70. AI in education: landscape, vision and critical ethical challenges in the 21st century

Daniel S. Schiff and Rinat B. Rosenberg-Kima

AI IN EDUCATION: PAST, PRESENT, AND FUTURE

Artificial intelligence in education (AIED) is entering a new stage. While the origin of artificial intelligence (AI) research is often attributed to the 1950s, AIED is a relatively more recent field involving the adoption of AI in educational settings or for educational purposes. In particular, AIED began to develop in the 1970s under the moniker of computer-assisted instruction (CAI) (Carbonell, 1970), which had the general ambition of simulating a human tutor. Along these lines, researchers in the 1980s began to develop intelligent tutoring systems aimed at providing personalized instruction and feedback based on a deep understanding of student learning as inferred from their interaction with the AI system (Corbett et al., 1997). Following the establishment of the International AIED Society in 1997, AI techniques were applied to other related domains, including through educational data mining (EDM) and learning analytics and knowledge (LAK), resulting in the annual international conference on EDM starting in 2008 and the LAK conference starting in 2010.

Further developments in deep learning techniques in the 21st century enabled researchers to advance work related to social AI abilities, including research on natural language processing (NLP) and identification of facial expressions, gestures, posture, and gaze, along with other indicators of human affect (Lemaignan et al., 2016; McDuff et al., 2013). This progress has helped realize the use of computerized social or virtual agents for educational purposes (Goel & Polepeddi, 2016; Rosenberg-Kima et al., 2008). Relatedly, while most implementations of AIED historically involved a computer as the primary interface, these new socially oriented AI abilities, along with progress in the field of robotics, have opened the way for research and development on human-robot interaction (HRI) in education, as demonstrated with the initiation of the annual international conference on HRI in 2006. Figure 70.1 offers a brief history of selected key milestones related to AIED and AI generally.

In the few decades since AIED's origins, research and applications have proliferated and are now increasingly being utilized in real-world settings. There is thus every reason to think that AIED will be as impactful as AI in sectors like transportation, manufacturing, and healthcare. Yet perhaps surprisingly, AIED has not received as much attention in mainstream AI policy discourse (Schiff, 2021a), and ethical issues associated with AIED have only begun to be examined (Holmes et al., 2021). Meanwhile, historical failures of educational technology and a growing number of scandals associated specifically with AIED urge the need for increased attention to this field. As such, and as recent advances in AI generally have triggered widespread discussion of AI's implications, governance, and purpose, it is likewise timely for AIED to be subject to similar consideration and scrutiny.

This article, therefore, seeks to critically examine the role of AIED in the present and future, including by questioning its underlying assumptions and presumed or unexamined

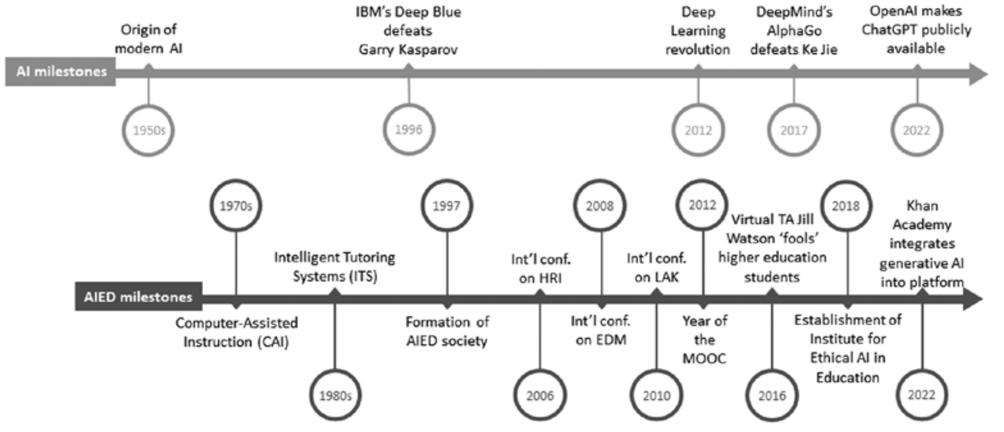

Source: Author.

