the use of *brainwaves* as correlates to self-reports of emotional states (pleasant to unpleasant, degree of arousal, and dominance/control) and were also able to determine whether a learner was guessing or answering randomly in a test based on brainwave activity. D'Mello and colleagues make use of *conversational cues* (dialogue features such as response verbosity), *gross body language* (posture features via pressure sensitive pads in seat and back of the learner's chair) and *facial expressions* monitored via a camera. Kapoor and colleagues have adopted a similar multi-modal approach. Also of note are Kleinsmith et al. (2005) who explored recognising emotional state from posture, Prendinger et al. (2005) who recognised emotional state from *body sensors* and Mozziconacci (2001) who explored how to interpret the facial expressions and *voice*.

This technology is now moving out of the laboratory and into the classroom: Arroyo and her colleagues have collected sensor data in a classroom environment and correlated it with learner self-reports of their confidence, frustration, excitement or interest. The sensor data was collected via a camera, the learner's chair, a wrist sensor and a pressure-sensitive mouse. They found that they could predict learner's self-reported emotional state from sensor data with a good level of reliability (Arroyo, Cooper, et al., 2009).

Feedback Reactions

There is very little work on reacting automatically to physiological data. For instance, one might imagine a system that changed the air temperature or angle of the learner's chair in response to a judgement about the learner's state of engagement or effort. Moreover, no system that we know of is concerned with what the learner knows and can articulate and regulate about her physiological responses, or can accept comments from the learner about her physiological state, or can conduct conversations with the learner about her physiological response to past, present or anticipated learning experiences.

Motivation and meta-motivation

Diagnostic Input

MORE was one of the first motivationally intelligent systems to be developed (del Soldato, 1994; del Soldato & du Boulay, 1995). The domain was Prolog debugging and the system was designed for one-to-one undergraduate use. It was based on an ARCS model of motivational pedagogy and monitored the amount of effort, persistence, use of help, self-reports and quality of answers to determine the learner's state of motivation expressed in terms of her confidence, independence and effort. For example, MORE sought information from the learner about how she felt about tackling a new problem, after just an initial glimpse. However the way that input was gathered blurred the distinction as to whether the learner was making a meta-affective or metacognitive comment when either accepting to solve the new problem or rejecting it.

McQuiggan and Lester investigated ways of automatically determining a learner's *self-efficacy*. They looked for correlations between learners' answers to a standardised questionnaire on self-efficacy with a range of static data as well as with a variety of dynamic measures. They found an interesting relation between heart rate and skin galvanic response and self-efficacy for a group of adult learners where the pattern of change of heart rate during problem-solving of those with high self-efficacy differed markedly from those with low self-efficacy (McQuiggan & Lester, 2006).

Feedback Reactions

MORE reacted both at the domain level e.g. by adjusting the difficulty of the next problem or by changing the degree of help that would be offered and at the motivational level by praising either performance or effort. In a similar fashion, Blanchard & Frasson (2004) chose to focus their design aspirations on dealing with the learner's sense of autonomy, though they exemplify this largely in terms maintaining the learner's sense of engagement with the activity by adjusting the nature of the activity if need be. Rebolledo-Mendez's extension of Ecolab II (Rose Luckin & Hammerton, 2002)

also introduced new elements into the educational interaction, such as a quiz and a treasure hunt, when motivation was detected to be flagging. This was coupled with reactions at the affective level by adjusting the facial expression and tone of voice of two pedagogical agents (Rebolledo-Mendez, et al., 2006). A similar approach was used by Kelly and Weibelzal to maintain the learner's interest and particularly to "making the learner confident that effort and performance are closely coupled with consequences" (Kelly & Weibelzhal, 2006, p. 537).

Games

There is increasing interest in developing intelligent systems using either an educational interaction based on games or exploiting the graphic and interaction technology that underpins games. It is a truism that many students find computer games compelling and enjoyable, have "flow" experiences while playing and exhibit high degrees of motivation to continue. For example, Bader-Natal and Pollack developed a system that supports pairs of learners (Grades 3-7) playing a competitive spelling game. They were concerned to increase the degree of *challenge* in the interactions by differentially rewarding well-chosen words for the other person to spell. However the most sophisticated contemporary example of this approach are the language tutors developed by Johnson (see e.g., Johnson, 2007). The domain of language learning fits well with the ethos of computer games, but in other domains (such as writing skills) there may be tension between the motivational effect of the game and its cognitive and particularly metacognitive effects unless special care is taken (Howland, du Boulay, & Good, 2009).

