Foundational Principles and Design of a Hybrid Tutor

Andrew J. Hampton(✉) and Arthur C. Graesser

University of Memphis, Memphis, TN 38152, USA
andrew.hampton@memphis.edu

Abstract. This paper describes the concept of a *hybrid tutor* as a type of adaptive instructional system (AIS). A hybrid tutor is a confederation of several digital learning resources and human interactions so that the right resource is available to the learner at the right time. We discuss a method for combining several existing educational technologies into a unified platform that tracks progress on learning a subject matter across several constituent parts and offers recommendations on what to do next. There is a learning record store that keeps track of progress and enables intelligent recommendations at several levels: broad topics, specific knowledge components, material difficulty, and mode of instruction. The fine-grain adaptability allows the incorporation of several cognitive learning principles, such as multiple representations and modalities, mental model construction, item spacing, and support for self-regulated learning. The proposed web-based learning environment can function as a stand-alone instructional platform that is integrated into classrooms with topics assigned according a curriculum-based calendar or an adaptive learning environment that suggests learning activities that are generated by an intelligent recommender system. As a proof of concept, we developed *ElectronixTutor*, a hybrid tutor designed for introductory and intermediate electrical engineering education. This paper describes the rationale for its design and preliminary results.

Keywords: Adaptive instructional systems · Intelligent tutoring systems · Learning principles

1 Introduction

Adaptive instructional systems (AISs) are computer-based systems that guide learning experiences by tailoring instruction and recommendations based on the goals, needs, and preferences of each learner in the context of domain learning objectives [1]. These advanced computer learning environments help students master knowledge and skills by implementing algorithms that adapt to students and that are informed by scientific principles of learning [2]. Typically, this type of instruction focuses on one student at a time to be sensitive to individual differences relevant to the topic at hand or instruction generally. It is also possible to have an automated tutor or mentor interact with small teams of learners in collaborative learning and problem-solving environments [3, 4].

Many of these systems go far beyond the capabilities of conventional computer training systems. Adaptivity in conventional systems often consists of no more than

R. A. Sottilare and J. Schwarz (Eds.): HCII 2019, LNCS 11597, pp. 96–107, 2019.
https://doi.org/10.1007/978-3-030-22341-0_8

coarse-grained signal-response using primitive learning principles. For example, a learner may study static material (e.g., text), take a multiple-choice assessment, receive a score, and iterate through the same process until achieving a threshold performance. Progression through topics often follows a predetermined order. Advanced AISs can drastically improve upon this approach by implementing fine-grained adaptivity. This can include providing feedback within individual problems to work toward a correct answer or directing learners to a subsequent problem suited to their level of mastery (as determined by previous performance). This is known as two-loop adaptivity [5]. Intelligent tutoring systems, a subset of AISs, track detailed learner characteristics such as knowledge, skills, and other psychological attributes and apply computational models based on the combination of artificial intelligence and cognitive science [2, 6, 7].

1.1 Contributions of AISs

The evolution of cognitive learning principles and models of learning have produced a range of pedagogically advanced AIS environments. Several mature systems have demonstrated significant learning outcomes. Some examples cover well-defined subject matters such as algebra and geometry, including Cognitive Tutors [8–10] and ALEKS [11]. Other efforts in electronics (SHERLOCK [12], BEETLE-II [13]) and digital information technology (Digital Tutor [14]) also have successful use cases.

The inclusion of verbal interaction with conversational agents can scaffold more natural engagement with the subject matter [15, 16] and open doors to less well-defined domains. Conversational systems encourage learners to explain concepts in their own words and thereby engender reflection and reasoning. Mixed initiative dialogues allow learners to direct the conversation to personally relevant areas of the topic. Paralinguistic cues (e.g., pointing, facial expressions) increase realism and allow for visual reinforcement of information.