Figure 70.1 Key milestones in the history of AIED and AI

trajectory. We begin by introducing AIED and providing an overview of the growing number of AIED applications in real-world settings. Next, we discuss why AIED applications have successfully scaled or failed to do so, reviewing key challenges related to implementation, governance, research, and ethics. Finally, we take stock of how AI's role in educational systems is unfolding, what ultimate goals and limitations are imagined, and whether the current trajectory of AIED is prudent or begs modification in light of this vision.

EVOLVING APPLICATIONS OF AIED: TEACHING, LEARNING, AND BEYOND

The techniques, applications, and roles of AIED systems have increased since the origin of the field three decades ago. Yet a useful place to start is with the intelligent tutoring system (ITS)—among the most commonly researched tools (Zawacki-Richter et al., 2019) and arguably the "holy grail" of AIED. At their core, ITSs are designed to computerize the teaching and learning process, making explicit the hidden dynamics of education and allowing for the automation of instruction (Self, 1998). They operate by representing content in a discipline (often STEM courses), assessing student engagement and performance on tasks and questions, and identifying the appropriate pedagogical response such as recommending subsequent content to learners (Zhang & Aslan, 2021). ITSs thus theoretically perform many of the core functions associated with teaching. Above all, ITSs are often lauded for their capacity for low-cost and easily-scalable personalization, functionally allowing for one-to-one instruction and the educational benefits that entails (Vincent-Lancrin & Vlies, 2020). Meta-analyses (Kulik & Fletcher, 2016; VanLehn, 2011; Xu et al., 2019) tend to support the capacity of ITS to solve the "2 sigma problem," referring to the ability of individualized tutoring to help students achieve their maximum potential, as compared with conventional group-based instruction that renders differentiation and mastery learning approaches less feasible (Bloom, 1984).

Various elements of ITS—and AIED more broadly—have been adopted to serve learners in a variety of settings, ranging from formal and in-class education to informal, self-guided, after-school, app-based, and online settings. For example, in addition to automated instruction, related functions aimed at learners include automated and personalized assessment (e.g., United States GREs) and feedback (e.g., Duolingo), recommendation of courses (e.g., Coursera), assistive technologies such as speech-to-text, and even academic and career counseling (Schiff, 2021b; Woolf et al., 2013; Zhang & Aslan, 2021). Further, AIED tools might be instantiated quite visibly in computers or online platforms via chatbots and other virtual agents (e.g., Jill Watson) or indeed physically embodied in social robots (e.g., Ozobot); alternatively, they may function more invisibly in the background of learning management systems (e.g., Canvas), intelligent textbooks, and massively online open courses (MOOCs).

Another set of AI systems is targeted at educators and educational administrators, construed broadly. Perhaps most centrally, AIED tools can support teachers by monitoring student learning progress, providing an individual- and classroom-level picture of educational needs (e.g., ASSISTments). Other such functions are largely logistical, like supporting course scheduling and automating teacher communication with parents. Notably, AIED tools are now increasingly used for indirect non-instructional purposes as well as instructional ones, such as to assess student dropout risk and intervention needs (e.g., Course Signals), to assign teachers to schools, inform admissions decisions in higher education, and optimize school operations related to transportation, security, or maintenance (T. Baker & Smith, 2019; Conner & Nelson, 2021; Diebold & Han, 2022).

Figure 70.2 provides an overview of prominent applications and implied purposes of AIED for learners and for teachers and administrators, as well as an incomplete list of underlying techniques, tools, and impacted stakeholders.

IMPLEMENTATION BARRIERS, FAILURE MODES, AND SCANDALS IN AIED

What, if anything, connects these diverse AIED applications? And what accounts for the relative success or failure of different kinds of AIED systems in advancing toward real-world usage? There are several candidate responses. AIED systems may advance because underlying techniques become technically feasible (capability-based reasons), because they fill important gaps in educational systems (needs-based reasons), because they become commercially viable (economic-based reasons), or because they align with certain idealized pathways associated with AI (vision-based). The tension between these drivers and constraints helps to explain differential progress in AIED adoption. Most notably, exciting technical advances, the promise of financial opportunity, and the pursuit of idealized visions may drive the development of AIED that is not especially targeted to serving the needs of students, leading to mismatches and even public-facing failures.

