International Journal of Artificial Intelligence and EvolveTech
ISSN: 0000-0000
Volume 1, Issue 1 | PP 9-17 | www.ijaiet.com |

Cognitive Computing for Adaptive Learning Environments: A Review of Intelligent Tutoring Systems in Higher Education

Dr Arpit Jain
Department of Computer Science and Engineering, Koneru Lakshmaiah Educational Foundation,Vaddeswaram, Vijayawada, Andhra Pradesh, 522502

*Corresponding author
Dr Arpit Jain
Department of Computer Science and Engineering, Koneru Lakshmaiah Educational Foundation,Vaddeswaram, Vijayawada, Andhra Pradesh, 522502, India Emaill: arpitjain@kluniversity.in

Received : 21-06-2026| Revised : 15-08-2026| Accepted : 26-08-2026 | Published : 31-08-2026

Abstract

The need for high scalability of personalised instruction has been a constant demand in higher education, which is often not met by traditional lectures. The path to adaptive learning environments that customize instruction to the individual student lies with cognitive computing systems that sense, reason, learn, and adapt through machine learning, natural-language processing, and knowledge representation. The most developed realization of this vision is the intelligent tutoring system (ITS), which integrates a model of the subject, a model of the learner and a model of pedagogy to help students solve problems with individualized hints and feedback. This paper reviews the cognitive computing for adaptive learning with the intelligent tutoring in higher education background. It brings together the field along four-axis: the four parts of the ITS architecture, the adaptive loop that connects them; the cognitive-computing techniques that enable adaptivity, including Bayesian and deep knowledge tracing, content adaptation and sequencing, natural-language dialogue and automated feedback, and affect and engagement detection; the applications and evidence of effectiveness in mathematics, programming, science, language, and writing instruction; and the open challenges of scalability, data privacy, generalisability, equity, and evaluation. The literature reviewed indicates that there is consistently evidence, based on meta-analysis, that intelligent tutoring is close to the effectiveness of human one-to-one tutoring and significantly more effective than conventional instruction, provided implementation is robust and fidelity is high, and provided measures are appropriate. The single central thesis is that the utility of cognitive computing in education is not any one specific technique, rather it is the precision of the learner model upon which the teaching changes are based; and that the potential contribution of this utility in higher education will rely on integration, assessment, and equity as much as algorithmic progress.

Keywords: cognitive computing; intelligent tutoring systems; adaptive learning; higher education; knowledge tracing; learner modelling; educational technology

1. Introduction

One-to-one tutoring is a long standing education finding that results in dramatically improved outcomes as compared to traditional classroom instruction, and this effect is often described in the education literature as the "two-sigma problem" – how to leverage the power of individual human tutoring to every student at scale. In higher education, it is seen vividly: large class sizes, a variety of preparation, and time limited for staff, leave individual attention limited where it is most wanted [1]. Personalisation at scale is a potential computational path provided by the cognitive computing class of systems, that sense, reason, learn and adapt through machine learning, natural-language processing, and knowledge representation.

The most developed form of this concept is the intelligent tutoring system (ITS). An ITS is based upon the theories of artificial intelligence and cognitive science, and goes beyond the early computer-assisted instruction with its single, fixed branching to reason about what a specific student knows and then adapt the guidance to that specific student [2], [3]. In contemporary systems, the learner's developing knowledge is tracked and then, a new problem/hint is chosen that is appropriate for the learner's current level of knowledge, and there is natural-language dialogue between the learner and system, increasingly responding to a learner's motivation and affect. The capabilities make a static courseware into a dynamic learning environment that adapts constantly to the individual [4].

More and more people have begun to take an interest in these systems since the development of cognitive-computing techniques, and as higher education has rapidly turned to online and hybrid learning, allowing adaptive software to deliver at scale. However, literature exists all over the place in the fields of artificial intelligence, cognitive science, educational psychology, and learning analytics and is a combination of architectural description, algorithmic progress, and effectiveness studies, where the results are context and measure dependent. It is valuable to have a consolidated review that is focused on higher education that links the architecture, the underlying cognitive-computing techniques, the application evidence and the open challenges.

