Intelligent Tutoring Systems

- *One-to-one human tutoring has long been thought to be the most effective approach to teaching and learning (since at least Aristotle's tutoring of Alexander the Great!) Unfortunately, one-to-one tutoring is untenable for all students. Not only will there never be enough human tutors; it would also never be affordable.*[1]

So-called Intelligent Tutoring Systems (or ITS) are among the most common applications of AI in education (in any case, as we have seen, they have probably been around the longest). Generally speaking, ITS provide step-by-step tutorials, individualised for each student, through topics in well-defined structured subjects such as mathematics or physics. Drawing on expert-knowledge about the subject and about pedagogy, and in response to the individual student's misconceptions and successes, the system determines an optimal step-by-step pathway through the learning materials and activities. As the student proceeds, the system automatically adjusts the level of difficulty and provides hints or guidance, all of which aim to ensure that the student is able to learn the given topic effectively.

ITS come in many shapes although typically they involve several AI models, an approach that we will unpack here. As we saw in our earlier discussion of AI technologies, AI models are highly-simplified computational representations of specific knowledge about the real world (just like a model car is a simplified representation of a real car). The models used by ITS represent knowledge specific to teaching and learning: typically, knowledge about the topic to be learned is represented in what is known as a *domain* model, knowledge about effective approaches to teaching is represented in a *pedagogical* model, and knowledge about the student is represented in a *learner* model. The ITS algorithm draws on these three models in order to adapt

[1] Luckin et al., 'Intelligence Unleashed. An Argument for AI in Education.', 24.

a sequence of learning activities for each individual student. A fourth model found in some ITS is the *open learner* model, to which we will return later.

The domain model

A *domain model* represents knowledge about the subject that the ITS aims to help the students learn (much like the subject knowledge in a standard, non-educational, expert system). This might, for example, be knowledge about mathematical procedures, genetic inheritance or the causes of World War I. In fact, over the years, mathematics for primary and secondary school students has dominated ITS (mathematics, along with physics and computer science, are AIED's low-hanging fruits because they are, at least at a basic level, well-structured and clearly defined).

The pedagogy model

The ITS pedagogy model represents knowledge about effective approaches to teaching and learning that have been elicited from teaching experts and from research in the learning sciences (although it should be acknowledged that some ITS developers appear to assume that expertise in pedagogy is unnecessary). Pedagogical knowledge that has been represented in many ITS include knowledge of instructional approaches[2], the zone of proximal development[3], interleaved practice[4], cognitive overload[5], and formative feedback[6]. For example, a pedagogical model that implements Vygotsky's zone of proximal development will ensure that activities provided by the

---

[2] Carl Bereiter and Marlene Scardamalia, 'Intentional Learning as a Goal of Instruction', *Knowing, Learning, and Instruction: Essays in Honor of Robert Glaser*, 1989, 361–392.

[3] L. S. Vygotsky, *Mind in Society: Development of Higher Psychological Processes* (Cambridge, MA: Harvard University Press, 1978).

[4] Doug Rohrer and Kelli Taylor, 'The Shuffling of Mathematics Problems Improves Learning', *Instructional Science* 35, no. 6 (6 October 2007): 481–98, https://doi.org/10.1007/s11251-007-9015-8.

[5] R.E. Mayer and R. Moreno, 'Nine Ways to Reduce Cognitive Load in Multimedia Learning', *Educational Psychologist* 38, no. 1 (2003): 43–52.

[6] V. J. Shute, 'Focus on Formative Feedback', *Review of Educational Research* 78, no. 1 (1 March 2008): 153–89, https://doi.org/10.3102/0034654307313795.

system to the student are neither too easy nor too challenging, while one that implements formative feedback will ensure that feedback is provided to the student whenever it might support the student's learning.

The learner model

As we have seen, some CAI effectively (although usually by another name) implemented versions of both domain and pedagogical models: knowledge of what was to be learned and knowledge of how to teach what was to be learned (for example, using linear or branching programmed instruction). However, what distinguishes AI-driven ITS is that, as foreshadowed by Pask's SAKI, they also include a learner model: "a representation of the hypothesized knowledge state of the student"[7]. In fact, many ITS incorporate a wide range of knowledge about the student – such as their interactions, material that has challenged the student, their misconceptions, and their emotional states while using the system – all of which can be used to inform what is being taught and how, together with what support needs to be provided and when. In fact, most ITS go much further. The knowledge stored about the individual student is augmented with knowledge of all the students who have used the system so far, from which the system machine learns in order to predict which pedagogical approach and which domain knowledge is appropriate for any particular student at any specific stage of their learning. It is the learner model that enables ITS to be adaptive, and the machine learning that makes this adaptivity especially powerful.