For example, despite numerous efforts to advance sophisticated, affect-aware, human-like tutors, this ambition has not been realized and remains in doubt. AIED systems are still limited in their ability to interpret social and emotional behavior (Belpaeme et al., 2018), and at worst, are accused of perpetuating pseudoscientific practices with historical analogs in physiognomy (Sloane et al., 2022). These limitations are especially apparent when AIED tools like ITS have been scaled beyond demonstration settings to impact large numbers of students. As

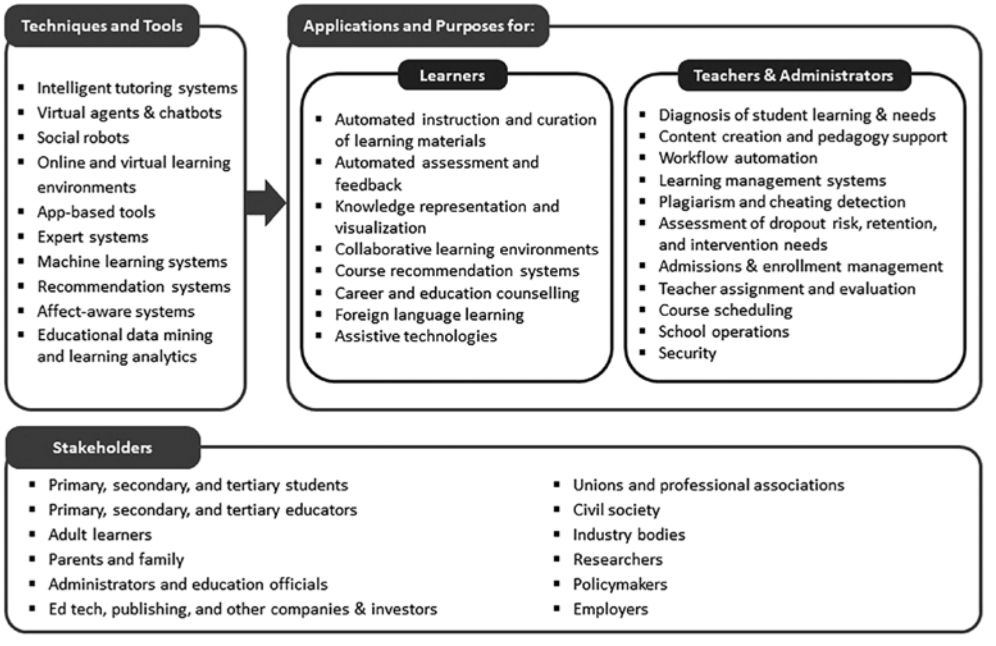

Source: Author.

Figure 70.2 An overview of AIED stakeholders, techniques, and purposes

Baker (2016, p. 601) notes, despite decades of research and significant advances employing techniques like Bayesian knowledge tracing and knowledge space theory, the most prominent systems today that model student knowledge rely on simple "heuristics to assess student mastery, such as whether the student gets three [answers] right in a row." Advanced systems, relative to simpler ones (e.g., adaptive assessments) seem to fail to provide a sufficient value to justify the complexities their adoption requires and the risks entailed.

The history of educational technology and AI's promises and overpromises are instructive toward understanding this development trajectory. Notably, many education technology (ed-tech) initiatives, especially in the context of low-income countries, have been marked by efficiency promises and techno-utopian visions that failed to appreciate relevant implementation and policy contexts (Sancho-Gil et al., 2020). Failed projects like One Laptop per Child (OLPC) and newer failures in school system adoption of MOOCs or other ed-tech systems (Rivard, 2013; Warschauer & Ames, 2010) underline this point. Even in countries with more advanced education infrastructure, the bottom-up nature of ed-tech adoption and the prominent role of teacher resistance (Hannafin & Savenye, 1993) delimit the willingness of educators to spend precious time on unproven technologies. Relatedly, while some national education systems are more centralized, allowing for coordinated adoption, many are not. AIED adoption thus often ultimately depends on the individual buy-in, capacity, and vision of individual educators, school leaders, and school systems, who must consider AIED in light of current pedagogical and curricular design as well as funding and policy constraints.