This paper gives such a review. It covers four questions: (RQ1) what is the architecture of an intelligent tutoring system and how do parts of it form an adaptive loop; (RQ2) what cognitive-computing techniques underpin adaptivity in such systems, and how can they be organised; (RQ3) in which higher-education domains are such systems applied, and what are their effectiveness and (RQ4) what are the challenges for their adoption. It is the result of the ITS architecture (Figure 1), a taxonomy of the cognitive-computing techniques that drive adaptation (Figure 2), a comparison of learner-modelling approaches (Table 3), a mapping of application domains and their evidence (Table 4), and a discussion of open challenges and design guidance.

2. Review Approach

2.1 Method. This is a structured narrative review, which was based on a predetermined protocol with research questions, literature sources, inclusion criteria and extraction and synthesis of the literature fixed before screening. The review includes architectural, algorithmic, and empirical work which report incommensurable outcomes, and where effectiveness is discussed, it is based on already available meta-analyses, not on the pooling of primary effects anew, as a narrative is appropriate for this review.

2.2 Scope and eligibility. The review addresses studies related to the development of intelligent tutoring system and adaptive learning environment based on the cognitive-computing techniques, focusing the analysis on higher education level but using the foundation literature in ITS where findings are applicable to the tertiary level. Those who are satisfied with non-adaptive courseware and/or just administrative educational technology are outside the scope. Table 1 shows all of the criteria.

2.3 Information sources and search strategy. The databases searched were the ACM Digital Library, IEEE Xplore, Scopus, Web of Science, and ERIC, which were supplemented by citation chasing of key reviews and meta-analyses. The search included three concept blocks (intelligent tutoring / adaptive learning; cognitive-computing techniques; and the higher-education setting) and Boolean operators optimized for each database, as shown in Table 2.

2.4 Selection and extraction. Records retrieved were de-duplicated and screened on title/abstract and then full text against the eligibility criteria, and the number of records at each stage should be documented in the final manuscript. A single template was developed for each of the retained works to summarize the system architecture, cognitive-computing technique(s), learner-modelling approach, application domain, and evaluation and reported outcomes.

2.5 Synthesis. The findings were themed based on the 4 research questions. The architecture is synthesized in Figure 1, the techniques in the taxonomy in Figure 2 and the comparison in Table 3, while the application domains and the evidence for them are mapped in Table 4.

Table 1. *Scope and eligibility criteria for the review.*

Dimension	Inclusion criteria	Exclusion criteria
Focus	ITS or adaptive learning driven by cognitive-computing / AI techniques	Non-adaptive courseware; administrative ed-tech
Setting	Higher education, or foundational ITS work that generalises to it	Settings with no relevance to tertiary education
Contribution	Architecture, technique, application, or review / meta-analysis	Non-technical opinion pieces
Type	Peer-reviewed papers, books, and well-cited preprints	Non-archival abstracts and posters
Language	Published in English	Full text unavailable in English
Window	Foundational to current ITS literature	Out of scope of the review

Table 2. *Search concept blocks and representative terms; syntax was adapted to each database.*

Concept block	Representative search terms (OR-combined within block; blocks AND-combined)
A · Tutoring / adaptive	"intelligent tutoring system" OR ITS OR "adaptive learning" OR "adaptive educational" OR "cognitive tutor"
B · Cognitive computing	"cognitive computing" OR "artificial intelligence" OR "machine learning" OR "knowledge tracing" OR "learner model" OR "student model" OR NLP
C · Higher education	"higher education" OR university OR "post-secondary" OR undergraduate OR tertiary OR college

3. Findings

3.1 Cognitive computing and adaptive learning. Learning is most effective when instruction is tailored to the learner's prior knowledge, and that can be automated when the system can keep an updated model of the learner and can perform actions on it [5]. Cognitive computing provides the ingredients that machine learning uses to make inferences about what a student knows based on their actions, knowledge representation to encode the content of the discipline and its prerequisites, natural-language processing to interpret and generate explanations (and reasoning to determine what to do next). All of these are utilized within the intelligent tutoring system, and it is within the ITS that cognitive computing has most clearly made an impact in education. It is explored in regard to its architecture, its techniques, its applications, and its limitations through the four research questions below.