Context and meta-context

We have added context and meta-context as two areas that potentially provide both input data (such as the nature of the physical environment or the social context within which the learning is taking place) and also possibilities for feedback reaction (getting two individually bored students to work as a pair). While many systems are carefully designed, tuned and integrated into the actual classroom of use (see e.g., Koedinger, et al., 1997) no system that we know of is able to make such adjustments automatically nor is concerned explicitly with what the learner knows and can articulate and regulate about the context for her learning, can accept comments from the learner about the context, or can conduct conversations with the learner about the context, see Luckin (2010).

CONCLUSIONS

This paper has provided a conceptual framework that describes the space between traditional intelligent systems and systems that care. A partial hierarchy of systems has been identified distinguishing metacognition, motivation and affect. The paper has described individual systems in terms of the way such systems (i) derive input data and (ii) the manner in which they react in order to maintain or develop the learner's motivational and/or affective state. This analysis suggests that there are overlaps and confusion particularly between affect and motivation; and that not all areas have been tackled as yet, particularly those for meta-affect and meta-motivation. The analysis also points to the lack of work on context, both in terms of the system understanding and reasoning about the context of use and then engaging the learner at the meta-level about this. Four areas of further work are

identified: theory development, the role of the meta level, tradeoffs between categories and plausibility.

While the psychological and pedagogic literature offers a rich set of theories in the areas of motivation and affect, the range of working operational models of such theories exploitable by motivationally intelligent systems is more limited. This has led to different teams of researchers operating with different motivational ontologies, looking at different aspects of educational interactions over different timescales and with different granularities, and categorising learners in terms of individual differences in a variety of ways. This has meant that it is hard to accumulate the findings from one piece of work with another so as to build the kind of comprehensive "pay-off" matrix that one might like to embody as a default motivational pedagogy. For an example of part of such a matrix derived from observation of human expert tutors, see (Lehman, et al., 2008). A more comprehensive matrix might contain rules something like as follows:

IF a learner who has personality characteristics A, B and C, and expectancy in this context of D, is feeling E & F just now and holds values G & H, and has just solved a problem without much help, but with a lot of effort, and \ldots

THEN the system should comment in terms of I, choose a next problem of type J, adjust the learning context to K, and update the student model as per L, M and N, and \dots

Of course such a set of rules is an ideal that will be hard to realise, given the huge variability of learners and contexts within which they might be assisted to learn (for a discussion of this issue see e.g., Zakharov, Mitrovic, & Johnston, 2008). An additional complicating factor is that not only are the clues offered by the learner about their motivational state often ambiguous, so are the reactions by the system. A comment from the system about the learner's effort (say) may be interpreted by the learner in number of ways, and depending on that interpretation have different effects.

Some researchers are partially circumventing the lack of operational theory (but adding to the fragmentation) by building different kinds of probabilistic model that link observables of various kinds, via motivational and affective models, to learning outcomes. While this is a sensible way forward for optimising particular systems, more work needs to be done to derive models that apply across a range of systems and in identifying the key variables that best determine how the system should react.

A particular factor that has been identified in each category is its associated meta-level. The importance of these meta-levels emerges from motivational theory. Few systems have attempted to interact with the learner in the meta-affective, meta-physiological category or in the meta-motivational category, i.e. discussing with the learner the kinds of feelings that they are likely to experience in future learning interactions or inviting self-reflection from learners about how past learning experiences actually felt or indeed what kinds of social context they expect to make most progress within. In just the same way that there are clear benefits in bringing meta-cognition and self-reflection to the fore, we argue that meta-affective and meta-motivational reflection can produce similar benefits in terms of increases in learning at the domain level as well as more mature attitudes to future learning. There may well be mileage in assisting students to reflect on each of the meta-levels themselves even if the educational system has limited capability to take input and react at that level (Avramides & du Boulay, 2009).

While work in the classroom is starting to provide data on what works motivationally and affectively and what does not, there remain many similar questions in the design of systems. For

example, what are the tradeoffs for diagnosing and reacting in the different categories? Imagine that you detect that a learner has disengaged somewhat and has started gaming the system. What should the system do? At the level of the domain, it could make adjustments to the educational activity; at the meta-cognitive level it could offer advice about effective learning; at the affective level, it could offer an affective diversion -a joke possibly; at the meta-affective level, it could try to find out how the learner is feeling; at the physiological level, it could suggest a screen break; at the motivational level, it could change the nature of the educational activity and make it a collaborative or cooperative one; and so on. What is for the best? The response from the system will be determined both by considerations of what caused the issue, of what might work best as well as by more pragmatic concerns. It is possible that a multi-category approach might be best. These kinds of trade-off are a relatively new area of research that will help build up the matrix of rules mentioned above. For example, Robison and her colleagues have captured learners' preferences for task-based vs affectbased feedback from a pedagogical agent and induced a model that explains a high percentage of the expressed preferences (Robison, et al., 2009). Observing non-expert human tutors working with computer science students (via terminals), Boyer et al. investigated the balance between cognitive and motivational scaffolding in tutorial dialogue. Among other results they found that "positive cognitive feedback may prove an appropriate strategy for responding to questionable student problem-solving action in task-oriented tutorial situations" and also that "direct standalone encouragement" helped students of low self-efficacy but not those with high self-efficacy (Bover, Phillips, Wallis, Vouk, & Lester, 2008, p. 247). Barrow et al. (2008) investigated the trade-off between positive and negative feedback against purely negative feedback in the SQL database query language tutor and found that learning was more efficient for the group who received both positive and negative feedback.