Conversational systems with two or more agents (e.g., a teacher agent and a student agent) allow multiple kinds of interactions and encourages greater social involvement [15, 17, 18]. Some examples of conversational AISs include AutoTutor [16, 19], Betty's Brain [20], Coach Mike [21], Crystal Island [22], and Tactical Language and Culture System [23], all with demonstrated advantages over conventional instructional techniques.

Analyses of the effectiveness of AISs (and more specifically intelligent tutoring systems) have demonstrated value added over more conventional approaches like classroom instruction or reading static materials. While the effect sizes vary substantially from $d = 0.05$ [24, 25] to an impressive $d = 1.08$ [26], most converge on relatively large values between $d = 0.40$ and $d = 0.80$ [5, 27]. Together, AISs cover a wide range of topics, and often the same topic from a variety of pedagogical angles.

1.2 Some Practical Challenges of AISs

Several practical problems have challenged widespread creation and adoption of these AISs despite their individual successes and collective contribution to our understanding of educational technologies more generally. AISs require large, diverse teams to work together effectively. For example, AutoTutor problems in electrical engineering required multiple experts in the fields of computer science, cognitive psychology, natural language processing, and, of course, electrical engineering [28]. These systems require a major investment in time and resources. Smaller, less expensive systems focus on a small band of pedagogical methods, learning principles, modalities, and content; but this runs the risk of yielding smaller learning gains and fewer learners who benefit from the AIS.

1.3 Multiple Representations

Learning functions on many levels (e.g., [29, 30] that can benefit from varied forms of instruction. Levels of learning also tend to build on one another, such that higher levels often assume competency on lower ones. While individual AISs may suffer from lack of breadth or prohibitive development time, they each potentially provide a valuable representation of the information. A staged algebraic approach can offer a concrete mathematical complement to the conceptual focus of conversational AISs. Simple word problems provide remedial representation of key concepts and relationships when targeted at problem areas. Easily accessible definitions and functional descriptions lower the barrier to interaction with information. And functional comprehension follows from representations integrating concepts, components, and relationships. Leveraging multiple representations in AISs can provide staged advancement when made concurrently available. This can also help ensure that learners stay within Vygotsky's zone of proximal development [31].

The National Academy of Sciences, Engineering, and Medicine [32] identified affordances of learning technologies. The affordances include interactivity, where the technology responds to learners' actions, and adaptivity, where information is contingent on the past behavior, knowledge, or characteristics of the learner. Taken together, these present a baseline technical qualification for AISs. Other affordances include providing feedback on quality of performance, offering a choice on what to learn next, and allowing nonlinear access to content for self-directed learners. Learning technology also affords linked representations that emphasize different conceptual viewpoints, open-ended learner input to encourage self-expression, and communication with others.

Leveraging all of these affordances into a single AIS presents a daunting challenge. However, given the potential for concurrently developed and available AISs in the same domain, a possible solution is to combine existing AISs into a larger, complex system. The resulting system would provide diverse modes of interacting with content, strengthening learners understanding by reinforcing through varied, stratified repetition. It could also foster ownership on the learner's part by allowing choice of not just content, but order, representation, and difficulty. The resulting confederated system would include a human instructor that orchestrates learning, together with the recommendations from an intelligent recommender system. This essentially is what we mean by a hybrid tutor: the best of accomplished human tutors and digital intelligence from an ensemble of digital resources.

The hybrid tutor would still require substantial investment by experts in diverse fields to create content, but the advantages stated should mean those hours are more likely to yield fruitful interactions with students. The pressing challenge becomes developing a way of translating progress in one system to progress in another, and in making intelligent recommendations across an array of resources on several levels (e.g., topic, modality, difficulty, sources, human vs. computer).

2 ElectronixTutor

The promise of a hybrid tutor spurred the development of *ElectronixTutor*. ElectronixTutor is a hybrid tutor, designed to supplement classroom instruction by leveraging multiple AISs (and conventional static learning resources) in a single platform. Critically, all individual learning resources contribute to a unified learner record store. This store translates progress among the many resources on several discrete levels.

