Modelling Human Teaching Tactics and Strategies for Tutoring Systems: 14 years on

Benedict du Boulay, *Human-Centred Technology, Department of Informatics,* University of Sussex, BN1 9QJ, UK **Rosemary Luckin,** The London Knowledge Lab, Institute of Education, London, WC1N 3QS, UK

Abstract: Our original paper tried to characterize the richness of the teaching repertoire of expert human teachers and to give a sense of how far there still was to go in the development of pedagogic expertise in AIED systems. It considered three ways in which more expert teaching strategies and tactics might be developed. These were via (i) the observation of human expert teachers, (ii) theoretical derivation from learning theories, and (iii) empirical observation of human and simulated students. This paper reflects on the original paper and briefly sketches progress since 2001.

Keywords: Pedagogy, teaching strategy, teachers

INTRODUCTION AND MOTIVATION

We are very pleased to have the opportunity to review our paper "Modelling Human Teaching Tactics and Strategies for Tutoring Systems" some 14 years after its initial publication (du Boulay & Luckin, 2001). Our impression is that it has served as a good jumping off point for PhD students and others interested in pedagogy and getting into the field of Artificial Intelligence in Education.

By 2001 AIED systems had continued to evolve in rather a lopsided way. While the domain representations, student models and interfaces had developed strongly, it seemed that their teaching knowledge and skill had not developed so well. Indeed, in the 1980s there had been critiques of the impoverished repertoire of teaching tactics and strategies available in AIED systems compared with human expert teachers (Carroll & McKendree, 1987; Ohlsson, 1987; Ridgway, 1988). So the purpose of our paper was to explore what kind of progress, if any, had been made since those earlier critiques.

APPROACH

Our approach was first to try to characterize the richness of the teaching repertoire of expert human teachers working one-to-one with their students as a goal towards which AIED system might aspire (see e.g., Bloom, 1984). We used the theoretical model proposed by Ohlsson (1987) as a starting point. This model, even though limited to the kinds of teaching actions associated with teaching a symbolic procedure (such as multiplying fractions), demonstrated at least in part, the complexity one would expect in an expert machine teacher.

Our paper was not a comprehensive review. Rather it attempted to provide a characterization of the issues. It considered three ways in which more expert teaching strategies and tactics might be developed. These were via (i) the observation of human expert teachers, (ii) theoretical derivation from learning theories, and (iii) empirical observation of human and simulated students. In looking at these three sources of ideas, particular attention was paid to the techniques of "dealing with student errors" and "motivating students" as these two areas are central issues in any pedagogy.

The paper then went on to describe two systems developed by the authors: one on the pedagogy of motivation (del Soldato & du Boulay, 1995), and the other on the pedagogy of Vygotsky's notion of the Zone of Proximal Development (Luckin & du Boulay, 1999). We note with pleasure that the papers

describing these two pieces of work in detail in IJAIED are also part of this special issue (REFS to special issue if these papers are accepted).

Observation of human expert teachers

While the field of education had studied the skill of teaching for centuries, much of its work was couched at a general level that was hard to implement in AIED systems (see e.g. Rutter, Maughan, Mortimore, & Ouston, 1979), with perhaps Socratic Tutoring as an interesting counter-example (Collins, Warnock, Aiello, & Miller, 1975). But there was an increasing body of work that had observed and codified expert teaching at a fine level of granularity. From these records it was possible to extract general teaching strategies and specific teaching tactics, as well as to compare and contrast these with those available in AIED systems (see e.g. Graesser, Person, Harter, & the Tutoring Research Group, 2000; Lajoie, Wiseman, & Faremo, 2000; Leinhardt & Greeno, 1991; Lepper, Aspinwall, Mumme, & Chabay, 1991).

Derived from learning theory

The paper considered a number of learning theories from which AIED teaching strategies had been derived. First were *epistemological* theories where the focus was on the subtle way that information is transformed into knowledge and then knowledge into understanding and skill. For example, Contingent Teaching (Wood, Wood, & Middleton, 1978) acknowledges the learner's need for independence and agency within a scaffolded learning experience. This approach involves the Vygotskian notion of the learner being able to achieve with the help of scaffolding more than he or she could achieve unaided. The collaborative success drives the acquisition of greater skill and understanding. From this view derives a regime to guide the teacher in offering such help in the most efficacious manner. From ACT and ACT* (Anderson, 1990) came the idea of the transformation of declarative knowledge into procedural skill and thus the value of goal setting, graded exposure to more complex aspects of the domain via model-tracing and knowledge-tracing, and immediate feedback on errors so that they could be acted upon soon after an error had been committed (Anderson, Corbett, Koedinger, & Pelletier, 1995).