For instance, many ITS are pedagogically premised on variations of mastery learning, where students are allowed sufficient time in a flexible fashion (and with sufficient support) to foster certain competencies before progressing onto subsequent ones (Corbett et al., 1997). While such a strategy pairs nicely with the computerization logic of AIED (and digital platforms like Khan Academy), adoption of mastery learning in school systems is still limited, with modest and sometimes mixed evidence of effectiveness, and a plurality of barriers, including a lack of educator familiarity with mastery approaches (Kulik et al., 1990; Pane et al., 2017). Many schools still rely on whole-class (rather than personalized) learning, traditional textbooks, and the structure of grade bands to demarcate progress. Simply, we lack knowledge on how to even achieve mastery learning and personalized learning using *traditional* tools, much less intelligent ones.

Still, additional challenges are introduced when increasingly socially and emotionally aware AIED systems are introduced. While there is growing interest in these systems (Smakman et al., 2021), research on long-term interaction between students and robots is limited (Woo et al., 2021) and beset by various ethical, technical, and contextual challenges. Stakeholders have raised concerns not only about the potentially pseudoscientific nature of emotion recognition technologies but also about privacy concerns of capturing student physiological data (e.g., tracking eye movements or heart rates) or "nudging" students behaviorally, among numerous other ethical issues (Baker & Hawn, 2021). Additionally, human-human relationships are argued to be essential in educational settings for fostering personal growth, mentorship, productive comparison, and even the capacity to fail safely (Schiff, 2021b). Thus, it is far from clear that even socially sophisticated AIED is suitable to "replace" student-teacher or student-student relationships and the benefits they entail, either now or in the foreseeable future.

Overall, a simplistic notion that AIED systems will deterministically drive low-cost, scalable, highly personalized instruction to upend the current paradigm is based on faulty assumptions; *there is no short path to computerizing away the complexity of educational systems, the social-emotional dynamics of teaching and learning, the broader policy context, the current state of teacher preparation, and so on*. Both technical and sociotechnical trajectories for research have revealed that developing effective and responsible AIED is far more complex than assumed. AIED adoption, if blind to these issues, risks perpetuating old failures associated with educational technology and utopian AI imaginaries, as well as inducing new ones. Indeed, some of these failures have already manifested in the early years and decades of expanded AIED adoption in real-world settings.

One notable example is the A-levels algorithm scandal in the UK's Office of Qualifications and Examinations Regulation (Ofqual). During the Covid-19 pandemic, Ofqual employed an algorithm to adjust teacher predictions about grades that students *would* have received if they had taken examinations. After these adjustments were found to have disproportionately disadvantaged higher-need students, Ofqual's chief resigned and the algorithm was abandoned (Bedingfield, 2020). While some have noted that the algorithm was not especially sophisticated (and thus not really "AI") or have emphasized human failures in decision-making, it is critical to understand that these human decisions are *intrinsic* to the process of adopting and implementing AIED systems.

Other examples of recent AIED experimentation include the use of AIED in K-12 classrooms to examine student facial expressions, moods, and learning behaviors; the widespread adoption of eye-tracking to deter cheating on online exams during the pandemic; and the advancement of AI used in college admissions that draw on student interaction and even social media data (Newton, 2021), despite concerns raised about similar systems used to hire workers. And while some associated harms (e.g., gender or racial/ethnic bias) are relatively salient to many stakeholders, others are relatively subtle and require exploration to understand and demonstrate. Thus, while concerns about some of these systems have come to light, *a likely implication is that an increasing number of AIED systems are leading to ethical violations and producing harms, with these impacts either undetected or unreported.*

CAN ETHICAL FRAMEWORKS AND EMERGING AI GOVERNANCE HELP?

In light of the extended adoption of AIED systems and recent technical advances, and knowledge of the history of mishaps, what is being done to manage the ethical risks and otherwise govern AIED? Though the technical community of scholars engaged in AIED paid relatively little attention to ethics as recently as 2018 (Holmes et al., 2018), a positive development is that the community has now begun to call for ethics in the research agenda, lagging only slightly behind the broader AI ethics community (Holmes et al., 2021). New efforts like The Institute for Ethical Al in Education (formed in 2018) have drawn on stakeholder discussions and interviews to define key ethical principles and associated criteria and requirements for AIED, covering the design, procurement, and implementation of AIED systems. This increased attention is evident in the growing number of scholarly workshops and journal issues focused on ethics, such as a 2021 special issue on the fairness, accountability, transparency, and ethics (FATE) of AIED in the *International Journal of AI in Education*.