This is a significant shift from the previous technological learning. First generation computer-assisted instruction was an approach that used an unalterable, standard decision tree in which the student was guided to the next screen depending on a single correct response. Cognitive computing supersedes these fixed-branches with an ever-changing inference: a probabilistic state of the learner that is updated at every moment, and that can adapt to new situations beyond those it was specifically designed for. It is this shift from the authored branching to the model-driven, learned branching that makes the learner model the central concern of the field and not just a side effect of the design.

3.2 RQ1 The intelligent tutoring system architecture. The classic ITS design separates the system into four interacting components that are illustrated in Figure 1. The domain model captures the concepts, procedures, relationship of prerequisites and, frequently, an executable expert with which to solve problems and produce correct steps. The learner model (or student model) keeps track of the learner's knowledge, skills and more importantly his/her affective state, which is constantly revised from the learner's feedback. The pedagogical model (or tutoring model) determines what to do next, which problem to ask, when to prompt and what to say next in terms of feedback. The interface provides instruction and records student interactions. These components create an adaptive loop with the learner interacting with the interface, the learner model being updated, the pedagogical model choosing a response based on the domain model, and the adapted content being fed back to the learner [6]. The learner model's accuracy is key, as all adaptive decisions depend on the accuracy of the learner model.

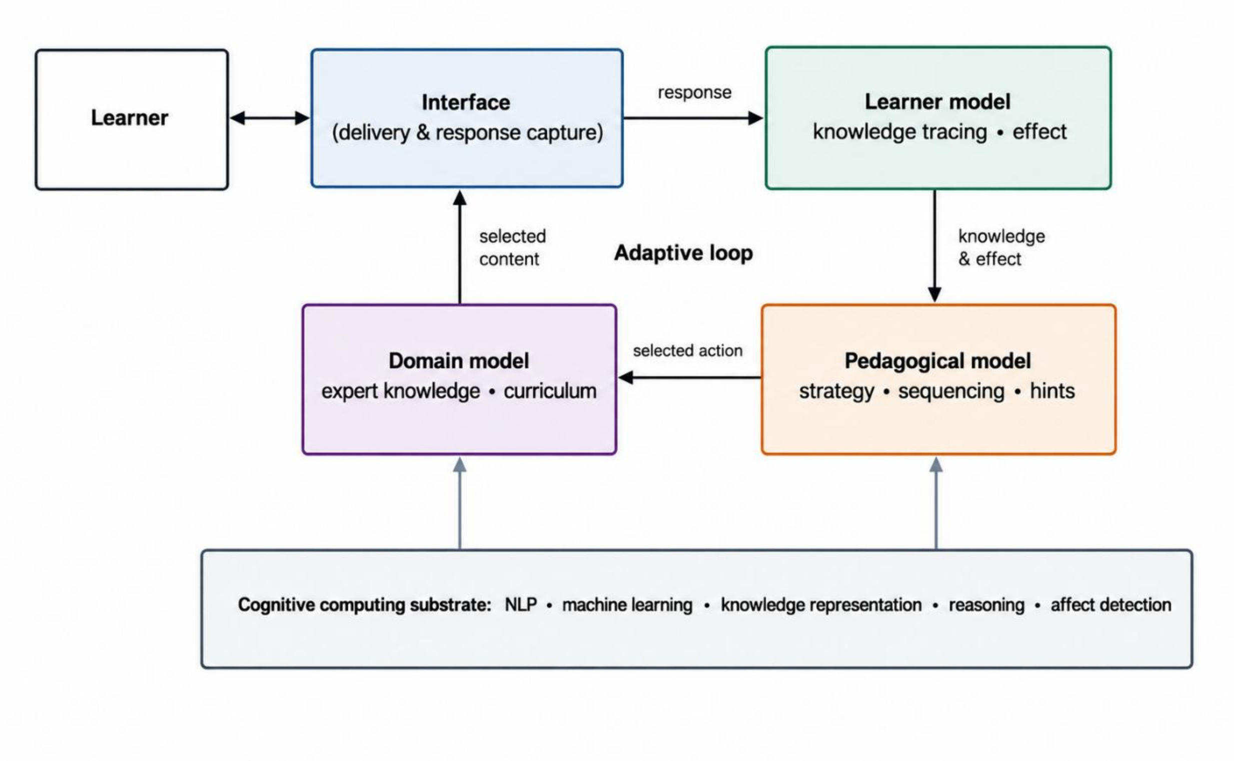