Second were those *reflective* theories that in their different ways embodied the idea that *two* complementary psychological processes operate within a learner: the one focusing on the domain itself and the other reflecting on how far that primary focus on the domain is leading to secure understanding. From this insight various ways for the teacher to support each process were derived, but particularly the second, the metacognitive one. These theories included Conversation Theory (Pask, Kallikourdis, & Scott, 1975), Reciprocal Teaching (Chan & Chou, 1997), Self-Explanation (Chi, Bassok, Lewis, Reimann, & Glaser, 1989) and Self-regulation (Winne, 1997). Pask's Conversation Theory described the interaction of the two processes in formal terms, essentially as a contribution to a cybernetic theory of understanding. His work was applied in the design and development of various computer-based learning systems some of whose interactions had a reflective component similar to reciprocal teaching, where the learner and the teacher take turns to explain to each other the subject matter to be understood and learned (for a brief guide to Pask's educational work, see e.g., Entwistle, 1978).

Of course there were many overlaps between teaching strategies derived from epistemological theories of learning and those derived from reflective theories of learning. For example, under Akhras and Self's (2000) view of Constructivism, teaching was aimed at emphasising:

"the process rather than the product of learning. The theoretical models that constitute our approach enable intelligent learning environments to evaluate learning according to four properties of constructivist learning processes: cumulativeness, constructiveness, self-regulatedness, and reflectiveness, and to make decisions about the learning opportunities to be provided to the learners, taking into consideration the affordances of learning situations regarding these properties" (Akhras & Self, 2000, page 344).

They describe four properties of sequences of *constructivist* learning interactions. Cumulativeness refers to the property of a sequence of learning interactions that involves the same entity being experienced more than once during the sequence. Constructiveness refers to the property of a sequence of learning interactions where "entities experienced by the learner in one situation are in some way related to new entities that the learner generates in a later situation." Self-regulatedness refers to that property of a sequence of their earlier actions with

a view to guiding what they do next. Finally Reflectiveness refers to the property of a sequence of learning interactions that involves a learner engaging in reflective activities on earlier episodes in that interaction.

Observing students

The brief section in the paper on work with real students looked at research on individual differences, such as gender (Arroyo, Beck, Woolf, Beal, & Schultz, 2000) and ability (Shute, 1993), and their consequences for differentiated teaching, typically via macro-adaptation to the overall style of interaction. Macro-adaptation is where a teaching or learning system is adapted, or adapts itself, to one particular group of learners as opposed to another group, for example novices versus semi-experts. This is contrasted with micro-adaptation where the adjustment is on an individual and typically dynamic basis.

There was also the intriguing work by VanLehn with simulated students, i.e. modelling human learning behaviour and then trying out different kinds of teaching on such a model to see which method worked best. From these VanLehn derived "felicity conditions" for the optimal structure and sequence of worked examples, i.e. the rules for maximising the educational effectiveness of a sequence of worked-examples (VanLehn, Ohlsson, & Nason, 1994).

CORE CONTRIBUTIONS AND LIMITATIONS

The core contribution of the paper was to emphasise the rich variability of human teaching with its roots in general communicative competence. While there are some specialized tactics that human teachers apply effectively, good teaching derives from the conversational and social interactive skills used in everyday settings such as listening, eliciting, intriguing, motivating, cajoling, explaining, arguing, persuading, enthralling, leading, pleading and so on. Implicitly the message was that neither learners nor teachers are disembodied cognitive entities engaged in symbolic knowledge sharing but rather are feeling and thinking beings living and working in a particular educational, social and cultural context. A secondary contribution was to show how far there was still to go before we could reasonably designate any AIED system as modelling expert teaching capability.

The core limitation of the paper was that it was not a comprehensive review of the state of the art in the implementation of teaching tactics and strategies in AIED systems. So some of the subsections were rather patchy in their coverage of the literature. A second limitation, probably a consequence of the first, was that we did not stand far enough back from the work we were describing to be able to organize the learning and teaching theories into clear categories focusing on different issues.