Finally there is the issue of plausibility (du Boulay & Luckin, 2001). As we have noted various researchers start from observations of human tutors when designing their systems. This is clearly a sensible place to start, as we indicated in the Introduction, but reactions from a human teacher and a human peer do not always have the same effect as "identical" reactions delivered from systems. For example, while a human tutor might get away with refusing a help request when asked, it is questionable as to whether this would be acceptable behaviour from a system. In a similar vein much effort is focused on ensuring that animated pedagogical agents display reactions that mimic their human counterparts. Again this remains somewhat of an open question as is the utility of the realism of the rendered face of the agent (see e.g., Haake & Gulz, 2009).

ACKNOWLEDGEMENTS

We thank many people for helpful comments on drafts of this paper including Paul Brna, the anonymous referees for the International Journal of Artificial Intelligence in Education, Kaska Porayska-Pompsta and various members of the IDEAS lab at Sussex University.

REFERENCES

- Aleven, V., McLaren, B., Roll, I., & Koedinger, K. (2006). Toward meta-cognitive tutoring: A model of help seeking with a Cognitive Tutor. *International Journal of Artificial Intelligence in Education*, 16(2), 101-128.
- Alexander, S. (2005). Modeling Emotions from Non-verbal Behaviour in an Affective Tutoring System. 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. 493-495): Springer.

- Ames, C. (1992). Classrooms: Goals, Structures, and Student Motivation. *Journal of Educational Psychology*, 84(3), 261-271.
- Arroyo, I., Cooper, D. G., Burleson, W., Woolf, B. P., Muldner, K., & Christopherson, R. (2009). Emotion Sensors Go to School. In V. Dimitrova, R. Mizoguchi, B. du Boulay & A. Grasser (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.
- Arroyo, I., & Woolf, B. P. (2005). Inferring learning and attitudes from a Bayesian Network of log file data. In C.-K. Looi, G. McCalla, B. Bredeweg & J. Breuker (Eds.), *Artificial Intelligence in Education, Supporting Learning through Intelligent and Socially Informed Technology* (Vol. Frontiers in Artificial Intelligence and Applications 125, pp. 33-40): IOS Press.
- Arroyo, I., Woolf, B. P., Royer, J. M., & Tai, M. (2009). Affective Gendered Learning Companions. In V. Dimitrova, R. Mizoguchi, B. du Boulay & A. Grasser (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. 41-48): 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. S. J., Rodrigo, M. M. T., & Xolocotzin, U. E. (2007). The Dynamics of Affective Transitions in Simulation Problem-solving Environments. In A. Paiva, R. Prada & R. W. Picard (Eds.), Affective Computing and Intelligent Interaction: Second International Conference, ACII 2007, Lisbon, Portugal, Proceedings (Vol. Lecture Notes in Computer Science 4738, pp. 666-677): Springer.
- Balaam, M., Harris, E., & Fitzpatrick, G. (2009). Encouraging Classrooms that Care Artificial Intelligence in Education. Building Learning Systems that Care: From Knowledge Representation to Affective Modelling (Vol. Frontiers in Artificial Intelligence and Applications 200, pp. 798): IOS Press.
- Bandura, A. (1997). Self-efficacy: The exercise of control: New York: Freeman.
- Barrow, D., Mitrovic, A., Ohlsson, S., & Grimley, M. (2008). Assessing the Impact of Positive Feedback in Contsrint-Based Tutors. 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. 250-259): Springer.
- Baylor, A. L., & Kim, Y. (2005). Simulating Instructional Roles through Pedagogical Agents. *International Journal of Artificial Intelligence in Education*, 15(2), 95-115.
- Baylor, A. L., & Plant, E. A. (2005). Pedagogical agents as social models for engineering: The influence of agent appearance on female choice. In C.-K. Looi, G. McCalla, B. Bredeweg & J. Breuker (Eds.), Artificial Intelligence in Education, Supporting Learning through Intelligent and Socially Informed Technology (Vol. Frontiers in Artificial Intelligence and Applications 125, pp. 65-72). Amsterdam: IOS Press.
- Baylor, A. L., Warren, D., Park, S., Shen, E., & Perez, R. (2005). The impact of frustration-mitigating messages delivered by an interface agent. In C.-K. Looi, G. McCalla, B. Bredeweg & J. Breuker (Eds.), *Artificial Intelligence in Education, Supporting Learning through Intelligent and Socially Informed Technology* (Vol. Frontiers in Artificial Intelligence and Applications 125): IOS Press.
- Beal, C. R., & Lee, H. (2005). Creating a pedagogical model that uses student self reports of motivation and mood to adapt ITS instruction *Workshop on Motivation and Affect in Educational Software, at AIED2005, 12th International Conference on Artificial Intelligence in Education* (pp. 39-46). Amsterdam.
- Blanchard, E., & Frasson, C. (2004). An Autonomy-Oriented System Design for Enhancement of Learner's Motivation in E-learning. In J. C. Lester, R. M. Vicari & F. Paraguacu (Eds.), 7th International Conference on Intelligent Tutoring Systems, ITS 2004, Maceio, Brazil, Proceedings (Vol. Lecture Notes in Computer Science 3220, pp. 34-44): Springer.