Satisfying these calls is no easy task, however, especially for a primarily technical community. It will likely require (at least) new modes of interdisciplinary collaboration with ethics and policy-oriented researchers, a change in structural incentives regarding publication and research priorities, the ability to pilot and evaluate AIED systems in large-scale rather than laboratory settings, thoughtful participation of a diverse array of stakeholders, and venues to facilitate engagement with education technology companies and policymakers (Schiff, 2021a). This challenge of collective governance is exacerbated because AIED is typically driven by academic research followed by industry translation, and then adopted in a bottom-up fashion. Thus, even if both researchers and education technology companies express concern regarding the ethical implications of AIED (Kousa & Niemi, 2022), addressing these concerns will require concerted efforts to promote responsible practices *throughout* the educational ecosystem, especially to safeguard settings with limited capacity and regulatory oversight.

For example, can we expect a teacher or administrator serving low-income students in a resource-poor school to know that they should ask an AIED startup about how their product has been developed and tested, about its differential accuracy across subgroups, about privacy and fairness considerations, and so on? Should we expect the company to have done adequate testing and be sober about the limitations and risks as they promote their product? *There are unfortunately good reasons to be skeptical of the capacity of the broader educational ecosystem to adopt some of the thoughtful practices currently being proposed by the AIED ethics community*, especially given countervailing pressures such as profit motives for providers and efficiency-seeking motives of the public sector. Paradigm shifts are hard.

Might formal governance provide a more realistic alternative? That is, can AIED providers be mandated to perform certain ethical checks, meet certain standards, and provide sufficient transparency to users? Can those who procure AIED systems be similarly required to engage in their own processes for vetting new AI systems, as has been proposed by The Institute for Ethical AI in Education and in line with certain social sectors and countries currently enhancing their AI procurement processes? The development of AI regulation suggests some reasons for worry. Prominent efforts like the EU's developing AI Act only require conformity assessments for certain "high-risk" systems, currently limited in education to those "used for the purpose of determining access" to institutions or for "assessing students" (European Commission, 2021). The challenge of delineating between low-risk and high-risk systems is quite relevant here. It is not clear a priori whether certain AIED applications are truly low risk, and *it has hardly been explored whether even modest impacts from "low-risk" systems might cascade and aggregate over time*. For example, a student, over their lifetime, could be taught inadequately by an ITS developed based on training data from very dissimilar students, punished because of a somewhat faulty plagiarism detection tool, and recommended a less-than-ideal career path or online course. What is the net result of many AIED systems ("low risk" or otherwise) working in concert to a student's (or teacher's) potential detriment?

While emerging and existing regulations surrounding student data and privacy, for example, might also apply, it thus seems likely that many of the diverse usages of AIED presented previously will fall outside of the scope of formal regulation. This is even more true in less-regulated settings like the US, which seems poised to adopt a risk-based classification system but may rely primarily on mere self-regulation, such as voluntary self-assessment using standards (National Institute of Standards and Technology, 2021). Countries without strong regulatory schemes and protections, including many low-income countries where student needs are vast, are even more susceptible to the promotion of AIED systems by unscrupulous external actors (Tzachor et al., 2022). Finally, another reason for the neglect of AIED's ethical and regulatory implications by leading policymakers is the highly emphasized role of education as the sector responsible for producing more AI-capable workers and subsequently driving innovation. Such an emphasis and associated urgency for education's role in building AI capacity arguably draws attention away from the impacts of AI on education itself (Schiff, 2021a), while AI's impacts on sectors like healthcare or transportation are better appreciated. In sum, due to the historical decentralization and self-regulation involved in education technology governance, the failure to imagine AIED's implications by leading policymakers, and the current emphasis on high-risk systems, *it is hardly clear that emerging AI governance is poised to combat the many concerns that are unfolding in AIED.*

QUESTIONS AND RECOMMENDATIONS FOR AIED IN THE 21st CENTURY

This chapter has reviewed the increasingly diverse and complicated landscape of AIED systems, some of the pressures and assumptions driving AIED adoption, a growing list of scandals and ethical risks, and where current efforts in ethical oversight and governance stand. Despite the historical experience and advancing scholarly and practitioner understanding of risks and failures of AI, education technology, and AIED, current efforts to adopt AIED responsibly seem poised to fail to keep pace. In addition to some of these suggestions offered above, this chapter, therefore, offers additional possibilities.