PRACTICAL IMPACT AND PROGRESS SINCE 2001

The original review paper was a call to arms rather than a technical advance. While the paper was not a comprehensive review of the field, it turned out to be a useful starting point for many researchers starting out to understand what was known at the time. The first author's own later work on pedagogy has tended to focus on motivation and metacognition (see e.g., du Boulay, 2011; du Boulay et al., 2010), while the second author has focused on further use of Vygotsky's Zone of Proximal Development (Luckin, 2010). We note with pleasure the distinct shift towards taking the affective, motivational and metacognitive dimensions of pedagogy much more seriously in the last few years (see e.g., Narciss, Sosnovsky, & Andres, 2014). In this section we briefly review the work on teaching strategies over the last 14 years that have tried to match the performance of expert teachers in a one-to-one tutorial situation. As in the original paper, this is divided into subsections on observation of expert teachers, on recently elaborated teaching and learning theories and on the observation of students. Also, as for the first paper, it does not claim to be comprehensive. It is important to be clearer about what we mean by "expert teachers". Lepper and Woolverton (2002) provide an excellent empirically observed account of expert human teachers working one-to-one via their INSPIRE (Intelligent, Nurturant, Socratic, Progresseive, Indirect, Reflective, Encouraging) framework. They characterize the behaviour of such teachers, as compared to the less expert, along a number of dimensions including their use of questions rather than didactic statements, their

attention to motivating and encouraging their students, and the attention they paid to fostering learner reflection on process. Their account emphasizes the richness of the teaching repertoire of expert human teachers while enumerating some of the tactics that make such teachers effective:

"Our best tutors are those who are concerned *simultaneously* with students' learning on the one hand and their motivation on the other. Thus, these tutors do not consider their task to be merely the efficient provision of feedback and information as some early theories of learning might have implied Nor are they willing to sacrifice learning for the sake of motivation, as critics of the socalled" self-esteem" movement in schools have described Rather than "dumbing down" the instructional content by presenting easy problems or preventing student errors in an attempt to preserve students' self-esteem, these tutors demonstrate knowledge of a wide array of systematic techniques, both for presenting information to students and for encouraging student involvement and persistence at a task" (Lepper & Woolverton, 2002, page 151).

In terms of understanding recent progress, largely in the cognitive rather than the motivational aspects of modelling teaching, we note a meta-analysis of the learning outcomes of using Intelligent Tutoring Systems (Ma, Adesope, Nesbit, & Liu, 2014). They found that in general an ITS "outperformed, in aggregate, the other modes of instruction to which it was compared in evaluative studies", but note that this may be due in part to publication bias (in other wards, that negative results or null results do not get published). However, they are more positive that "in some situations ITS can successfully complement and substitute for other instructional modes."

In a similar vein, VanLehn (2011) provides meta-review comparing the differences and effectiveness of (typically non-expert) human and intelligent tutoring systems. His paper enumerates the possible reasons why human tutoring *might in principle* be more effective than computer tutors. These include: detailed diagnostic assessments, individualized task selection, sophisticated tutorial strategies, learner control of dialogues, broader domain knowledge, motivation, feedback, and the potential for the tutor to elicit effective learning behaviour in their students. The paper found that not all of the reasons above applied in practice, that intelligent tutoring systems were now in some respects approaching the effectiveness of "ordinary" human tutors and that the effectiveness of the latter were rather less than the gold-standard of individual expert tutoring (Bloom, 1984). Note however that both the meta-analysis and the meta-review rather underplay the motivational aspects of expert teaching behavior, as emphasised by Lepper and Woolverton (2002).

Observing expert teachers

Work on the direct observation of expert human teachers has continued. A major study observed a carefully selected sample of 12 expert maths and science teachers tutoring students who had difficulties in those subjects (Olney, Person, & Graesser, 2011). So the lessons they observed were essentially doing remedial work within a STEM context with a mixture of helping students with problem-solving as well as (re-)introducing topics that needed remediating. The researchers recorded fifty one-hour tutoring sessions and, after annotating and coding the transcripts, produced a three level analysis of the teaching tactics and strategies that they observed. At the base level they observed the teachers employing 12 different "motivational moves" and 15 different "pedagogical moves". Examples of motivation moves were the use of humour or providing negative, positive or neutral feedback. Examples of pedagogical moves were providing a counter example or paraphrasing what the student had just said. These tactics formed the basic building blocks of the expert teachers' behaviour. At the top level they observed 8 modes that the teachers might be in: introducing, lecturing, highlighting, modeling, scaffolding, fading, talking off topic, and concluding. Between the two levels Olney and his colleagues used data mining techniques in later studies to identify repeated patterns of a few teacher moves, as well as students moves, and identified in which modes they tended to occur (D'Mello, Olney, & Person, 2010; Lehman, D'Mello, Cade, & Person, 2012). So their overall analysis of expert human teaching behaviour could be expressed at three levels: mode, dialogue move, and move.

This work was used as the basis of a tutor, GURU, that embodied the above expert teacher behaviour, plus others from the literature such as the INSPIRE framework of Lepper and Woolverton (2002), to provide

conversational teaching in Biology (with typed input from the student and spoken output from the pedagogical agent in GURU). An evaluation of GURU compared standard classroom teaching *vs* classroom teaching augmented by small group sessions with a human (non-expert) teacher *vs* classroom teaching augmented by sessions with GURU. They researchers found that the two augmented modes of teaching produced similar learning gains, with both being better than classroom teaching alone (Olney et al., 2012).