- 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. Grasser (Eds.), Artificial Intelligence in Education. Building Learning Systes that Care: from Knowledge Representation to Affective Modelling (Vol. Frontiers in Artificial Intelligence and Applications 200, pp. 81-88): IOS Press.
- Boekaerts, M., de Koning, E., & Vedder, P. (2006). Goal-directed behavior and contextual factors in the classroom: An innovative approach to the study of multiple goals. *Educational Psychologist*, 41(1), 33-51.
- Boyer, K. E., Phillips, R., Wallis, M., Vouk, M., & Lester, J. (2008). Balancing Cognitive and Motivational Scaffolding in Tutoprial Dialogue. 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. 239-249): Springer.
- Bregman, N. J., & McAllister, H. A. (1982). Motivation and skin temperature biofeedback: Yerkes-Dodson revisited. *Psychophysiology*, 19(3), 282-285.
- Bull, S., Greer, J., & McCalla, G. (2003). The Caring Personal Agent. *International Journal of Artificial Intelligence in Education*, 13(1), 21-34.
- Chaffar, S., & Frasson, C. (2004). Inducing optimal emotional state for learning in Intelligent Tutoring Systems In J. C. Lester, R. M. Vicari & F. Paraguacu (Eds.), 7th International Conference on Intelligent Tutoring Systems, ITS 2004, Maceio, Brazil, Proceedings (Vol. Lecture Notes in Computer Science 3220, pp. 45-54): Springer.
- Chen, Z.-H., Chou, C.-Y., Deng, Y.-C., & Chan, T.-W. (2007). Active Open Models as Animal Companions: Motivating Children to Learn through Intercating with My-Pet and Our-Pet. *International Journal of Artificial Intelligence in Education*, 17(2), 145-167.
- Chen, Z.-H., Deng, Y.-C., Chou, C.-Y., & Chan, T.-W. (2005). Motivating Learners by Nurturing Animal Companions: My-Pet and Our-Pet. In C.-K. Looi, G. McCalla, B. Bredeweg & J. Breuker (Eds.), Artificial Intelligence in Education: Supporting Learning through Intelligent and Socially Informed Technology (Vol. Frontiers in Artificial Intelligence and Applications 125, pp. 136-143): 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.
- Conati, C., & Manske, M. (2009). Adaptive Feedback in an Educational Game for Number Factorisation. In V. Dimitrova, R. Mizoguchi, B. du Boulay & A. Grasser (Eds.), *Artificial Intelligenec in Education. Building Learning Systems that Care: from Knowledge Representation to Affective Modelling* (Vol. Frontiers in Artificial Intelligence and Applications 200, pp. 581-583): IOS Press.
- Conati, C., & Van Lehn, K. (2000). Toward Computer-Based Support of Meta-Cognitive Skills: a Computational Framework to Coach Self-Explanation *International Journal of Artificial Intelligence in Education*, 11(4), 389-415.
- Conati, C., & Zhou, X. (2002). Modeling Students' Emotions from Cognitive Appraisal in Educational Games. In S. A. Cerri, G. Guy & F. Paraguacu (Eds.), *Intelligent Tutoring Systems. 6th International Conference, ITS2002, Biarritz, France and San Sebastian, Spain, Proceedings* (Vol. Lecture Notes in Computer Science 2363, pp. 944-954). Berlin: Springer.
- Cooper, B. (2003). Care making the Affective Leap: More Than a Concerned Interest in a Learner's Cognitive Abilities. *International Journal of Artificial Intelligence in Education*, 13(1), 3-9.
- D'Mello, S., Graesser, A., & Picard, R. W. (2007). Toward an affect-sensitive AutoTutor. *IEEE Intelligent Systems*, 22(4), 53-61.
- 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.
- D'Mello, S., Person, N., & Lehman, B. (2009). Antecedent-Consequent realtionships amdf Cyclical Patterns between Affective States and Problem Solving Outcomes. In V. Dimitrova, R. Mizoguchi, B. du Boulay &

A. Grasser (Eds.), Artificial Intelligence in Education. Building Learning Systems that Care: from Knowledge Representation to Affective Modelling (pp. 57-64): IOS Press.