The disciplined classification of these resources and levels allows the learning record store to inform an integrated recommender system. Collectively, these components leverage the established benefits of AIS interaction and provide both detailed (i.e., individual) and composite (i.e., classroom or population) learner information to a human instructor who can then manually set assignments at the item or topic level. The inclusion of multiple learning resources, both adaptive and static, allows ElectronixTutor to present learning content in multiple modes. Learners then have detailed records for how they interact with each one (e.g., time, performance, self-selected versus recommended or assigned).

These resources all appear in a common user interface (see Fig. 1). In this example, an AutoTutor conversation (complete with optional dialogue history and scroll-over information from Point & Query) appears in the activity window that dominates the screen. The left-side navigation bar includes site navigation as well as all course content. Featured prominently is the "Topic of the Day" facility, where instructors set the area of content to be mastered. Below that appear "Recommendations", based on the learners' history of interaction with the system holistically. Finally, learners can self-select any of the available problems from a drop-down menu (though instructors can limit content availability for pedagogical reasons).

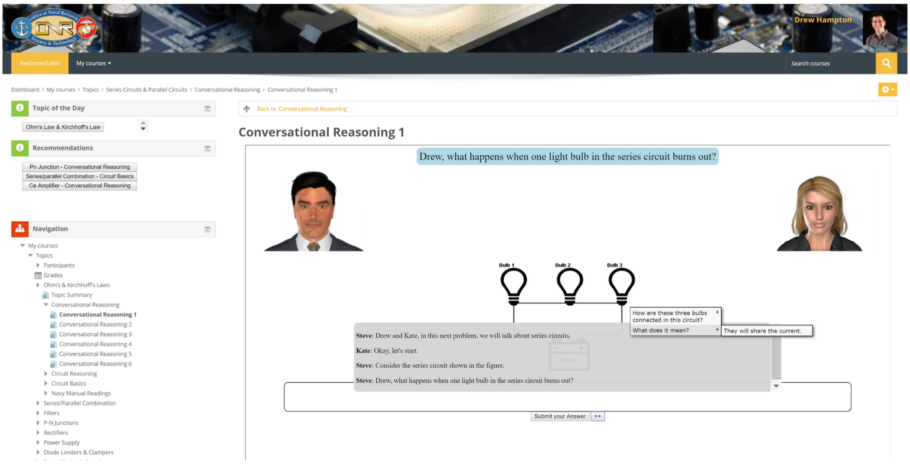

Fig. 1. The ElectronixTutor user interface, here showing an AutoTutor question with Point & Query engaged.

2.1 AutoTutor

AutoTutor [16, 19] presents conceptual questions on electrical circuits in a conversational exchange. Learners have both a tutor agent and a peer agent with whom to engage in a natural language discussion on a designated topic. The resulting "trialogues" always orient the learner to the topic, introduce an appropriate graphical representation, and directly address the learner by name when asking the main question. This allows learners to go from the concrete image to the deeper concept.

Further, each main question has several components of a full correct answer, with the AutoTutor Conversational Engine able to extract partial, as well as incorrect, responses from natural language input. This affords follow-up hints, prompts, or pumps from one or both conversational agents to elicit all information the learner knows about the topic at hand. In addition to the depth of understanding that AutoTutor examines, the analysis of breadth makes it an excellent diagnostic interaction with the learner to identify the appropriate next problem within the larger system. This approach has proven successful across numerous domains, including STEM topics such as computer literacy physics, biology, and scientific reasoning.

2.2 Point and Query

Within AutoTutor, Point & Query [28] aims to mitigate the difficulty many learners have in identifying appropriate questions to ask by offering a simple mouse-over interaction with circuit diagrams. The learner clicks on a hot spot, which launches a set of good questions to ask; the learner selects a question and immediately receives a good answer [33]. Research has demonstrated that this facility greatly increases the absolute number of interactions with the learning program. The low effort necessary to engage with content lowers the barrier and encourages further engagement by reinforcing question-asking behavior with immediate answers.