In addition to observing expert teachers dealing with domain content and motivational issues, recent work has included studies of the attention that teachers give to supporting students' learning processes. For example, Yeager and Dweck (2012) looked at ways that teachers help students to develop "resilience" in the face of inevitable learning setbacks. This is an example of a meta-motivation skill (or learning how to learn) in line with the work of Lepper and Woolverton (2002) mentioned earlier. In a similar vein, Maehr (2012) argues that a basic goal of teaching should be to encourage a "Continuing Personal Investment in Learning: Motivation As an Instructional Outcome."

New tools are being brought to bear on the annotated records of teaching episodes. Porayska-Pomsta and Mellish (2013) modeled best practice in human tutor feedback on the correctness or otherwise of a student's answer to a problem, taking in both contextual factors such as the amount of time left in the lesson, the student's aptitude and the difficulty of the material, while also respecting the student's need for autonomy and approval. Here best practice was determined by what an experienced teacher would advise. But it is also important to have some method of evaluating which methods are empirically effective, not least as methods that might work for human teachers might not work for machine teachers. For example, Boyer and her colleagues (Boyer et al., 2011) used a machine-learning technique to compare two human tutors' sets of dialogues. Each tutor utterance in each dialogue was annotated using about a dozen different categories such as, question, lukewarm content feedback, or negative feedback and so on (similar to the preparatory work on GURU, described above). Hidden Markov Models were used to model the underlying structure of the dialogues and infer differences in the fine-structure (or tactics) of the strategies and their impact on the effectiveness of two tutors.

In a similar fashion Chi and her colleagues (Chi, VanLehn, Litman, & Jordan, 2011) induced two tutorial tactics from annotated corpora: one (NormGain) designed specifically to assist learning and the other (invNormGain) not to assist learning, in the sense that it "enhanced those decisions that contribute less or even nothing to learning" (Page 88). A follow up evaluation showed that the groups exposed to NormGain learned better than other strategies that the literature had shown to be effective as well as doing better than groups exposed to invNormGain.

So the major change over the last decade or so has been the increased detail with which expert teacher behaviour has been analysed, as well as the more sophisticated use of natural language processing techniques to provide pedagogically, pragmatically and culturally appropriate questions and feedback.

Derived from learning theory

While both epistemological and reflective theories of learning retain their utility, there has been increasing emphasis on reflective issues, with much ongoing work in the area of tutoring metacognitive skills (for a comprehensive review, see e.g. Azevedo & Aleven, 2013).

A variation on reflective theories of learning has been the emergence of systems aimed to help learners learn by getting them to teach the material to be learned to someone else, typically a simulated fellow student such as Betty's Brain (Leelawong & Biswas, 2008). The issue of how much this helps a learner to learn has been explored in a series of studies using SimStudent (Matsuda et al., 2013). SimStudent is a simulated student that uses machine-learning techniques to build production rules expressing a skill such as solving a simple equation in algebra. It is a "knowledge-tracing" simulated student embodying the same way of representing skill as the cognitive tutors (Anderson et al., 1995). The researchers were able to identify factors such as the quality of the learners' explanations that both helped them to learn as well as helping SimStudent to learn well. But perhaps the largest change in pedagogy has been the increased focus on affective and motivational issues in learning and teaching (Calvo & D'Mello, 2011). So there is a greater understanding of the feelings that are associated with academic learning (Pekrun, 2011) and a

greater interest in detecting feelings (see e.g., Arroyo et al., 2009; Kaśka Porayska-Pomsta, Mavrikis, D'Mello, Conati, & Baker, 2013). From this has emerged ways of dealing with such states as boredom and confusion (see e.g., Baker, D'Mello, Rodrigo, & Graesser, 2010) and their surface manifestations such as gaming the system (Baker et al., 2008). Indeed rather than seeing learner confusion as a situation to be avoided, it should be harnessed as a driver of understanding. Thus Lehman and her colleagues set up a variation on Socratic tutoring in which there was an explicit contradiction and then helped the student to figure out the nature of it. Normally this happens when the student discovers a bug or is told that there is a bug in his or her answer, but in the case of the work of Lehman and her colleagues, the contradiction was deliberately introduced by the tutor, and then it was the system's role to support the student in dealing with it (Lehman et al., 2013). In addition to work directly on affect, there has also been progress in allied areas such as politeness (Kaśka Porayska-Pomsta, Mavrikis, & Pain, 2008), cultural norms (Johnson, 2010) and the differential feedback needs of different personality types (Dennis, Masthoff, Pain, & Mellish, 2011; Robison, McQuiggan, & Lester, 2010). Johnson's (2010) work is also interesting in that it exploits many of the techniques used in gaming systems to provide the learner with a sense of being embedded and so engaged within a situation that needs him or her to act.