- D'Mello, S., Taylor, R., Davidson, K., & Graesser, A. (2008). Self Versus Teacher Judgments of Learner Emotions During a Tutoring Session with AutoTutor In B. P. Woolf, E. Aïmeur, R. Nkambou & S. Lajoie (Eds.), *Intelligent Tutoring Systems, 9th International Conference, ITS 2008, Montreal, Canada, Proceedings* (Vol. Lecture Notes in Computer Science 5091, pp. 9-18): Springer.
- D'Mello, S. K., Craig, S. D., Sullins, J., & Graesser, A. C. (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.
- De Rosis, F. (2001). Preface: Towards adaptation of interaction to affective factors. User Modeling and User-Adapted Interaction, 11(4), 267-278.
- de Vicente, A., & Pain, H. (2002). Informing the detection of the students' motivational state: an empirical study. In S. A. Cerri, G. Guy & F. Paraguacu (Eds.), *Intelligent Tutoring Systems. 6th International Conference*, *ITS2002, Biarritz, France and San Sebastian, Spain, Proceedings* (Vol. Lecture Notes in Computer Science 2363, pp. 933-943): Springer.
- del Soldato, T. (1994). Motivation in tutoring systems. University of Sussex.
- 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.
- Dimitrova, V. (2003). STyLE-OLM: Interactive Open Learner Modelling. International Journal of Artificial Intelligence in Education, 13(1), 35-78.
- Dragon, T., Arroyo, I., Woolf, B. P., Burleson, W., Kaliouby, R., & Eydgahi, H. (2008). Viewing Student Affect and Learning through Classoom Observation and Physical Sensors. 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. 29-39): Springer.
- du Boulay, B. (in press). Towards a Motivationally-Intelligent Pedagogy: How should an intelligent tutor respond to the unmotivated or the demotivated? In R. A. Calvo & S. D'Mello (Eds.), *New Perspectives on Affect and Learning Technologies*. New York: Springer.
- du Boulay, B., & Luckin, R. (2001). The plausibility problem: An initial analysis. In M. Beynon, C. L. Nehaniv & K. Dautenhahn (Eds.), 4th International Conference on Cognitive Technology - Instruments of Mind, Coventry, England (pp. 289-300).
- du Boulay, B., Luckin, R., Martinez-Miron, E., Rebolledo-Mendez, G. R., & Harris, A. (2008). Motivationally Intelligent Educational Systems: Three Questions *Second International Conference on Innovations in Learning for the Future, Future e-Learning 2008* (Vol. 4793, pp. 1-10). Istanbul: University Rectorate Publication.
- 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.), Proceeding of the 2007 conference on Artificial Intelligence in Education: Building Technology Rich Learning Contexts that Work (Vol. Frontiers in Artificial Intelligence and Applications 158, pp. 563-565). Amsterdam: IOS Press.
- Dweck, C. S., & Leggett, E. L. (1988). A social-cognitive approach to motivation and personality. *Psychological review*, *95*(2), 256-273.
- Eccles, J. S. (1983). Expectancies, values, and academic behavior. In J. T. Spencer (Ed.), *Achievement and achievement motivation*. San Francisco: W.H. Freeman.
- Ekman, P., Friesen, W. V., & Ellsworth, P. C. (1972). Emotion in the Human Face. New York: Pergamon Press
- Flavell, J. H. (1979). Metacognition and cognitive monitoring: A new area of cognitive-developmental inquiry. *American Psychologist*, *34*(10), 906-911.
- Forbes-Riley, K., & Litman, D. (2009). Adapting to Student Uncertainty Improves Tutoring Dialogues. In V. Dimitrova, R. Mizoguchi, B. du Boulay & A. Grasser (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. 33-40): IOS Press.