2.3 Dragoon

Dragoon [34] has learners construct and manipulate dynamic models of circuits, ensuring functional understanding of interacting parts by fostering the development of appropriate mental models (see Fig. 2). These questions represent by far the most difficult problems available. The holistic perspective on structures, parameters, and relationships among them requires comprehensive understanding. This adds substantial value, both in ensuring mastery with a high degree of confidence, and in providing challenges to the most advanced, diligent learners.

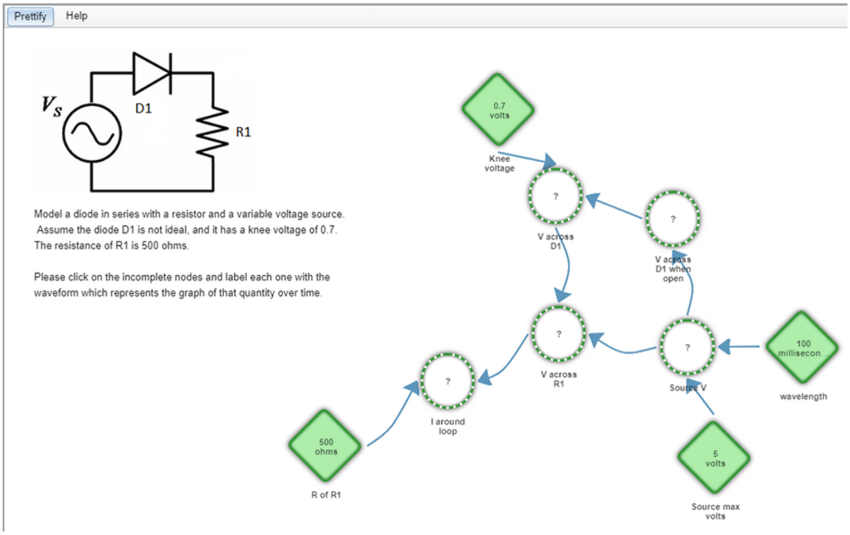

Fig. 2. A sample Dragoon problem, requiring detailed, comprehensive knowledge of the circuits.

2.4 LearnForm

LearnForm is oriented toward complex problem-solving on electronics circuit problems, with overarching problems that deconstruct into constituent parts and feedback (see Fig. 3). Mathematical reasoning and algebraic logic play an important and recurring role in electrical engineering. These problems ensure that learners have detailed knowledge of all required steps, with explanations provided in relatively simple mathematical sentences that build on one another until a complete, applied problem is complete.

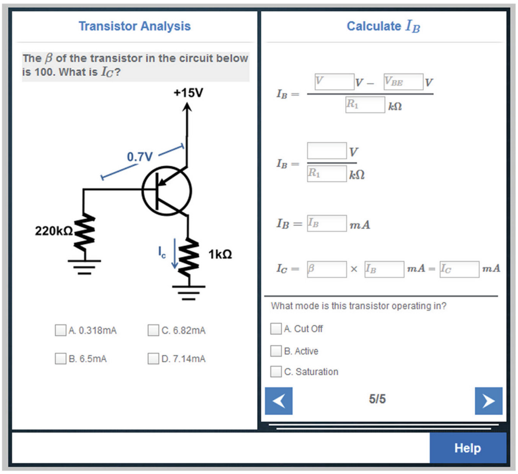

Fig. 3. A sample LearnForm problem, with algebraic formulas broken down into constituent parts while still relating to the original problem in total.

2.5 Beetle-II

BEETLE-II [13] addresses basic understanding of circuits, with a focus on introductory concepts such as voltage, current, open versus closed circuits, and how to find faults using voltage. These problems demonstrated learning gains, but only engage the macro-level of discourse and pedagogy as opposed to the micro-level language and content adaptation present in AutoTutor.