Figure 1. *The four-component ITS architecture and its cognitive-computing-driven adaptive loop.*

3.3 RQ2 Cognitive-computing techniques for adaptivity. There are four kinds of techniques involved in this loop, as are presented in Figure 2. Learner modelling is best developed. Bayesian knowledge tracing (BKT) views a skill as a latent binary variable, with the probability that a student has mastered the skill updated as they answer, a technique which has been used to power deployed tutors for decades [7]. In contrast to the hand-specified skills in BKT, deep knowledge tracing (DKT) uses a recurrent neural network to learn to predict performance directly from interaction sequences, and sometimes at the cost of interpretability [8]. Between the extremes are logistic and item-response models, as well as comparative overviews that stress the importance of choosing models based on interpretability and data requirements, in addition to accuracy [9], [10]. Once this is done, content adaptation acts on the learner model sequencing curricula, enforcing mastery criteria before moving to the next content, and in the recent work, employing reinforcement learning to optimise the sequence of activities. Conversational tutors that engage in mixed-initiative dialogues with students [11] are models of natural-language interaction that enable dialogue-based tutoring, automated feedback and the scoring of open responses. There is evidence that affective states are tightly coupled with learning [12], and affect and engagement techniques detect confusion, boredom, and frustration, and adjust accordingly. Table 3 compares the key aspects of the main techniques for modelling the learner, for deployment.

International Journal of Artificial Intelligence and EvolveTech
ISSN: 0000-0000
Volume 1, Issue 1 | PP 9-17 | www.ijaiet.com |