Safe experimentation. Piloting AIED systems is necessary to perform research on AIED's benefits and risks, and an overly precautionary approach can foreclose on urgently needed gains in educational systems. What is needed is *safe* experimentation. AIED research can be gradually applied to larger settings, deeply informed by ethical frameworks, social and policy contexts, and teacher and student needs. Methods like design-based implementation research (Fishman et al., 2013) and engagement with AIED-specific ethical frameworks like that from The Institute for Ethical Al in Education can underpin responsible research projects. Careful research projects, including experiments, can be undertaken before AIED systems are introduced at large scale by the private or public sectors. The AIED community has the opportunity to avoid a "wild west" reputation and instead be a leader in the AI space, a necessity given the vulnerable populations it affects and the scope of AIED's impacts.

Expanded understanding of social and ethical implications. The AI ethics community has devoted a significant portion of its effort to addressing ostensibly technically tractable issues like algorithmic bias, transparency, and privacy, issues for which there are imagined technical "solutions." Yet, there are many more ethical concerns that have received less attention, especially the kinds of murky sociotechnical issues that do not lend themselves to easy solutions. AIED researchers who have explored these issues have the opportunity to both learn from and help mature the broader AI ethics agenda. Associated topics might include the role of AIED in fostering justice within and across school systems or regions, the mental health of students, the well-being of teachers, the ethics of nudging and deception, and many more issues.

Reporting on risks and harms. Relatively little is known about the individual risks of the growing variety of AIED systems, much less their potential harms in combination, over time, for groups versus individuals, and so on. Prominent AIED or education actors should systematically track risks and harms from different types of AIED systems, and catalog this for research, development, implementation, and policy guidance purposes. For example, ethical harms and scandals related to AIED could be reported to the Partnership on AI's AI Incident Database (McGregor, 2020) or through the EU's developing AI database meant to track harms to well-being and human rights. This can foster a more holistic research agenda on AIED's impacts, as well as growing awareness and accountability.

Revisiting the vision for AIED. As noted, much of AIED's research trajectory—as with AI generally—derives from an idealized vision of human-like AI. Yet adopting AIED systems that spin out of this research program, and the program itself, may not be optimal for serving the needs of students, teachers, and other education stakeholders. It is unclear whether we would even want pseudo-conscious AI systems that pretend to portray emotions (much less fully conscious ones) teaching students or emulating teachers. Even so, it remains dubious that AIED will not be used to replace (rather than augment) educators, despite the best intentions. Using AIED to foster personalization and low-cost scaling may be laudable goals, but between the instrumentalized economic logic of efficiency and the idealized technical logic of computerizing human teachers, it may be time for AIED stakeholders to reconsider the long-term vision and course of action, perhaps cantered more closely on real-world contexts and needs.

AIED is a promising and dynamic field, increasingly touching the lives of students and educators, and mirroring the development of AI in key ways in the 21st century. Through historical knowledge, lived experience, and careful examination of the benefits and risks of the diversifying array of AIED systems, we can chart a more careful and beneficial course for the decades to come.

REFERENCES

Baker, R. S. (2016). Stupid Tutoring Systems, Intelligent Humans. *International Journal of Artificial Intelligence in Education*, *26*(2), 600–614. https://doi.org/10.1007/s40593-016-0105-0

Baker, R. S., & Hawn, A. (2021). Algorithmic Bias in Education. *International Journal of Artificial Intelligence in Education*. https://doi.org/10.1007/s40593-021-00285-9

Baker, T., & Smith, L. (2019). *Educ-AI-tion Rebooted? Exploring the Future of Artificial Intelligence in Schools and Colleges* (p. 56). Nesta. https://media.nesta.org.uk/documents/Future_of_AI_and_education_v5_WEB.pdf

Bedingfield, W. (2020, August 19). Everything That Went Wrong with the Botched A-Levels Algorithm. *Wired UK*. https://www.wired.co.uk/article/alevel-exam-algorithm

Belpaeme, T., Kennedy, J., Ramachandran, A., Scassellati, B., & Tanaka, F. (2018). Social Robots for Education: A Review. *Science Robotics*, *3*(21), eaat5954. https://doi.org/10.1126/scirobotics.aat5954

Bloom, B. S. (1984). The 2 Sigma Problem: The Search for Methods of Group Instruction as Effective as One-to-One Tutoring. *Educational Researcher*, *13*(6), 4–16. https://doi.org/10.3102/0013189X013006004

Carbonell, J. R. (1970). AI in CAI: An Artificial-Intelligence Approach to Computer-Assisted Instruction. *IEEE Transactions on Man-Machine Systems*, *11*(4), 190–202. https://doi.org/10.1109/TMMS.1970.299942

Conner, T., & Nelson, C. (2021). *Ensuring Equitable AI Adoption in Education: An Initial Framework and Proposed Next Steps* (No. 21-02538–02; p. 14). Oak Ridge Associated Universities and the MITRE Corporation.