2.6 NEETS and Topic Summaries

As mentioned above, quality static texts will retain their place on the pedagogical landscape for the foreseeable future. The Navy Electricity and Electronics Training Series (NEETS) is a hefty collection of documents encapsulating all essential training information for the several Navy specialties dealing with electronics. These documents are both irrefutably useful and irredeemably dry. We provide these in context as a necessary backstop for any well-rounded education in Navy electrical engineering, though relying on learners to voluntarily engage (or reliably engage when compelled) remains a challenge beyond our scope. ElectronixTutor has indexed these texts and provides hyperlinks to the appropriate section when listed or recommended to learners.

The considerable depth of the NEETS suggested the need for a more approachable static text resource that would still allow learners to peruse at their own pace. To that end, subject matter experts collaboratively created topic summaries for each of the 15 topics covered in ElectronixTutor. These summaries comprised between one and four pages of a high-level overview, including diagrams, important definitions or formulas, and links to external resources such as Wikipedia or university webpages deemed to be of value.

2.7 Unified Learner Model and Recommender System

The inclusion of learning resources that instantiate such divergent pedagogical strategies is an important first step. However, these resources need to be organized and launched to the right person at the right time according to a disciplined framework. Knowledge components [35] provide a common currency of content to support this (see Fig. 4). In ElectronixTutor, every problem is annotated by experts on various content topics and knowledge components. They determine whether each learning resource includes these topics and knowledge components. For each of the 15 topics in the system on devices and versus circuits, decisions are made by the experts as to whether particular knowledge components tap the structure, behavior, function, or a parameter. This 15 (topic) $\times$ 2 (device versus circuit) $\times$ 4 (structure, behavior, function, parameter) matrix yields a list of 120 combinations. The landscape of learning resources has questions/items that touches at least one of these, and potentially several.

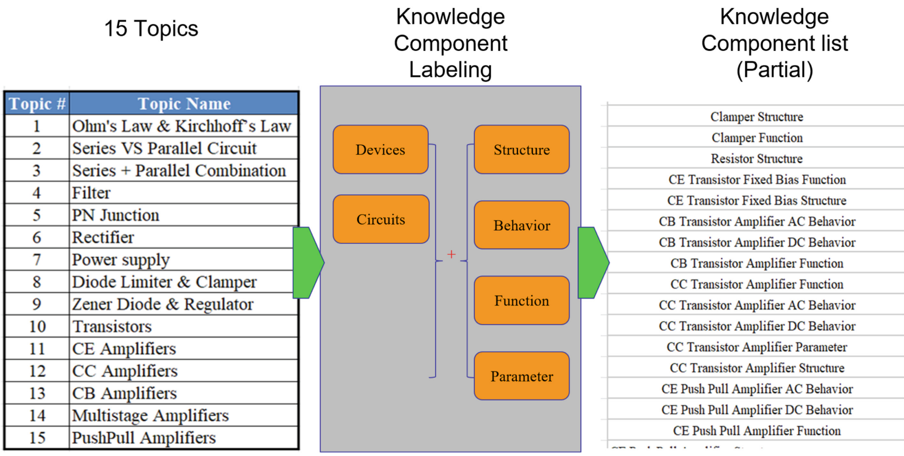

Fig. 4. Knowledge component mapping in ElectronixTutor

All learner interactions with learning content are discretized through knowledge components, formalized in xAPI format, and sent to the unified learner record store. An intelligent recommendation system ensures adequate coverage of each topic while keeping learners in the zone of proximal development, with problems not too difficult or too easy, but just within the learner's mastery. ElectronixTutor allows learners to choose problems in three ways. They can stay in the Topic of the Day, which completely controls their selection and guides them to what our algorithm judges to be the optimal problem for their advancement within a topic designated by a human instructor. Alternatively, they can also choose from among three recommended activities (item/question in a topic), providing a degree of control (and thereby psychological engagement) while staying within near-optimal parameters. Finally, learners can opt for completely self-directed learning, allowing full access to the suite of learning resources and range of topics. Progress in any of these will update the recommendations accordingly.