Work on engagement and disengagement has also progressed. For example, Forbes-Riley and Litman (2013) identified 6 kinds of student disengagement based on the most likely cause of that disengagement: (i) hard, because the work was too hard, (ii) easy, because the work was too easy, (iii) presentation, because the work was confusingly presented, (iv) NLP-gaming, because the student was attempting to get the system to reveal the answer, (v) NLP-distracted, because the student attempted to compensate for the system's own weakness in understanding a previous answer, and (vi) done, because the student was "bored, tired and/or not interested in continuing".

An ongoing and important issue in the field is illustrated by the fact that Forbes-Riley and Litman (2013) had less success in differentiating the tutor's best response to these six causes, and so chose two generic responses: first "disengagement associated with a correct answer", to which the tutor responded with "productive disengagement feedback" with a review of progress and praise if things had been going well; second, "disengagement associated with an incorrect answer", to which the tutor's response was intended to "remediate the negative learning correlation and target learning improvement". This links back to one of the findings in the review by VanLehn (2011), namely that human tutors did not normally develop and make use of insight into the individual strengths and weaknesses of their human students.

Observing students

Direct observation of individual students has become a core activity, especially in studies involving detecting student affect and goals (see e.g., Baker, Rodrigo, & Xolocotzin, 2007). The focus on how differences in individual students' attitudes to learning affect outcomes has been explored by Hulleman, Durik, Schweigert, and Harackiewicz (2008). They "examined the antecedents (initial interest, achievement goals) and consequences (interest, performance) of task value judgments" in two groups of students and found that "initial interest and mastery goals predicted subsequent interest, and task values mediated these relationships. Performance-approach goals and utility value predicted actual performance". Their work suggests that the overall strategy for how best to help an individual student depends to a large extent on what that student brings to the learning situation, and thus that individual differences are important.

Over the last 14 years a major addition to the observation of individual classes of students has been the development of large scale datasets of logs of student work with both AIED and e-learning systems. These datasets have provided a fruitful source for data-mining techniques (Mavrikis, D'Mello, Porayska-Pomsta, Cocea, & Graesser, 2010) including the use of learning curves (Martin, Mitrovic, Koedinger, & Mathan, 2011) and have enabled the analysis of many students' use of the same tutoring system in different educational contexts (see e.g., Mathews & Mitrovic, 2008).

FUTURE NEEDS

There has been steady but not startling progress in the development of pedagogy within an AIED context. Thus we now understand expert teacher behaviour in greater detail and at a finer level of granularity than before, but there is still some way to go before we are able to implement motivationally sophisticated tactics to deal with both the transient and more deep-seated impediments to learning experienced by students. Perhaps the use of simulated students, such as SimStudent (Matsuda et al., 2013), but incorporating more cognitively and motivationally mature models of human learning might be a route to developing more effective teaching strategies. Clearly there is a bit of a 'chicken and egg' situation here in terms of the relative difficulty of understanding learning *vs* understanding teaching.

The natural language capabilities of tutors both for typed (Olney et al., 2012) and spoken input (Johnson, 2010) as well as output has increased, though much remains to be achieved. Much more data is now available and more sophisticated methods of analyzing that data are also now available. The interfaces to systems are more complex (see e.g., Johnson, 2010), and the penetration of systems into the mainstream of education is greater (see e.g., Pane et al., 2014; Mitrovic, 2012). As indicated above there have also been great strides in detecting the affective states of learners such as engaged and disengaged, but greater progress is still needed in the pedagogy of how best to react to such more finely differentiated learner states, so the need for modelling teaching remains as strong as ever.

General progress in the technologies for learning have also developed alongside those for AIED systems. For example, MOOCs hold out much promise in terms opening up education to much wider audiences, but their developers are discovering the need for adaptivity to individual students, not least to improve drop out rates, and thus should take account of the work already done on student modeling and pedagogy within AIED over recent decades (Liyanagunawardena, Adams, & Williams, 2013).

ACKNOWLEDGEMENTS

We thank the anonymous reviewers of this article for their many helpful comments and suggestions.