- Gama, C. (2004). Metacognition in Interactive Learning Environments: The Reflection Assistant Model. In J. C. Lester, R. M. Vicari & F. a. Paraguacu (Eds.), 7th International Conference on Intelligent Tutoring Systems, ITS 2004, Maceio, Brazil, Proceedings (Vol. Lecture Notes in Computer Science 3220, pp. 668-677): Springer.
- 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.
- Gulz, A., & Haake, M. (2005). *Social and visual style in virtual pedagogical agents*. Paper presented at the Workshop on Adapting the Interaction Style to Affective Factors associated with the 10th International Conference on User Modeling.
- Haake, M., & Gulz, A. (2009). A Look at the Roles of Look & Roles in Embodied Pedagogical Agents -- A User Preference Perspective. *International Journal of Artificial Intelligence in Education*, 19, 39-71.
- Harackiewicz, J., Barron, K., & Elliot, A. (1998). Rethinking achievement goals: When are they adaptive for college students and why? *Educational Psychologist*, 33(1), 1-21.
- Heraz, A., & Frasson, C. (2009). Predicting Learner Ansers Correctness through Brainwaves Assessment and Emotional Dimensions. In V. Dimitrova, R. Mizoguchi, B. du Boulay & A. Grasser (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. 49-56): IOS Press.
- Hernandez, Y., & Noguez, J. (2005). Affective Behaviour in In telligent Tutoring Systems for Virtual Laboratories. In C.-K. Looi, G. McCalla, B. Bredeweg & J. Breuker (Eds.), Artificial Intelligence in Education, Supporting Learning through Intelligent and Socially Informed Technology (Vol. Frontiers in Artificial Intelligence and Applications 125, pp. 960-960): IOS Press.
- Herrington, J., & Herrington, T. (2005). *Authentic learning environments in higher education*: Information Science Pub.
- Howland, K., du Boulay, B., & Good, J. (2009). Fostering engaged and directed learning by actovity foregrounding and backgrounding. In V. Dimitrova, R. Mizoguchi, B. du Boulay & A. Grasser (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. 241-248): IOS Press.
- Johnson, L. W. (2007). Serious Use of a Serious Game for Language Learning. In R. Luckin, K. R. Koedinger & J. Greer (Eds.), *Proceeding of the 2007 conference on Artificial Intelligence in Education: Building Technology Rich Learning Contexts that Work* (Vol. Frontiers in Artificial Intelligence and Applications 158, pp. 67-74). Amsterdam: IOS Press.
- Johnson, L. W., Rickel, J. W., & Lester, J. C. (2000). Animated Pedagogical Agents: Face-to-Face Interaction in Interactive Learning Environments. *International Journal of Artificial Intelligence in Education*, 11(1), 47-78.
- Johnson, L. W., & Rizzo, P. (2004). Politeness in tutoring dialogues. "Run the factory. That's what I'd do". In J. C. Lester, R. M. Vicari & F. Paraguacu (Eds.), 7th International Conference on Intelligent Tutoring Systems, ITS 2004, Maceio, Brazil, Proceedings (Vol. Lecture Notes in Computer Science 3220, pp. 67-76): Springer.
- Kapoor, A., Burleson, W., & Picard, R. (2007). Automatic prediction of frustration. International Journal of Human-Computer Studies, 65(8), 724-736.
- Katz, S., Allbritton, D., & John, C. (2003). Going Beyond the Problem Given: How Human Tutors Use Post-Solution Discussions to Support Transfer. *International Journal of Artificial Intelligence in Education*, 13(1), 79-116.
- Kay, J., & McCalla, G. (2003). The Careful Double Vision of Self. International Journal of Artificial Intelligence in Education, 13(1), 11-18.
- Keller, J. M. (1983). Motivational design of instruction. In C. M. Reigluth (Ed.), *Instructional design theories and models: An overview of their current status* (pp. 386-434). Hillsdale, NJ: Lawrence Erlbaum.