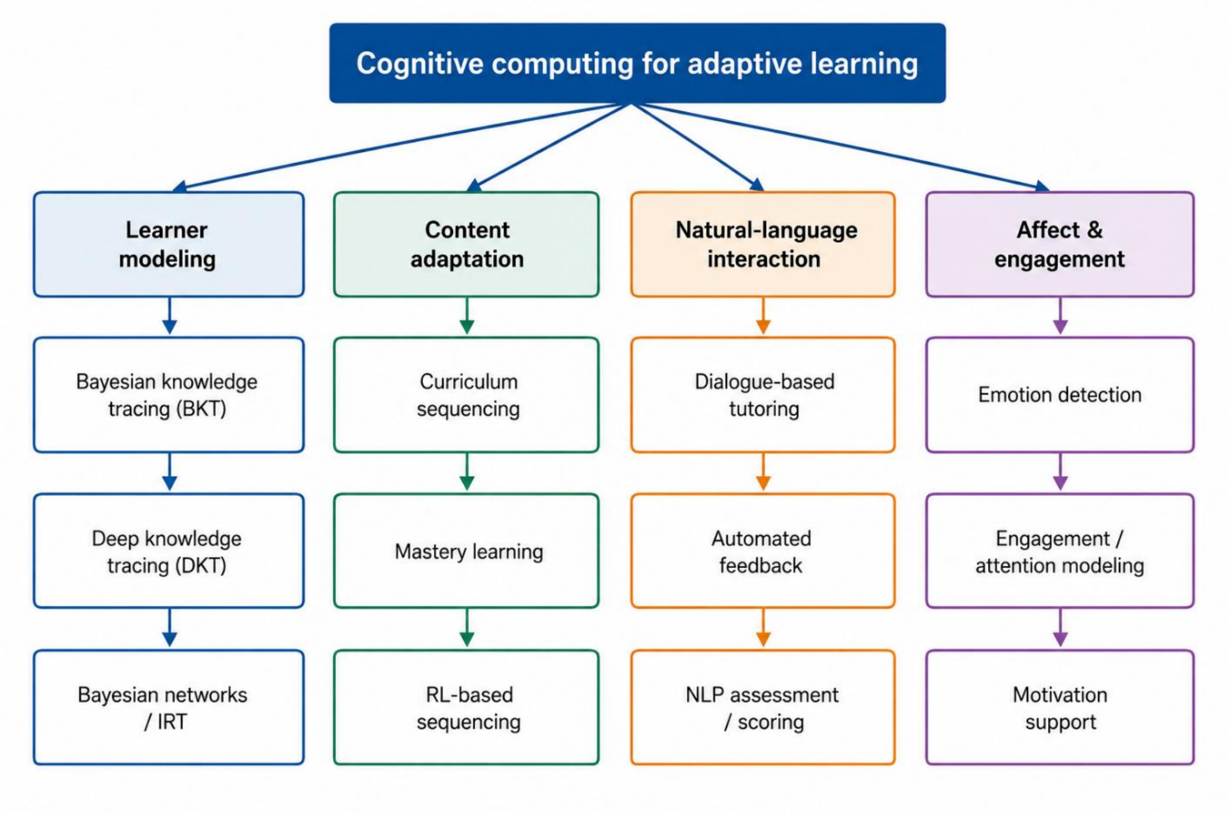

Figure 2. *A taxonomy of cognitive-computing techniques underpinning adaptive learning environments.*

Table 3. *Comparison of learner-modelling and adaptation techniques used in intelligent tutoring systems.*

Technique	What it models	Data needs	Interpretability	Key limitation
Bayesian knowledge tracing	Per-skill mastery probability	Modest; skill-tagged items	High	Assumes independent skills
Deep knowledge tracing	Performance from interaction sequences	Large interaction logs	Low	Opaque; data-hungry
Bayesian networks	Probabilistic knowledge structure	Moderate; expert structure	Medium	Structure hard to author
Logistic / IRT (e.g., PFA)	Skill difficulty & learning rate	Moderate	Medium–high	Limited temporal dynamics
Constraint / rule-based	Correctness against domain rules	Expert-authored rules	High	Authoring effort
RL-based sequencing	Optimal activity ordering	Large; reward signal	Low	Sample-inefficient; reward design

3.4 RQ3 Applications in higher education and their effectiveness. Intelligent tutoring has been used in the whole higher-education curriculum, as indicated in Table 4. The best-developed application area is mathematics and STEM problem solving, in which a cognitive tutor and/or model-tracing tutors lead the student step-by-step, and in which computer-based tutoring has been studied in large-scale real classrooms and observed to have a measurable impact on student learning [13], [14]. Programming tutors offer adaptive support for code, physics tutors reason about