REFERENCES

- Akhras, F. N., & Self, J. A. (2000). System Intelligence in Constructivist Learning. International Journal of Artificial Intelligence in Education, 11(4), 344-376.
- Anderson, J. R. (1990). *The Adaptive Character of Thought*. Hillsdale, New Jersey: Lawrence Erlbaum Associates.
- Anderson, J. R., Corbett, A. T., Koedinger, K. R., & Pelletier, R. (1995). Cognitive Tutors: Lessons Learned. *The Journal of the Learning Sciences*, 4(2), 167-207.
- Arroyo, I., Beck, J. E., Woolf, B. P., Beal, C. R., & Schultz, K. (2000). Macroadapting Animalwatch to Gender and Cognitive Differences with Respect to Hint Interactivity and Symbolism Intelligent Tutoring Systems: 5th International Conference, ITS2000, Montreal (pp. 574-583). Berlin: Springer.
- 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). Amsterdam: IOS Press.
- Azevedo, R., & Aleven, V. (Eds.). (2013). International Handbook of Metacognition and Learning Technologies: Springer.
- Baker, R., Walonoski, J., Heffernan, N., Roll, I., Corbett, A., & Koedinger, K. (2008). Why Students Engage in "Gaming the System" Behaviours in Interactive Learning Environments. *Journal of Interactive Learning Research*, 19(2), 185-224.
- Baker, R. S. J. d., D'Mello, S. K., Rodrigo, M. M. T., & Graesser, A. C. (2010). Better to be frustrated than bored: The incidence, persistence, and impact of learners' cognitive–affective states during

interactions with three different computer-based learning environments International Journal of Human-Computer Studies, 68(4), 223-241.

- Baker, R. S. J. d., 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.
- Bloom, B. S. (1984). The 2 Sigma Problem: The Search for Methods of Group Instruction as Effective as One-to-One Tutoring. *Educational Researcher*, 13(6), 4-16.
- Boyer, K. E., Phillips, R., Ingram, A., Ha, E. Y., Wallis, M., Vouk, M., & Lester, J. (2011). Investigating the Relationship Between Dialogue Structure and Tutoring Effectiveness: A Hidden Markov Modeling Approach International Journal of Artificial Intelligence in Education, 21(1-2), 65-81.
- Calvo, R. A., & D'Mello, S. K. (Eds.). (2011). New Perspectives on Affect and Learning Technologies. New York: Springer.
- Carroll, J., & McKendree, J. (1987). Interface Design Issues For Advice-Giving Expert Systems. *Communications of the ACM*, 30(1), 14--31.
- Chan, T.-W., & Chou, C.-Y. (1997). Exploring the Design of Computer Supports for Reciprocal Tutoring. International Journal of Artificial Intelligence in Education, 8(1).
- Chi, M., VanLehn, K., Litman, D., & Jordan, P. (2011). An Evaluation of Pedagogical Tutorial Tactics for a Natural Language Tutoring System: A Reinforcement Learning Approach. *International Journal* of Artificial Intelligence in Education, 21(2), 83-113.
- Chi, M. T. H., Bassok, M., Lewis, M. W., Reimann, P., & Glaser, R. (1989). Self-explanations: How students study and use examples in learning to solve problems. *Cognitive Science*, 13(2), 145-182.
- Collins, A., Warnock, E. H., Aiello, N., & Miller, M. L. (1975). Reasoning from Incomplete Knowledge. In D. G. Bobrow & A. Collins (Eds.), *Representation and Understanding* (pp. 383-415). New York: Academic Press.
- D'Mello, S., Olney, A., & Person, N. (2010). Mining Collaborative Patterns in Tutorial Dialogues. Journal of Educational Data Mining, 2(1), 1-37.
- 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.
- Dennis, M., Masthoff, J., Pain, H., & Mellish, C. (2011). Does Self-Efficacy Matter When Generating Feedback? Artificial Intelligence in Education: 15th International Conference, AIED 2011. Auckland, New Zealand (Vol. 6738, pp. 444-446). Berlin: Springer.
- du Boulay, B. (2011). Towards a Motivationally-Intelligent Pedagogy: How should an intelligent tutor respond to the unmotivated or the demotivated? In R. A. Calvo & S. K. D'Mello (Eds.), *New Perspectives on Affect and Learning Technologies* (pp. 41-54). New York: Springer.
- du Boulay, B., Avramides, K., Luckin, R., Martinez-Miron, E., Rebolledo-Mendez, G., & Carr, A. (2010). Towards Systems That Care: A Conceptual Framework based on Motivation, Metacognition and Affect. *International Journal of Artificial Intelligence in Education*, 20(3), 197-229.
- du Boulay, B., & Luckin, R. (2001). Modelling human teaching tactics and strategies for tutoring systems. International Journal of Artificial Intelligence in Education, 12(3), 235-256.
- Entwistle, N. (1978). Knowledge Structures And Styles Of Learning: A Summary Of Pask's Recent Research. *British Journal of Educational Psychology*, 48(3), 255-265.
- Forbes-Riley, K., & Litman, D. (2013). When Does Disengagement Correlate with Performance in Spoken Dialog Computer Tutoring? *International Journal of Artificial Intelligence in Education*, 22(1-2), 39-58.
- Graesser, A. C., Person, N., Harter, D., & the Tutoring Research Group. (2000). *Teaching Tactics in AutoTutor*. Paper presented at the Modelling Human Teaching Tactics and Strategies: Workshop W1 at ITS'2000, Montreal.
- Hulleman, C. S., Durik, A. M., Schweigert, S. A., & Harackiewicz, J. M. (2008). Task Values, Achievement Goals, and Interest: An Integrative Analysis. *Journal of Educational Psychology*, 100(2), 398-416.
- Johnson, W. L. (2010). Serious Use of a Serious Game for Language Learning. International Journal of Artificial Intelligence in Education, 20(2), 175-195.