3 Preliminary Data

We have conducted several small-scale studies in university and trade school settings, though we have not yet attained our optimal goal of full classroom integration. Available data indicate that learners' performance in ElectronixTutor was not significantly correlated with degree of problem difficulty (ρ [12] = –0.068, p = 0.818), and stayed relatively stable near 78% correct. This means learners stayed roughly in the zone of proximal development, which was intentionally in the design of the system. That is, the stability is likely explained by adaptivity of the recommender system, and in part due to allowing users to self-select problems commensurate with their experience and level of comfort with the material.

We saw a good distribution of engagement relative to available topics. In Table 1, the most advanced topics appear at the top of the list, and most basic at the bottom. With the first eleven progressing topics garnering generally increasing use, we feel confident in the content match between sample and target populations to a point. Minimal use for the four most advanced topics suggests the need for some advanced classes to be incorporated into future studies for full system evaluation.

Table 1. Relative time spent in each available topic.

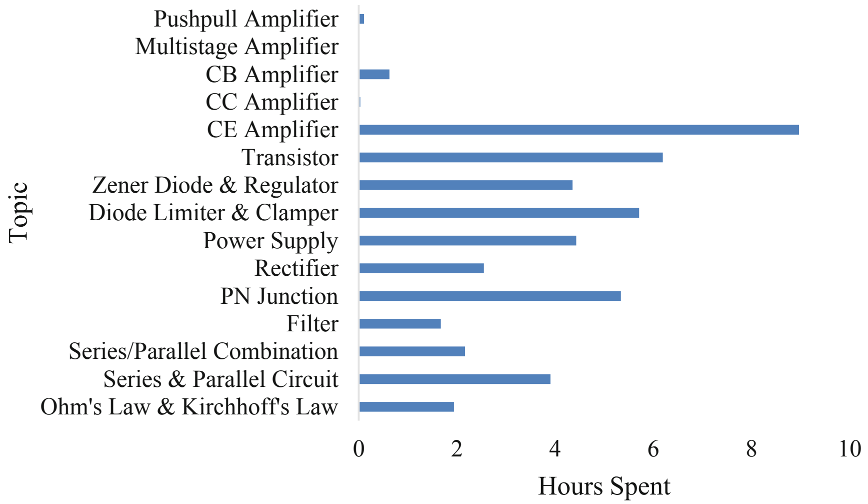

Survey data obtained from those who completed ten hours of interaction provide some promising results as well. While the total number of participants to complete the study was small (only 6 completed out of 50 to request log-on credentials), the majority indicated that they would continue to use ElectronixTutor without paid compensation. This indicates a certain start-up cost in learning the system that can be mediated by a more thorough introduction. Our initial attempt at classroom integration showed largely the same effect, with few able to move past the opening stages. It should be noted that these systems were used voluntarily, so adoptions were accepted to be low. The results are compatible with research on MOOCs, which are known to have high dropout rates.

We also identified some learners who became disengaged, particularly during the posttest assessment. Performance features like time on task and scores relative to historical performance make this detection relatively simple. In this way, human instructors can more readily intervene when learners become distracted or discouraged.

4 Conclusions, and Future Work

The integration of multiple existing AISs in a unified, cohesive platform is the essence of hybrid tutors. Integration enhances the range of interactions available, expanding the range of interaction types available to learners and leveraging more of the potential affordances of learning technology. Variability in representations is useful in fostering deep, lasting understanding of complex topics like electrical engineering. The multiple modes of content and strategy acquisition are expected to provide better transfer to new problems that the learner may encounter. Providing varying levels of control over content selection encourages engagement and investment in learning activities. Students can defer to classroom assignment (adaptive to their mastery level), evaluate a manageable number of recommendations (based on historical performance), or choose to explore freely. These overarching factors collectively suggest to a high likelihood that hybrid tutors like ElectronixTutor will yield a product greater than the sum of their parts.