International Journal of Artificial Intelligence and EvolveTech
ISSN: 0000-0000
Volume 1, Issue 1 | PP 9-17 | www.ijaiet.com |

problem-solving, language-learning systems employ natural-language processing to grammar and comprehension, and writing tutors give automated formative feedback, whose thoughtful design has been shown to be important for learning [15], [16]. The evidence is promising in these areas. The results of meta-analyses show that intelligent tutoring is significantly more effective than conventional instruction and is as effective as human one-to-one tutoring: A review of 50 controlled evaluations found intelligent tutoring positively impacted the students' learning by an average of approximately two-thirds of a standard deviation, compared with conventional instruction [17]; A different meta-analysis reached broadly similar conclusions [18], [1]. Through systematic reviews, researchers found that the use is extensive and also great variation in the rigor [19]. Importantly, measured effectiveness is very dependent on the alignment of the test and instruction, which is a strong determinant of effectiveness, and on the fidelity of application of the system; headline effect sizes should be interpreted with these caveats in mind.

Table 4. *Higher-education application domains, representative systems, and the reported evidence of effectiveness.*

Domain	Representative system / type	Cognitive-computing role	Reported effectiveness
Mathematics / STEM	Cognitive / model-tracing tutors	Knowledge tracing; step-level feedback	Well-evidenced classroom gains
Programming	Adaptive code tutors	Solution analysis; targeted hints	Positive; growing evidence
Physics / science	Problem-solving tutors	Expert model; dialogue	Positive in controlled studies
Language learning	Intelligent CALL	NLP for grammar & comprehension	Mixed to positive
Writing	Automated writing feedback	NLP scoring; formative feedback	Depends on feedback design
Clinical / professional	Case-based reasoning tutors	Domain reasoning; adaptivity	Emerging evidence

3.5 RQ4 Open challenges. Four challenges recur. Authoring and scalability: creating the domain and learner models for a new subject is labor-intensive, and finding ways to reduce this cost, such as using examples to author the domain or introduce models using data, are central concerns [16], [20]. Consent, security and appropriate use of fine-grained records of student behaviour raise questions of data privacy which are particularly salient in educational contexts. Generalisability and equity: models trained on one population, might not generalize to another, and adaptive systems can be the source of inequities if they do not perform equally across student populations. Evaluation: Effect sizes depend on measurement and implementation fidelity, and measures of outcomes that are self-reported or home-grown may overestimate benefits, making it critical to have rigorous, aligned evaluation. These problems, rather than algorithms, currently form the barrier to responsible growth in adaptive learning in higher education. Governance and design process, not just ability, is becoming the critical issue when addressing them: consent and data-handling practices are transparent and clear, subgroup-disaggregated evaluation to identify inequitable performance is present, and tools for authoring and audit are available that allow domain experts without deep technical expertise to address them together.

4. Discussion

4.1 Principal findings. Three conclusions emerge. First of all, the intelligent tutoring system is still the most successful implementation of cognitive computing in education and its 4-component system of components gives the field a long-lasting framework. Second, the effectiveness of intelligent tutoring is real and substantial across a large body of evaluations, and is comparable to human one-to-one tutoring and far greater than conventional instruction, though the size of the effect is highly dependent on measurement and implementation [17], [1]. Third, it is not a single algorithm, but the accuracy of the learner model that underlies adaptation: each and every hint, problem selection, and feedback decision is limited by the accuracy of the learner model on which it is based.

4.2 Interpretability vs. predictive power. There is a tension in the learning modelling literature, directly relevant to higher education. While deep methods like DKT offer more accurate predictions of performance, they also do so with models that are opaque and hard to explain, as opposed to the per-skill mastery estimates provided by BKT, which are transparent to instructors and students and can be inspected and corrected [8], [9]. In educational applications, where students and teachers gain insight into the reasoning behind a particular system's recommendations, the more interpretable model can be advantageous, albeit at a sacrifice of raw accuracy—a more general trade-off between accuracy and transparency in applied machine learning.