Corbett, A. T., Koedinger, K. R., & Anderson, J. R. (1997). Chapter 37—Intelligent Tutoring Systems. In M. G. Helander, T. K. Landauer, & P. V. Prabhu (Eds.), *Handbook of Human-Computer Interaction (Second Edition)* (pp. 849–874). North-Holland. https://doi.org/10.1016/B978-044481862-1.50103-5

Diebold, G., & Han, C. (2022). *How AI Can Improve K-12 Education in the United States* (p. 33). Center for Data Innovation. https://www2.datainnovation.org/2022-ai-education.pdf

European Commission. (2021). *Proposal for a Regulation on a European Approach for Artificial Intelligence* (COM(2021) 206 final; p. 108). European Commission. https://digital-strategy.ec.europa.eu/en/library/proposal-regulation-european-approach-artificial-intelligence

Fishman, B. J., Penuel, W. R., Allen, A.-R., Cheng, B. H., & Sabelli, N. (2013). Design-Based Implementation Research: An Emerging Model for Transforming the Relationship of Research and Practice. *National Society for the Study of Education*, *112*(2), 136–156.

Goel, A. K., & Polepeddi, L. (2016). *Jill Watson: A Virtual Teaching Assistant for Online Education*. 21.

Hannafin, R. D., & Savenye, W. C. (1993). Technology in the Classroom: The Teacher's New Role and Resistance to it. *Educational Technology*, *33*(6), 26–31.

Holmes, W., Bektik, D., Whitelock, D., & Woolf, B. P. (2018). Ethics in AIED: Who Cares? In C. Penstein Rosé, R. Martínez-Maldonado, H. U. Hoppe, R. Luckin, M. Mavrikis, K. Porayska-Pomsta, B. McLaren, & B. du Boulay (Eds.), *Artificial Intelligence in Education: 19th International Conference, AIED 2018, London, UK, June 27–30, 2018, Proceedings, Part II* (Vol. 10948). Springer International Publishing. https://doi.org/10.1007/978-3-319-93846-2

Holmes, W., Porayska-Pomsta, K., Holstein, K., Sutherland, E., Baker, T., Shum, S. B., Santos, O. C., Rodrigo, M. T., Cukurova, M., Bittencourt, I. I., & Koedinger, K. R. (2021). Ethics of AI in Education: Towards a Community-Wide Framework. *International Journal of Artificial Intelligence in Education*. https://doi.org/10.1007/s40593-021-00239-1

Kousa, P., & Niemi, H. (2022). AI Ethics and Learning: EdTech Companies' Challenges and Solutions. *Interactive Learning Environments*, 1–12. https://doi.org/10.1080/10494820.2022.2043908

Kulik, C.-L. C., Kulik, J. A., & Bangert-Drowns, R. L. (1990). Effectiveness of Mastery Learning Programs: A Meta-Analysis. *Review of Educational Research*, *60*(2), 265–299. https://doi.org/10.3102/00346543060002265

Kulik, J. A., & Fletcher, J. D. (2016). Effectiveness of Intelligent Tutoring Systems: A Meta-Analytic Review. *Review of Educational Research*, *86*(1), 42–78. https://doi.org/10.3102/0034654315581420

Lemaignan, S., Garcia, F., Jacq, A., & Dillenbourg, P. (2016). From Real-Time Attention Assessment to "with-me-ness" in Human-Robot Interaction. *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 157–164. https://doi.org/10.1109/HRI.2016.7451747

McDuff, D., Kaliouby, R., Senechal, T., Amr, M., Cohn, J., & Picard, R. (2013). *Affectiva-MIT Facial Expression Dataset (AM-FED): Naturalistic and Spontaneous Facial Expressions Collected*, 881–888. https://www.cv-foundation.org/openaccess/content_cvpr_workshops_2013/W16/html/McDuff_Affectiva-MIT_Facial_Expression_2013_CVPR_paper.html