Future work in this field will focus on tight integration into classrooms, with calendar functions determining assignments and system participation built into the syllabus. This is the ideal application of hybrid tutors. Interim goals include improvements based on learner feedback, notably improvements in the early stages when many participants dropped off. Further, deployment of ElectronixTutor "in the wild" (that is, made available widely to any who are interested in using it) will opportunistically recruit participants with motivation to learn, thus providing data and opportunities for iterative improvement. The University of Memphis library offers a "sandbox" learning tools interface that may facilitate this branch of inquiry. Further, Shelby County Schools, the school district surrounding the University of Memphis, has expressed interest in supplementing conventional classroom instruction with ElectronixTutor, potentially expanding the learner base.

References

1. Sottilare, R., Brawner, K.: Exploring standardization opportunities by examining interaction between common adaptive instructional system components. In: Proceedings of the First Adaptive Instructional Systems (AIS) Standards Workshop, Orlando, Florida (2018)
2. Graesser, A.C., Hu, X., Sottilare, R.: Intelligent tutoring systems. In: Fischer, F., Hmelo-Silver, C.E., Goldman, S.R., Reimann, P. (eds.) International Handbook of the Learning Sciences, pp. 246–255. Routledge, New York (2018)
3. Gilbert, S., et al.: Creating a team tutor using GIFT. Int. J. Artif. Intell. Educ. **28**, 286–313 (2018)
4. Sottilare, R., Graesser, A.C., Hu, X., Sinatra, A. (eds.): Design Recommendations for Intelligent Tutoring Systems: Team Science, vol. 6. U.S. Army Research Laboratory, Orlando (2018)
5. VanLehn, K.: The relative effectiveness of human tutoring, intelligent tutoring systems and other tutoring systems. Educ. Psychol. **46**, 197–221 (2011)
6. Sottilare, R., Graesser, A. C., Hu, X., Holden, H. (eds.): Design Recommendations for Intelligent Tutoring Systems, Volume 1—Learner modeling. U.S. Army Research Laboratory, Orlando (2013)