4.3 Implications for practice. The review indicates that for institutions, adaptive learning is not simply about selecting the best performing algorithm, but also around investing time and resources in the creation of the accurate learner model, alignment of assessment, and faithful implementation. Formative feedback should be carefully designed so that the quality of this feedback strongly conditions learning gains [15]; systems should be assessed on measures that are consistent with the goals of instruction, rather than on the "easy" measures; and the adoption of formative feedback should be regarded as an integration and change-management effort for existing platforms and curriculum, not as a mere drop-in tool.

4.4 The emerging role of large language models. One major advancement in recent times is the introduction of “large language models” into adaptive learning. Their command of natural language greatly reduces the long-standing limitation on dialogue-based tutoring and open-ended feedback, which was previously restricted to elaborate systems of authoring, and they are capable of producing explanations, questions and worked examples on the fly. They do not, however, on their own provide the disciplined learner model that is critical for effective adaptation, and they raise novel concerns regarding factual reliability, the need to deliver answers rather than to guide students in their struggle, and the interpretability questions posed by any other opaque model. The most promising way forward is not to abandon the ITS architecture in favor of a language model, but rather to take language models as a powerful interaction and content generation layer on top of the ITS, and to have an explicit learner model to control the adaptation of the system, and why it adapts.

4.5 Limitations of this review. This review is a narrative synthesis and editorial judgements were made regarding the selection and organisation of a wide literature that was cross-disciplinary. The search was limited to English-language literature. Many of the best pieces of evidence come from mathematics and from contexts where there is both secondary and tertiary education, so caution should be used when generalising to all higher-education sectors. The taxonomy and architecture are not meant to be complete enumerations. These limitations must be clearly indicated in the final manuscript.

5. Conclusion

Personalised instruction at scale is a viable option in higher education with the advent of cognitive computing, and intelligent tutoring systems are the most developed form of it. The review has synthesized the four components of the ITS, the adaptive loop connecting them, and categorized the techniques for affect detection, natural-language interaction, content sequencing, and adaptation knowledge tracing into a taxonomy of techniques for the former, while mapping application domains and associated meta-analytic evidence for the latter three; it has also identified challenges in authoring, privacy, generalisability, equity, and evaluation that now stand in the way of progress. The lesson is that the effectiveness of adaptive learning stems from the correctness of the learner model, not any single technique, and that the ability to reliably deliver benefits to students relies on the alignment of assessment, fine-tuned feedback design, equity, and institutionalization, as well as additional algorithmic developments. Sound learner modelling and responsible implementation are what make the difference between this cognitive computing providing the means and improving learning.

Declarations

Conflict of Interest

The authors have no conflict of interest to declare.

Funding Statement

This work was not supported by any specific grant from the public, commercial or not-for-profit sectors.

Ethical Approval

There were no human or animal subjects in this study and no need for ethical approval.

References

[1] K. VanLehn, "The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems," *Educational Psychologist*, vol. 46, no. 4, pp. 197–221, 2011. doi: 10.1080/00461520.2011.611369.

[2] H. S. Nwana, "Intelligent tutoring systems: An overview," *Artificial Intelligence Review*, vol. 4, no. 4, pp. 251–277, 1990. doi: 10.1007/BF00168958.

[3] J. R. Anderson, A. T. Corbett, K. R. Koedinger, and R. Pelletier, "Cognitive tutors: Lessons learned," *The Journal of the Learning Sciences*, vol. 4, no. 2, pp. 167–207, 1995. doi: 10.1207/s15327809jls0402_2.

[4] B. P. Woolf, *Building Intelligent Interactive Tutors: Student-Centered Strategies for Revolutionizing E-Learning*. Morgan Kaufmann, 2009.

[5] M. Vandewaetere, P. Desmet, and G. Clarebout, "The contribution of learner characteristics in the development of computer-based adaptive learning environments," *Computers in Human Behavior*, vol. 27, no. 1, pp. 118–130, 2011. doi: 10.1016/j.chb.2010.07.038.

[6] R. Nkambou, J. Bourdeau, and R. Mizoguchi, Eds., *Advances in Intelligent Tutoring Systems*. Springer, 2010. doi: 10.1007/978-3-642-14363-2.