A typical ITS architecture

Figure 1 shows how the domain, pedagogy and learner models might be connected in a typical ITS.

[7] Self, 'Student Models in Computer-Aided Instruction'.

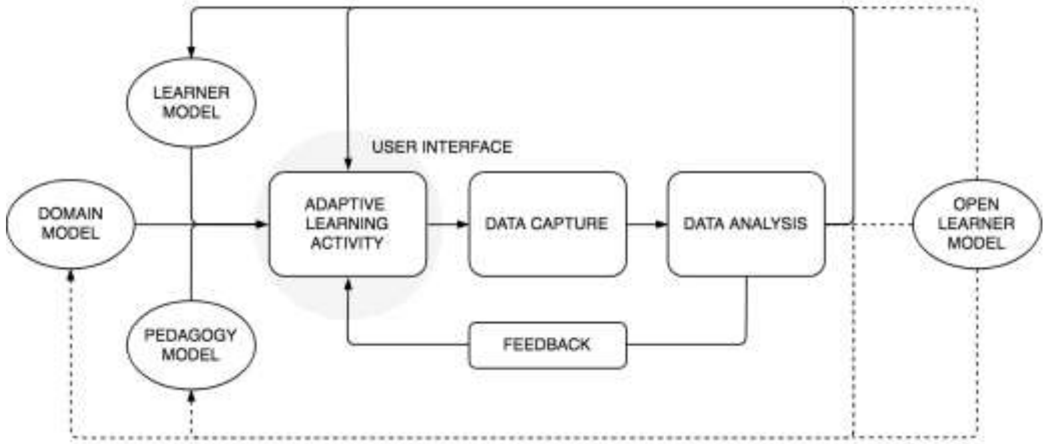

Figure 1. Flowchart representing a typical ITS architecture, including the pedagogy, domain, learner and open learner models.

In this exemplar architecture, the ITS algorithm draws on the domain, pedagogy and learner models to determine what adaptive learning activity should be presented to the individual student and how it should be adapted to that student's needs and capabilities. For example, in a mathematics ITS, the DOMAIN MODEL might contain knowledge about quadratic equations, the PEDAGOGY MODEL might contain knowledge of an effective way to teach quadratic equations, and the LEARNER MODEL might contain knowledge about the student's experience learning about quadratic equations in this ITS (for example, a misconception that they exhibited, or the fact that this topic caused them some anxiety). The learner model will also contain knowledge of all students who have ever used this ITS to learn about quadratic equations.

Drawing all of this together, the ITS algorithm will determine what ADAPTIVE LEARNING ACTIVITY to present to the student in the USER INTERFACE – in other words, which specific aspect of quadratic equations to deliver (perhaps factorising or completing the square) and what approach to use to best help the student learn about those aspects of quadratic equations (perhaps some instructional text, an image or video, or an interleaved practice activity), all of which is also dependent on the learner model (knowledge of the individual's and all students' experience of learning quadratic equations in this ITS).

While the student engages with the adaptive learning activity selected by the system, DATA CAPTURE involves the system capturing thousands of data points representing each individual interaction (what is clicked on the screen and what is typed, and possibly even how rapidly they move the mouse around the screen), together with the student's achievements (which tasks they have answered correctly or partially) and any misconceptions that they have demonstrated. Some advanced ITS also capture other data such as the student's speech and an inference of their affective (emotional) state.

The next step involves DATA ANALYSIS, in which all of the captured data is automatically examined, possibly using machine learning or Bayesian network techniques, both to provide the student with individualised formative FEEDBACK (to support their learning according to their individual needs) and to update the learner model (to inform the system's decision of which adaptive learning activity to deliver next, and to contribute to the model of all students). The analysis might also, in some circumstances, update the pedagogy model (identifying which approaches have been shown to support student learning most and least effectively, in particular circumstances) and the domain model (perhaps with previously unknown misconceptions that have become apparent from the student interactions).

Over time, this ITS cycle – (a) drawing on the domain, pedagogy and learner models, (b) delivering adaptive learning activities, (c) data capture, (d) data analysis, and (e) updating the models – means that each individual student will experience their own unique personalised learning pathway through the available learning activities. If their interactions suggest that they find quadratic equations particularly challenging, perhaps they will spend more time engaging with multiple relevant activities; whereas if their interactions suggest otherwise, perhaps they made few errors along the way, they will work through fewer activities for this topic and will more quickly move onto another topic deemed to be more appropriate for their particular needs.