- Lajoie, S. P., Wiseman, J., & Faremo, S. (2000). Tutoring Strategies for Effective Instruction in Internal medicine. Paper presented at the Modelling Human Teaching Tactics and Strategies: Workshop W1 at ITS'2000, Montreal.
- 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., D'Mello, S., Cade, W., & Person, N. (2012). How do they do it? Investigating Dialogue Moves within Dialgue Modes in Expert Human Tutoring Intelligent Tutoring Systems: 11th International Conference, ITS 2012, Chania, Crete, Greece, June 14-18, 2012. Proceedings (pp. 557-562). Berlin: Springer.
- Lehman, B., D'Mello, S., Strain, A., Mills, C., Gross, M., Dobbins, A., . . . Graesser, A. C. (2013). Inducing and Tracking Confusion with Contradictions during Complex Learning. *International Journal of Artificial Intelligence in Education*, 22(1-2), 85-105.
- Leinhardt, G., & Greeno, J. G. (1991). The Cognitive Skill of Teaching. In P. Goodyear (Ed.), *Teaching Knowledge and Intelligent Tutoring* (pp. 233-268). Norwood, New Jersey: Ablex Publishing Corporation.
- Lepper, M. R., Aspinwall, L., Mumme, D., & Chabay, R. W. (1991). 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 Sixth Ontario Symposium in Social Psychology (pp. 217-237). Hillsdale, New Jersey: Lawrence Erlbaum Associates.
- Lepper, M. R., & Woolverton, M. (2002). The wisdom of practice: Lessons learned from the study of highly effective tutors. In J. M. Aronson (Ed.), *Improving Academic Achievement: Impact of Psychological Factors on Education* (pp. 135-158). New York: Academic Press.
- Liyanagunawardena, T. R., Adams, A. A., & Williams, S. A. (2013). MOOCs: A Systematic Study of the Published Literature 2008-2012. *International Review of Research in Open and Distance Learning*, 14(3), 202-227.
- Luckin, R. (2010). Re-designing Learning Contexts. London: Routledge.
- Luckin, R., & du Boulay, B. (1999). Ecolab: The development and evaluation of a Vygotskian design framework. *International Journal of Artificial Intelligence in Education*, 10(2), 198-220.
- Ma, W., Adesope, O. O., Nesbit, J. C., & Liu, Q. (2014). Intelligent Tutoring Systems and Learning Outcomes: A Meta-Analysis. *Journal of Educational Psychology, Advance online publication*.
- Maehr, M. L. (2012). *Encouraging a Continuing Personal Investment in Learning: Motivation As an Instructional Outcome*. Charlotte, NC, USA: Information Age Publishing.
- Martin, B., Mitrovic, A., Koedinger, K. R., & Mathan, S. (2011). Evaluating and improving adaptive educational systems with learning curves. User Modeling and User-Adapted Interaction, 21(3), 249-283.
- Mathews, M., & Mitrovic, A. (2008). Do Students Who See More Concepts in an ITS Learn More? Paper presented at the 1st International Conference on eData Mining, Montreal, Quebec.
- Matsuda, N., Yarzebinski, E., Keiser, V., Raizada, R., William, W. C., Stylianides, G. J., & Koedinger, K. R. (2013). Cognitive anatomy of tutor learning: Lessons learned with SimStudent. *Journal of Educational Psychology*, 105(4), 1152-1163.
- Mavrikis, M., D'Mello, S., Porayska-Pomsta, K., Cocea, M., & Graesser, A. (2010). Modeling Affect by Mining Students' Interactions within Learning Environments, . In C. Romero, S. Ventura, M. Pechenizkiy, & R. S. J. d. Baker (Eds.), *Handbook of Educational Data Mining*: Chapman & Hall/CRC.
- Mitrovic, A. (2012). Fifteen years of Constraint-Based Tutors: What we have achieved and where we are going. User Modeling and User-Adapted Interaction, 22(1-2), 39-72.
- Narciss, S., Sosnovsky, S., & Andres, E. (2014). Adapting Tutoring Feedback Strategies to Motivation. In C. Rensing, S. de Freitas, T. Ley, & P. J. Munoz-Merino (Eds.), Open Learning and Teaching in Educational Communities: 9th European Conference on Technology Enhanced Learning, EC-TEL 2014, Graz, Austria, September 16-19, 2014, Proceedings (pp. 288-301). Switzerland: Sringer International Publishing.
- Ohlsson, S. (1987). Some Principles of Intelligent Tutoring. In R. W. Lawler & M. Yazdani (Eds.), Learning Environments and Tutoring Systems: Learning Environments and Tutoring Systems (Vol. 1, pp. 203-237). Norwood, New Jersey: Ablex Publishing.
- Olney, A. M., D'Mello, S., Person, N., Cade, W., Hays, P., Williams, C., ... Graesser, A. (2012). Guru: A Computer Tutor That Models Expert Human Tutors *Intelligent Tutoring Systems: 11th*