[7] A. T. Corbett and J. R. Anderson, "Knowledge tracing: Modeling the acquisition of procedural knowledge," *User Modeling and User-Adapted Interaction*, vol. 4, no. 4, pp. 253–278, 1994. doi: 10.1007/BF01099821.

[8] C. Piech, J. Bassen, J. Huang, S. Ganguli, M. Sahami, L. J. Guibas, and J. Sohl-Dickstein, "Deep knowledge tracing," in *Advances in Neural Information Processing Systems 28*, pp. 505–513. Curran Associates, 2015.

[9] R. Pelánek, "Bayesian knowledge tracing, logistic models, and beyond: An overview of learner modeling techniques," *User Modeling and User-Adapted Interaction*, vol. 27, nos. 3–5, pp. 313–350, 2017. doi: 10.1007/s11257-017-9193-2.

[10] K. Chrysafiadi and M. Virvou, "Student modeling approaches: A literature review for the last decade," *Expert Systems with Applications*, vol. 40, no. 11, pp. 4715–4729, 2013. doi: 10.1016/j.eswa.2013.02.007.

[11] A. C. Graesser, P. Chipman, B. C. Haynes, and A. Olney, "AutoTutor: An intelligent tutoring system with mixed-initiative dialogue," *IEEE Transactions on Education*, vol. 48, no. 4, pp. 612–618, 2005. doi: 10.1109/TE.2005.856149.

[12] S. D'Mello and A. Graesser, "Dynamics of affective states during complex learning," *Learning and Instruction*, vol. 22, no. 2, pp. 145–157, 2012. doi: 10.1016/j.learninstruc.2011.10.001.

[13] K. R. Koedinger, J. R. Anderson, W. H. Hadley, and M. A. Mark, "Intelligent tutoring goes to school in the big city," *International Journal of Artificial Intelligence in Education*, vol. 8, pp. 30–43, 1997.

[14] J. F. Pane, B. A. Griffin, D. F. McCaffrey, and R. Karam, "Effectiveness of cognitive tutor Algebra I at scale," *Educational Evaluation and Policy Analysis*, vol. 36, no. 2, pp. 127–144, 2014. doi: 10.3102/0162373713507480.

[15] V. J. Shute, "Focus on formative feedback," *Review of Educational Research*, vol. 78, no. 1, pp. 153–189, 2008. doi: 10.3102/0034654307313795.

[16] V. Aleven, B. M. McLaren, J. Sewall, and K. R. Koedinger, "A new paradigm for intelligent tutoring systems: Example-tracing tutors," *International Journal of Artificial Intelligence in Education*, vol. 19, no. 2, pp. 105–154, 2009.

[17] J. A. Kulik and J. D. Fletcher, "Effectiveness of intelligent tutoring systems: A meta-analytic review," *Review of Educational Research*, vol. 86, no. 1, pp. 42–78, 2016. doi: 10.3102/0034654315581420.

[18] W. Ma, O. O. Adesope, J. C. Nesbit, and Q. Liu, "Intelligent tutoring systems and learning outcomes: A meta-analysis," *Journal of Educational Psychology*, vol. 106, no. 4, pp. 901–918, 2014. doi: 10.1037/a0037123.

[19] E. Mousavinasab, N. Zarifsanaiey, S. R. Niakan Kalhori, M. Rakhshan, L. Keikha, and M. Ghazi Saeedi, "Intelligent tutoring systems: A systematic review of characteristics, applications, and evaluation methods," *Interactive Learning Environments*, vol. 29, no. 1, pp. 142–163, 2021. doi: 10.1080/10494820.2018.1558257.

[20] R. S. J. d. Baker and K. Yacef, "The state of educational data mining in 2009: A review and future visions," *Journal of Educational Data Mining*, vol. 1, no. 1, pp. 3–17, 2009.