Finally, as also shown in *Figure 1*, a few ITS also feature a fourth model, known as an *open learner model*[8]. Open learner models aim to make visible or explicit, for both the students and their teachers to inspect, both the teaching and learning that has taken place and the decisions that have been taken by the system. This is especially important if the system uses a neural network approach, in which as noted earlier it can be otherwise difficult to decipher how a decision has been made. The *open learner model* enables learners to monitor their achievements and personal challenges, supporting their metacognition, and enables teachers to better understand each individual learner's learning (their approach, any misconceptions and their learning trajectories) in the context of the whole class.

Evaluating ITS

Over the years there have been countless examples of ITS, many of which have been evaluated in schools or universities. Usually these evaluations have focused on learning gains, comparing one or other ITS with traditional teaching methods, such as whole-class or one-to-one teaching by a human teacher, or with CAI systems. In fact, as detailed by du Boulay and colleagues, there have also now been several meta-reviews[9] (i.e. review papers that aim to draw some general

---

[8] Vania Dimitrova, Gord Mccalla, and Susan Bull, 'Preface:" Open Learner Models: Future Research Directions" Special Issue of the IJAIED (Part 2).', *International Journal of Artificial Intelligence in Education*, 2007, http://psycnet.apa.org/psycinfo/2007-13116-001.

[9] Kurt VanLehn, 'The Relative Effectiveness of Human Tutoring, Intelligent Tutoring Systems, and Other Tutoring Systems', *Educational Psychologist* 46, no. 4 (2011): 197–221, https://doi.org/10.1080/00461520.2011.611369; Wenting Ma et al., 'Intelligent Tutoring Systems and Learning Outcomes: A Meta-Analysis.', *Journal of Educational Psychology* 106, no. 4 (2014): 901; John C. Nesbit et al., 'How Effective Are Intelligent Tutoring Systems in Computer Science Education?', in *Advanced Learning Technologies (ICALT), 2014 IEEE 14th International Conference On* (IEEE, 2014), 99–103, http://ieeexplore.ieee.org/abstract/document/6901409/; James A. Kulik and J. D. Fletcher, 'Effectiveness of Intelligent Tutoring Systems A Meta-Analytic Review', *Review of Educational Research*, 17 April 2015, 0034654315581420, https://doi.org/10.3102/0034654315581420; Saiying Steenbergen-Hu and Harris Cooper, *A Meta-Analysis of the Effectiveness of Intelligent Tutoring Systems on K–12 Students' Mathematical Learning.* (American Psychological Association, 2013), http://psycnet.apa.org/journals/edu/105/4/970/; Saiying Steenbergen-Hu and Harris Cooper, *A Meta-Analysis of the Effectiveness of Intelligent Tutoring Systems on College Students' Academic Learning.* (American Psychological Association, 2014), http://psycnet.apa.org/journals/edu/106/2/331/.

conclusions by combining several individual evaluations). Pooling the outcomes of the meta-studies suggests that ITS have yet to achieve parity with one-to-one teaching (when combined, the meta-reviews show an average small negative effect size of -0.19[10]). However, for ITS that are compared with whole-class teaching, the outcomes of the meta-reviews have been very positive. They show a weighted average effect size of 0.47[11] (in educational intervention research, effect sizes above 0.4 are thought to be "worth having"[12]), with one meta-analysis noting: "Developers of ITSs long ago set out to improve on the success of CAI tutoring and to match the success of human tutoring. Our results suggest that ITS developers have already met both of these goals."[13].

As we mentioned at the start of this section, generally speaking ITS focus on well-defined domains such as mathematics or physics. However, it is also worth noting that over recent years ITS for ill-defined problems (such as legal argumentation, intercultural skills acquisition, and dispute resolution) have also been explored[14]. One reason for the relatively low-levels of interest in ITS for ill-defined domains probably stems from the fact that ill-defined problems often require students to apply cognitively complex skills, while the contexts can be uncertain and dynamic, all of which can be challenge to model in traditional ITS. The relative lack of structure also makes it difficult to scaffold effective learning pathways without artificial constraint, to provide appropriate feedback, and to evaluate what learning is actually happening. ITS in ill-

[10] Although one meta-review did find that ITS were "just as effective as adult, one-to-one tutoring", VanLehn, 'The Relative Effectiveness of Human Tutoring, Intelligent Tutoring Systems, and Other Tutoring Systems', 214.

[11] The effect size measures how far the mean of the experimental group is from the mean of the control group measured in terms of the standard deviation of the control group scores.

[12] John Hattie, *Visible Learning*, 1st ed. (Routledge, 2008).

[13] Kulik and Fletcher, 'Effectiveness of Intelligent Tutoring Systems A Meta-Analytic Review', 67.