- Kelly, D., & Weibelzhal, S. (2006). Raising Confidence Levels Using Motivational Contingency Design Techniques. In M. Ikeda, K. D. Ashley & T.-W. Chan (Eds.), *Intelligent Tutoring Systems: 8th International Conference, ITS 2006, Jhongli, Taiwan, Proceedings* (Vol. Lecture Notes in Computer Science 4053, pp. 535-544): Springer.
- Kim, Y. (2005). Pedagogical Agents as Learning Companions: Building Social Relations with Learners. In C.-K. Looi, G. McCalla, B. Bredeweg & J. Breuker (Eds.), *Artificial Intelligence in Education: Supporting Learning through Intelligent and Socially Informed Technology* (Vol. Frontiers in Artificial Intelligence and Applications 125, pp. 362-369): IOS Press.
- Kim, Y. (2007). Desirable Characteristics of Learning Companions. *International Journal of Artificial Intelligence in Education*, 17(4), 371-388.
- Kleinsmith, A., De Silva, P. R., & Bianchi-Berhouze, N. (2005). Recognizing Emotion from Postures: Cross-Cultural Differences in User Modelling. 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. 50-59): Springer.
- Koedinger, K. R., Anderson, J. R., Hadley, W. H., & Mark, M. A. (1997). Intelligent Tutoring Goes to School in the Big City. *International Journal of Artificial Intelligence in Education*, 8(1), 30-43.
- Kort, B., & Reilly, B. (2002). An Affective Module for an Intelligent Tutoring System. In S. A. Cerri, G. Guy & F. Paraguacu (Eds.), *Intelligent Tutoring Systems. 6th International Conference, ITS2002, Biarritz, France and San Sebastian, Spain, Proceedings* (Vol. Lecture Notes in Computer Science 2363, pp. 955-962): Springer.
- Leelawong, K., & Biswas, G. (2008). Designing Learning by Teaching Agents: The Betty's Brain System. International Journal of Artificial Intelligence in Education, 18(3), 181-208.
- 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. 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. 50-59): Springer.
- Lepper, M. R., Aspinwall, L. G., Mumme, D. L., & Chabay, R. W. (1990). Self perception and social-perception processes in tutoring: Subtle social control strategies of expert tutors In J. M. Olson & M. P. Zanna (Eds.), *Self-Inference Processes: The Ontario Symposium* (Vol. 6, pp. 217-237). Hillsdale, New Jersey: Lawrence Erlbaum Associates.
- Luckin, R. (2010). Re-designing Learning Contexts. london: Routledge.
- Luckin, R., & Hammerton, L. (2002). Getting to know me: helping learners understand their own learning needs through metacognitive scaffolding *Intelligent Tutoring Systems*. 6th International Conference, ITS2002, Biarritz, France and San Sebastian, Spain, Proceedings (Vol. Lecture Notes in Computer Science 2363, pp. 759-771): Springer.
- Malone, T. W., & Lepper, M. R. (1987). Making Learning Fun: A taxonomy of intrinsic motivations for learning. Aptitude, learning and instruction. In R. E. Snow & M. J. Farr (Eds.), *Aptitude Learning and Instruction* (Vol. Volume 3: Conative and Affective Process Analyses, pp. 223-253). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.
- Manske, M., & Conati, C. (2005). Modelling learning in an educational game. In C.-K. Looi, G. McCalla, B. Bredeweg & J. Breuker (Eds.), *Artificial Intelligence in Education: Supporting Learning through Intelligent and Socially Informed Technology* (Vol. Frontiers in Artificial Intelligence and Applications 125, pp. 411-418): IOS Press.
- Marsella, S. C., Johnson, L. W., & LaBore, C. M. (2003). Interactive Pedagogical Drama for Health Interventions. In U. Hoppe, F. Verdejo & J. Kay (Eds.), (Vol. Frontiers in Artificial Intelligence and Applications 97, pp. 341-348). Artificial Intelligence in Education: Shaping the Future of Learning through Intelligent Technologies: IOS Press.
- Martinez-Miron, E., Harris, A., du Boulay, B., Luckin, R., & Yuill, N. (2005). The role of learning goals in the design of ILEs: Some issues to consider. In C.-K. Looi, G. McCalla, B. Bredeweg & J. Breuker (Eds.),

Artificial Intelligence in Education: Supporting Learning through Intelligent and Socially Informed Technology (Vol. Frontiers in Artificial Intelligence and Applications 125, pp. 427-434): IOS Press.