7. Woolf, B.P.: Building Intelligent Interactive Tutors. Morgan Kaufmann Publishers, Burlington (2009)
8. Aleven, V., Mclaren, B.M., Sewall, J., Koedinger, K.R.: A new paradigm for intelligent tutoring systems: example-tracing tutors. Int. J. Artif. Intell. Educ. **19**(2), 105–154 (2009)
9. Koedinger, K.R., Anderson, J.R., Hadley, W.H., Mark, M.: Intelligent tutoring goes to school in the big city. Int. J. Artif. Intell. Educ. **8**, 30–43 (1997)
10. Ritter, S., Anderson, J.R., Koedinger, K.R., Corbett, A.: Cognitive tutor: applied research in mathematics education. Psychon. Bull. Rev. **14**, 249–255 (2007)
11. Falmagne, J., Albert, D., Doble, C., Eppstein, D., Hu, X.: Knowledge Spaces: Applications in Education. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-35329-1
12. Lesgold, A., Lajoie, S.P., Bunzo, M., Eggan, G.: SHERLOCK: a coached practice environment for an electronics trouble-shooting job. In: Larkin, J.H., Chabay, R.W. (eds.) Computer Assisted Instruction and Intelligent Tutoring Systems: Shared Goals and Complementary Approaches, pp. 201–238. Erlbaum, Hillsdale (1992)
13. Dzikovska, M., Steinhauser, N., Farrow, E., Moore, J., Campbell, G.: BEETLE II: deep natural language understanding and automatic feedback generation for intelligent tutoring in basic electricity and electronics. Int. J. Artif. Intell. Educ. **24**, 284–332 (2014)
14. Fletcher, J.D., Morrison, J.E.: DARPA Digital Tutor: Assessment data (IDA Document D-4686). Institute for Defense Analyses, Alexandria (2012)
15. Johnson, W.L., Lester, J.C.: Face-to-face interaction with pedagogical agents, Twenty years later. International Journal of Artificial Intelligence in Education **26**(1), 25–36 (2016)
16. Nye, B.D., Graesser, A.C., Hu, X.: AutoTutor and family: a review of 17 years of natural language tutoring. Int. J. Artif. Intell. Educ. **24**(4), 427–469 (2014)
17. Craig, S.D., Twyford, J., Irigoyen, N., Zipp, S.A.: A Test of spatial contiguity for virtual human's gestures in multimedia learning environments. J. Educ. Comput. Res. **53**(1), 3–14 (2015)
18. Graesser, A.C., Li, H., Forsyth, C.: Learning by communicating in natural language with conversational agents. Curr. Dir. Psychol. Sci. **23**, 374–380 (2014)
19. Graesser, A.C.: Conversations with AutoTutor help students learn. Int. J. Artif. Intell. Educ. **26**, 124–132 (2016)
20. Biswas, G., Jeong, H., Kinnebrew, J., Sulcer, B., Roscoe, R.: Measuring self-regulated learning skills through social interactions in a teachable agent environment. Res. Pract. Technol. Enhanced Learn. **5**, 123–152 (2010)
21. Lane, H.C., Noren, D., Auerbach, D., Birch, M., Swartout, W.: Intelligent tutoring goes to the museum in the big city: a pedagogical agent for informal science education. In: Biswas, G., Bull, S., Kay, J., Mitrovic, A. (eds.) AIED 2011. LNCS (LNAI), vol. 6738, pp. 155–162. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-21869-9_22
22. Rowe, J.P., Shores, L.R., Mott, B.W., Lester, J.C.: Integrating learning, problem solving, and engagement in narrative-centered learning environments. Int. J. Artif. Intell. Educ. **21**, 115–133 (2011)
23. Johnson, L.W., Valente, A.: Tactical language and culture training systems: Using artificial intelligence to teach foreign languages and cultures. AI Magazine **30**, 72–83 (2009)
24. Dynarsky, M., et al.: Effectiveness of Reading and Mathematics Software Products: Findings from the First Student Cohort. U.S. Department of Education, Institute of Education Sciences, Washington (2007)
25. Steenbergen-Hu, S., Cooper, H.: A meta-analysis of the effectiveness of intelligent tutoring systems on college students' academic learning. J. Educ. Psychol. **106**, 331–347 (2013)
26. Dodds, P.V.W., Fletcher, J.D.: Opportunities for new "smart" learning environments enabled by next generation web capabilities. J. Educ. Multimedia Hypermedia **13**, 391–404 (2004)

27. Kulik, J.A., Fletcher, J.D.: Effectiveness of intelligent tutoring systems: a meta-analytic review. Rev. Educ. Res. **85**, 171–204 (2015)
28. Graesser, A.C., et al.: ElectronixTutor: an intelligent tutoring system with multiple learning resources. Int. J. STEM Educ. **5**(15), 1–21 (2018)
29. Bloom, T.M.E.: Bloom's Taxonomy of Educational Objectives. Longman, New York (1965)
30. Kyllonen, P.C., Shute, V.J.: Taxonomy of Learning Skills. Universal Energy Systems Inc., Dayton (1988)
31. Chaiklin, S.: The zone of proximal development in Vygotsky's analysis of learning and instruction. Vygotsky's Educ. Theor. Cult. Context **1**, 39–64 (2003)
32. National Academy of Sciences, Engineering, and Medicine: How people learn II: Learners, contexts, and cultures. National Academies Press, Washington, D.C. (2018)
33. Graesser, A.C., Hu, X., Person, N.K., Jackson, G.T., Toth, J.: Modules and information retrieval facilities of the human use regulatory affairs advisor (HURAA). Int. J. E-Learn. **3** (4), 29–39 (2004)
34. VanLehn, K., Chung, G., Grover, S., Madni, A., Wetzel, J.: Learning science by constructing models: can dragoon increase learning without increasing the time required? Int. J. Artif. Intell. Educ. **26**(4), 1033–1068 (2016)
35. Koedinger, K.R., Corbett, A.C., Perfetti, C.: The Knowledge-Learning-Instruction (KLI) framework: bridging the science-practice chasm to enhance robust student learning. Cogn. Sci. **36**(5), 757–798 (2012)