[14] Collin Lynch et al., 'Defining Ill-Defined Domains; a Literature Survey', in *Proceedings of the Workshop on Intelligent Tutoring Systems for Ill-Defined Domains at the 8th International Conference on Intelligent Tutoring Systems*, 2006, 1–10, http://people.cs.pitt.edu/~collinl/Papers/Ill-DefinedProceedings.pdf#page=7; Beverly Woolf, 'Social and Caring Tutors ITS 2010 Keynote Address', 2010, 12–14; H. Chad Lane et al., 'Intelligent Tutoring for Interpersonal and Intercultural Skills', 2007.

defined domains can also require additional pedagogical approaches, such as non-didactic Socratic dialogue, collaborative activities, or exploratory learning (which we look at in more detail later).

Mathia

We conclude our brief introduction to ITS by highlighting some prominent examples, focusing on Carnegie Learning's *Mathia*, Worcester Polytechnic Institute's *Assistments* and Knewton's *alta*.

Building on research at Carnegie Mellon University, *Mathia* (previously known as *Cognitive Tutor*) aims to mirror a human tutor by delivering AI-driven personalised mathematics instruction for K12 students. As the students work through carefully structured mathematics tasks, the system acts as a personal coach, monitoring their progress (their successes and misconceptions) and directing them along individualised learning pathways. It also provides automatic feedback that aims to explain not just why the individual student got something wrong but also how they might get it right. Interestingly, Carnegie Learning argues that *Mathia* is most effective when it is used as part of a blended learning approach (i.e. they acknowledge that, on its own, it is insufficient), which includes the use of both print and digital resources, and involves students learning collaboratively in groups as well as individually.

Assistments

Our second instructional ITS example is *Assistments*[15], which overall uses a similar approach to *Mathia*. However, *Assistments* also aims to address a key issue for ITS, that they by definition lead to students progressing at different rates, meaning that in any one classroom the students can be at increasingly diverging levels of attainment (potentially making the classroom teacher's job more, rather than less, challenging). Accordingly, Assistments is designed to help students catch up in the evenings, working independently at home, so that in the classroom everyone's

[15] https://www.assistments.org

progress remains roughly aligned. Both *Mathia*[16] and *Assistments*[17] have strong, although not definitive[18], evidence for their effectiveness.

alta

Our third example ITS is Knewton's *alta*[19], which is unusual in focusing on adaptive learning for Higher Education students across a range of subjects, including college-level mathematics, economics, chemistry and statistics. Nonetheless, like most ITS, *alta* aims to function like a 1:1 tutor, with personalised step-by-step instruction, assessment, feedback and just-in-time remediation while a student engages with an assignment. The *alta* approach clearly maps onto the typical ITS architecture outlined earlier. For each subject, it has a *domain model*, which uses open educational resources[20] (OER) and includes tutor-selectable learning objectives, together with a semantic network (or knowledge graph[21]) of relationships between the content and objectives. The domain models also include databases of relevant questions, together with data about the difficulty of those questions (based on how previous students have performed when responding to them). Alta's *pedagogy model* is based on Item Response Theory[22] (i.e., it works at the granularity of individual questions, taking into account both the question's difficulty and their representativeness of the underlying concepts), while adopting a mastery level approach (i.e.,

---

[16] John F. Pane et al., 'Continued Progress. Promising Evidence on Personalized Learning.' (Santa Monica, CA: RAND Corporation, 2015), https://www.rand.org/content/dam/rand/pubs/research_reports/RR1300/RR1365/RAND_RR1365.pdf.

[17] J. Roschelle et al., 'How Big Is That? Reporting the Effect Size and Cost of ASSISTments in the Maine Homework Efficacy Study' (Menlo Park, CA: SRI International, 2017).

[18] Holmes et al., 'Technology-Enhanced Personalised Learning. Untangling the Evidence.'

[19] https://www.knewtonalta.com

[20] Jan Hylén, 'Open Educational Resources: Opportunities and Challenges', *Proceedings of Open Education*, 2006, 49–63.

[21] Heiko Paulheim, 'Knowledge Graph Refinement: A Survey of Approaches and Evaluation Methods', ed. Philipp Cimiano, *Semantic Web* 8, no. 3 (6 December 2016): 489–508, https://doi.org/10.3233/SW-160218.

[22] Susan E. Embretson and Steven P. Reise, *Item Response Theory* (Psychology Press, 2013).

students do not move onto new learning objectives until they have achieved mastery of earlier learning objectives). In particular, the model assumes that if a student masters one of two learning objectives that are (according to the domain model's knowledge graph) related, there is a good chance that they have also mastered the other one. Meanwhile, *alta*'s *learner model* represents a student's level of mastery in terms of the learning objectives at any given point in time. This is based on the observed history of the individual student's interactions, and of all student interactions, including which questions the students have answered correctly and incorrectly, giving more weight to an individual student's most recent responses. Finally, the ongoing *data analysis* is based on a Bayesian modelling framework.