- Mavrikis, M., Maciocia, A., & Lee, J. (2007). Towards Predictive Modelling of Student Affect from Web-Based Interactions. In R. Luckin, K. R. Koedinger & J. Greer (Eds.), *Proceeding of the 2007 conference on Artificial Intelligence in Education: Building Technology Rich Learning Contexts that Work* (Vol. Frontiers in Artificial Intelligence and Applications 158, pp. 169-176). Amsterdam: IOS Press.
- McQuiggan, S. W., Lee, S., & Lester, J. C. (2007). Early Prediction of Student Frustration. In A. Paiva, R. Prada & R. W. Picard (Eds.), Affective Computing and Intelligent Interaction: Second International Conference, ACII 2007, Lisbon, Portugal, Proceedings (Vol. Lecture Notes in Computer Science 4738): Springer.
- McQuiggan, S. W., & Lester, J. C. (2006). Diagnosing Self-efficacy in Intelligent Tutoring Systems: An Empirical Study. In M. Ikeda, K. D. Ashley & T.-W. Chan (Eds.), *Intelligent Tutoring Systems: 8th International Conference, ITS 2006, Jhongli, Taiwan, Proceedings* (Vol. Lecture Notes in Computer Science 4053, pp. 565-573): Springer.
- Mozziconacci, S. J. L. (2001). Modeling Emotion and Attitude in Speech by Means of Perceptually Based Parameter Values. User Modeling and User-Adapted Interaction, 11(4), 297-326.
- Murphy, P., & Alexander, P. (2000). A motivated exploration of motivation terminology. Contemporary Educational Psychology, 25(1), 3-53.
- Ortony, A., Clore, G. L., & Collins, A. (1988). *The Cognitive Structure of Emotions*. : Cambridge University Press.
- Paiva, A., Prada, R., & Picard, R. W. (Eds.). (2007). Affective Computing and Intelligent Interaction: Second International Conference, ACII 2007, Lisbon, Portugal, Proceedings (Vol. Lecture Notes in Computer Science 4738): Springer.
- Pekrun, R., Goetz, T., Titz, W., & Perry, R. P. (2002). Academic Emotions in Students' Self-Regulated Learning and Achievement: A Program of Qualitative and Quantitative Research. *Educational Psychologist*, 37(2), 91-105.
- Picard, R. W. (2000). Affective Computing. Cambridge, Massachusetts: MIT Press.
- Pintrich, P. (2003). Motivation and Classroom Learning. *Handbook of Psychology: Educational Psychology*, 7, 103-122.
- Porayska-Pomsta, K., Mavrikis, M., & Pain, H. (2008). Diagnosing and acting on student affect: the tutor's perspective. User Modeling and User-Adapted Interaction, 18, 125-173.
- Porayska-Pomsta, K., & Pain, H. (2004). Providing cognitive and affective scaffolding through teaching strategies: applying linguistic politenss to the educational context. In J. C. Lester, R. M. Vicari & F. Paraguacu (Eds.), *7th International Conference on Intelligent Tutoring Systems, ITS 2004, Maceio, Brazil, Proceedings* (Vol. Lecture Notes in Computer Science 3220, pp. 77-86): Springer.
- Prendinger, H., Mori, J., & Ishizuka, M. (2005). Recognizing, Modeling, and Responding to Users' Affective States. 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. 60-69): Springer.
- Rebolledo-Mendez, G., du Boulay, B., & Luckin, R. (2006). *Motivating the learner: An empirical evaluation*. Paper presented at the 8th International Conference on Intelligent Tutoring Systems, Jhongli, TAIWAN.
- Robison, J., McQuiggan, S. W., & Lester, J. C. (2009). Modeling Task-Based vs. Affect-based Feedback Brhaviour in Pedagogical Agents: An Inductive Approach. In V. Dimitrova, R. Mizoguchi, B. du Boulay & A. Grasser (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. 25-32): IOS Press.
- 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.
- Ryan, R. M., & Deci, E. L. (2000). Intrinsic and Extrinsic Motivations: Classic Definitions and New Directions. *Contemporary Educational Psychology*, 25, 54-67.

- Schunk, D. H., Pintrich, P. R., & Meece, J. L. (2008). *Motivation in Education: Theory, Research and Applications* (3rd ed.): Pearson, Merrill, Prentice Hall.
- Self, J. (1999). The defining characteristics of intelligent tutoring systems research: ITSs care, precisely. *International Journal of Artificial Intelligence in Education*, 10, 350-364.
- Soller, A., & Lesgold, A. (2003). A computational approach to analyzing online knowledge sharing interaction. In U. Hoppe, F. Verdejo & J. Kay (Eds.), *Artificial Intelligence in Education: Shaping the Future of Learning through Intelligent Technologies* (Vol. Frontiers in Artificial Intelligence and Applications 97, pp. 253-268): IOS Press.
- Song, S. H., & Keller, J. M. (2001). Effectiveness of motivationally adaptive computer-assisted instruction on the dynamic aspects of motivation. *Educational Technology Research and Development*, 49(2), 5-22.
- Tedesco, P. A. (2003). MArCo: Building an Artificial Conflict Mediator to Support Group Planning Interactions. International Journal of Artificial Intelligence in Education, 13(1), 117-155.
- Uresti, J. A. R., & du Boulay, B. (2004). Expertise, Motivation and Teaching in Learning Companion Systems. International Journal of Artificial Intelligence in Education, 14(2), 193-231.
- Wentzel, K. R. (2002). Are effective teachers like good parents? Teaching styles and student adjustment in early adolescence. *Child Development*, 287-301.
- Yussof, M. Z., & du Boulay, B. (2009). Can Relaxtion Exercises Improve Learning? In V. Dimitrova, R. Mizoguchi, B. du Boulay & A. Grasser (Eds.), Artificial Intelligence in Education. Building Learning Systes that Care: from Knowledge Representation to Affective Modelling (Vol. Frontiers in Artificial Intelligence and Applications 200): IOS Press.
- 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.
- Zhang, L., Barden, J. A., & Hendley, R. J. (2005). Affect detection and metaphor in e-drama Workshop on Motivation and Affect in Educational Software, at AIED2005, 12th International Conference on Artificial Intelligence in Education (pp. 56-58). Amsterdam.
- Zhang, L., Gillies, M., Dhaliwal, K., & Gower, A. (2009). E-Drama: Facilitating Online Role-play using an AI Actor and Emotionally Expressive Characters. *International Journal of Artificial Intelligence in Education*, 19, 5-38.