Commentary on Kurt VanLehn's "The behaviour of tutoring systems"

Benedict du Boulay

Publication date
01-01-2006

Published in
International Journal of Artificial Intelligence in Education

Commentary on Kurt VanLehn's "The behaviour of tutoring systems"

Article (Unspecified)

du Boulay, Benedict (2006) Commentary on Kurt VanLehn's "The behaviour of tutoring systems". International Journal of Artificial Intelligence in Education, 16 (3). pp. 267-270. ISSN 1560-4292

Commentary on Kurt VanLehn's "The Behaviour of Tutoring Systems"

Benedict du Boulay, *IDEAs Lab, HCT Research Group, University of Sussex, Brighton, UK*

The paper by Kurt VanLehn has been designed for two audiences: one who does not know very much about ITSs and the other who does. The paper has much to offer both audiences. For the former it gives an indication of the educational performance and complexity that ITSs are now able to offer in a way that focuses on the pedagogy and not on the underlying mechanisms. For the latter audience it offers a way of thinking about the nature of ITSs that helps to unify the field and demonstrates common themes across a range of research programmes. The paper is interesting and useful to read and makes many wise points about ITSs and their development.

This commentary on the paper is divided into two parts. The first deals with the basic ontology set out in VanLehn's paper and offers some comments on it. The second goes through the paper sequentially and makes some comments on individual sections and subsections (but not all of them).

THE BASIC ONTOLOGY

Intelligent Tutoring Systems

It is pleasing that the term "Intelligent Tutoring System" (ITS) is being applied to a wider class of systems than just the various Cognitive Tutors coming out of the work of Anderson, Koedinger and their colleagues. It is also pleasing to read the comments made about the nature of the relationship of the technology itself to the use made of that technology by instructors in the classroom. But we need to acknowledge that the analysis offered by VanLehn applies only to a certain subclass of ITSs. These are ones whose intelligence is devoted to maximising the chance that students learn how to solve a certain kind of multi-step problem in technical domains. Of course, such systems are central to the concept of ITSs and probably form the largest subset of all systems. But there is a range of other systems that also attempt to help students develop by deploying their intelligence in different ways. Let us offer three examples:

i. First are systems that are concerned to deal as much with motivational issues as cognitive issues, including the "frustration" mentioned by VanLehn. These systems may well come to rather different conclusions about dealing with individual steps and about task selection. For example, they may deliberately offer a task that the system believes the learner can easily solve so as to provide that learner with the experience of success, where the increase in self-confidence will prepare the learner for domain difficulties ahead (del Soldato & du Boulay, 1995).

ii. Second there are systems where the goal of the system is to focus on the metacognitive rather than the cognitive. Again these systems may come to different conclusions about steps and tasks than a tutor concerned just with knowledge components and learning events (Luckin & Hammerton, 2002).

iii. Finally there are systems which deploy their intelligence to develop and manage the collaboration and interaction between students in a more open-ended way: for example by "refereeing" the degree of contribution made by each student to a group solution. Such a system may concentrate on one learner to get them to contribute more, or on the rest of the group to get them to critique one student's contribution (Vizcaino, 2005).

In each of these cases, the two loop structure can still be imposed on the behaviour of the system but the nature of what counts as a step and a task may differ from the analysis offered.

The Two Loop Structure

There are problems, e.g. in learning programming, where the student is asked to design and produce a complex, multi-part structure (e.g. a program) but where the order of working on those parts does not need to correspond either to a beginning-to-end traversal or to a top-down decomposition. In these cases the learner may develop the solution by moving around the developing structure opportunistically (as an expert would), and so it may not be possible for the system to comment sensibly on an individual step until sufficient of the context for that step has been developed. So while the student is working step by step, the system may not be in a position to operate the inner loop quite in the way described.

While the two loop structure captures in a compelling way much of the behaviour of many ITSs, we are now seeing the emergence of systems that essentially have a three loop structure. The extra loop is an outer one and is used by the system to update itself over many lessons with respect to the probabilities of various events. VanLehn refers to the fact that "considerable data from human students is needed in order to set up the parameters in its model of student learning": but this process can be undertaken in parallel with the use of the system via a third outermost loop. Such systems are a kind of development of O'Shea's "self-improving teaching system" from the 1970s but nowadays use complex statistical techniques to try to figure out the probabilities underpinning the "utilities" of possible tutor moves, or of the mapping between tutor moves, student moves and hidden internal states of the learner (see e.g. Soller & Lesgold, 2003). That is, rather than having the designer do all this once and for all prior to the system being built, the system employs an outer loop to adapt itself dynamically between lessons or sets of lessons.

Terminology

The term "learning event" captures much of the meaning of what is intended, but perhaps the events being talked about are better described as problem-solving events which act as *learning opportunities*: opportunities that are often, but not always, grasped.