International Conference, ITS 2012, Chania, Crete, Greece, June 14-18, 2012. Proceedings (pp. 256-261): Springer.

- Olney, A. M., Person, N. K., & Graesser, A. C. (2011). GURU: Designing a Conversational Expert Intelligent Tutoring System. In C. Boonthum-Denecke, P. M. McCarthy, & T. A. Lamkin (Eds.), Cross-Disciplinary Advances in Applied Natural Language Processing: Issues and Approaches (Vol. 11, pp. 156-171).
- Pane, J. F., Griffin, B. A., McCaffrey, D. F. & Karam, R. (2014). Effectiveness of Cognitive Tutor Algebra I at Scale. *Educational Evaluation and Policy Analysis*, 36(2), 127-144.
- Pask, G., Kallikourdis, D., & Scott, B. C. E. (1975). The representation of knowables. *International Journal of Man-Machine Studies*, 7(1), 15-134.
- Pekrun, R. (2011). Emotions as Drivers of Learning and Cognitive Development. In R. A. Calvo & S. K. D'Mello (Eds.), New Perspectives on Affect and Learning Technologies (Vol. 3). New York: Springer.
- Porayska-Pomsta, K., Mavrikis, M., D'Mello, S., Conati, C., & Baker, R. S. J. D. (2013). Knowledge Elicitation Methods for Affect Modelling in Education. *International Journal of Artificial Intelligence in Education*, 22(3), 107-140.
- 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(1-2), 125-173.
- Porayska-Pomsta, K., & Mellish, C. (2013). International Journal of Human-Computer Studies, 71(6), 703-724.
- Ridgway, J. (1988). Of Course ICAI Is Impossible Worse Though, It Might Be Seditious. In J. Self (Ed.), Artificial Intelligence and Human Learning (pp. 28-48). London: Chapman and Hall Computing.
- Robison, J., McQuiggan, S., & Lester, J. (2010). Developing Empirically Based Student Personality Profiles for Affective Feedback Models. In V. Aleven, J. Kay, & J. Mostow (Eds.), *Intelligent Tutoring Systems: 10th International Conference, ITS 2010, Pittsburgh PA.* (pp. 285-295). Berlin: Springer.
- Rutter, M., Maughan, B., Mortimore, P., & Ouston, J. (1979). *Fifteen Thousand Hours: Secondary Schools* and Their Effects on Children: Paul Chapman Publishing.
- Shute, V. J. (1993). A Macroadaptive Approach to Tutoring. *Journal of Artificial Intelligence in Education*, 4(1), 61-93.
- VanLehn, K. (2011). The Relative Effectiveness of Human Tutoring, Intelligent Tutoring Systems, and Other Tutoring Systems. *Educational psychologist*, 46(4), 197-221.
- VanLehn, K., Ohlsson, S., & Nason, R. (1994). Applications of Simulated Students. Journal of Artificial Intelligence in Education, 5(2), 135-175.
- Winne, P. H. (1997). Experimenting to Bootstrap Self-Regulated Learning. *Journal of Educational Psychology*, *89*(3), 397-410.
- Wood, D. J., Wood, H. A., & Middleton, D. J. (1978). An Experimental Evaluation of Four Face-to-Face Teaching Strategies. *International Journal of Behavioural Development*, 1, 131-147.
- Yeager, D. S., & Dweck, C. S. (2012). Mindsets That Promote Resilience: When Students Believe That Personal Characteristics Can Be Developed. *Educational psychologist*, 47(4), 302-314.