It may also be helpful to align the two loop structure more explicitly with two other terms in wide circulation, namely "model tracing" and "knowledge tracing". Now these terms refer to mechanisms more than to behaviours but they are widely understood and correspond to the inner and outer loops to some extent.

VanLehn makes some excellent comments about the unhelpful connotations associated with the term "tutor". The term can be used in a pedagogically neutral way to denote the person/system charged with helping the student develop, irrespective of the way in which they discharge that role. As VanLehn argues, it is the way that systems actually get used in their context which is the main issue. What might look like a one-to-one tutor in terms of its design can be used by an imaginative human instructor in a variety of ways that may differ greatly from the classic one-to-one style. It is arguable that the technology is "quite neutral" though: it will have certain affordances that make some kinds of use in context easier than others.

SOME POINTS, SECTION BY SECTION

The Outer Loop –Task Selection

As already indicated, the focus of the paper is on knowledge and learning events. The affective, motivational and metacognitive state of the student is only fleetingly referred to: and to some extent it appears to be the responsibility of the human instructors employing the ITS to deal with these issues. Perhaps there should be a fifth bullet point to the four common types of outer loops to cover these other dimensions of learning, see e.g. du Boulay et al. (in preparation).

The Inner Loop – Feedback

The discussion of feedback in terms of both its timing and kind is very helpful, but there are some missing issues.

First is an issue that tends to get underplayed in many systems and that is what to do when the student solves the problem *correctly*. This paper, like many others, focuses to some extent on that word "incorrect", but the learning opportunities associated with "correct" also need emphasis. Perhaps there are better ways to solve the problem, perhaps there are much worse ways to solve this problem that would be better for a slightly different class of problems and so on. Exploring the space of correct and incorrect solutions, or making judgements about the quality of the correct solution are just two possible ways forward from a correct solution.

Second is the issue of the timing and nature of feedback. Sometimes (in other work) a false dichotomy is established between immediate and delayed feedback. Clearly if someone is very busy on a task such that complex feedback would be disruptive, then that feedback should be delayed (unless the consequences would be dire). But delayed feedback should not necessarily be the same as immediate feedback that is delayed. By the time someone has finished a task, the context has changed and the feedback can and should take rather a different form e.g. with respect to the sequence of the points to be made, compared to the corresponding sequence of immediate feedback comments and hints. It may be that a point that would have been made early in an immediate feedback situation is better expressed late in delayed feedback. There is also the point that delayed feedback can look at the whole solution and not just the partial solution developed so far, and thus make more general points than could have been made at individual steps.

Finally, there are moments in learning when the student is likely to be much more receptive to learning than others. Now this is an argument for certain kinds of immediate feedback. But when someone has just finished giving a talk (say), it probably is not helpful to immediately point out all the shortcomings in their performance (even in private) as emotionally that is just the wrong moment to do it. Later when they are calmer may be a better time.

Assessment

One of the puzzles about human problem-solving is that sometimes what appears to be a complex problem can be solved more accurately than the sub-problems that comprise it. The classic cases revolve around algebra “word problems” where the contextual information provided by the situation described in the word problem leads to more correct answers than if the students were simply asked just to solve the bare equations that they need to derive in order to solve the overall problem. This underlines the difficulty of assessing not just what the student knows but their ability to use that knowledge in a wide variety of contexts.

CONCLUSIONS

The analysis of systems in terms of the two loop structure is a very helpful development. It will be interesting to see whether the three loop structure, suggested above, becomes a reality. It will also be interesting to see the analysis extended to systems whose agenda is broader than simply the student’s domain level understanding and skill but also, to paraphrase Lepper and his colleagues (1990), to mimic expert human teachers who 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.”(p. 219).

REFERENCES

del Soldato, T., & du Boulay, B. (1995). Implementation of motivational tactics in tutoring systems. *International Journal of Artificial Intelligence in Education,* 6, 337-378.

du Boulay, B., Luckin, R., Martínez-Mirón, E., & Rebolledo Méndez, G. (Submitted). Motivationally Intelligent Educational Systems (submitted to *User Modeling and User-Adapted Interaction*).

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, NJ: Lawrence Erlbaum Associates.

Luckin, R., & Hammerton, L. (2002). Getting to know me: helping learners understand their own learning needs through Metacognitive scaffolding. In S. A. Cerri, G. Gouardères & F. Paraguaçu (Eds.) *Intelligent Tutoring Systems: 6th International Conference, ITS 2002*. Berlin: Springer.

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* (pp. 253-268). Amsterdam: IOS Press.

Vizcaino, A. (2005). A Simulated Student Can Improve Collaborative Learning. *International Journal of Artificial Intelligence in Education,* 15(1), 3-40.