McGregor, S. (2020). Preventing repeated real world AI failures by cataloging incidents: The AI Incident Database. ArXiv:2011.08512 [Cs]. http://arxiv.org/abs/2011.08512

National Institute of Standards and Technology. (2021). *Summary Analysis of Responses to the NIST Artificial Intelligence Risk Management Framework (AI RMF)—Request for Information (RFI)* (p. 10). National Institute of Standards and Technology. https://www.nist.gov/system/files/documents/2021/10/15/AI%20RMF_RFI%20Summary%20Report.pdf

Newton, D. (2021, April 26). *Artificial Intelligence is Infiltrating Higher Ed, from Admissions to Grading*. The Hechinger Report. http://hechingerreport.org/from-admissions-to-teaching-to-grading-ai-is-infiltrating-higher-education/

Pane, J., Steiner, E., Baird, M., Hamilton, L., & Pane, J. (2017). *Informing Progress: Insights on Personalized Learning Implementation and Effects* (p. 59). RAND Corporation. https://doi.org/10.7249/RR2042

Rivard, R. (2013, July 18). *Udacity Project on "Pause."* Inside Higher Ed. https://www.insidehighered.com/news/2013/07/18/citing-disappointing-student-outcomes-san-jose-state-pauses-work-udacity

Rosenberg-Kima, R. B., Baylor, A. L., Plant, E. A., & Doerr, C. E. (2008). Interface Agents as Social Models for Female Students: The Effects of Agent Visual Presence and Appearance on Female Students' Attitudes and Beliefs. *Computers in Human Behavior*, *24*(6), 2741–2756.

Sancho-Gil, J. M., Rivera-Vargas, P., & Miño-Puigcercós, R. (2020). Moving Beyond the Predictable Failure of Ed-Tech Initiatives. *Learning, Media and Technology*, *45*(1), 61–75. https://doi.org/10.1080/17439884.2019.1666873

Schiff, D. (2021a). Education for AI, Not AI for Education: The Role of Education and Ethics in National AI Policy Strategies. *International Journal of Artificial Intelligence in Education*. https://doi.org/10.1007/s40593-021-00270-2

Schiff, D. (2021b). Out of the Laboratory and Into the Classroom: The Future of Artificial Intelligence in Education. *AI & Society*, *36*(1), 331–348. https://doi.org/10.1007/s00146-020-01033-8

Self, J. (1998). The Defining Characteristics of Intelligent Tutoring Systems Research: ITSs Care, Precisely. *International Journal of Artificial Intelligence in Education (IJAIED)*, *10*, 350–364.

Sloane, M., Moss, E., & Chowdhury, R. (2022). A Silicon Valley Love Triangle: Hiring Algorithms, Pseudo-Science, and the Quest for Auditability. *Patterns*, *3*(2), 100425. https://doi.org/10.1016/j.patter.2021.100425

Smakman, M., Vogt, P., & Konijn, E. A. (2021). Moral Considerations On Social Robots in Education: A Multi-Stakeholder Perspective. *Computers & Education*, *174*, 104317. https://doi.org/10.1016/j.compedu.2021.104317

Tzachor, A., Devare, M., King, B., Avin, S., & Ó hÉigeartaigh, S. (2022). Responsible Artificial Intelligence in Agriculture Requires Systemic Understanding of Risks and Externalities. *Nature Machine Intelligence*, *4*(2), 104–109. https://doi.org/10.1038/s42256-022-00440-4

VanLehn, K. (2011). The Relative Effectiveness of Human Tutoring, Intelligent Tutoring Systems, and Other Tutoring Systems. *Educational Psychologist*, *46*(4), 197–221.

Vincent-Lancrin, S., & Vlies, R. van der. (2020). *Trustworthy Artificial Intelligence (AI) in Education: Promises and challenges*. OECD. https://doi.org/10.1787/a6c90fa9-en

Warschauer, M., & Ames, M. (2010). Can One Laptop per Child Save the World's Poor? *Journal of International Affairs*, *64*(1), 33–51.

Woo, H., LeTendre, G. K., Pham-Shouse, T., & Xiong, Y. (2021). The Use of Social Robots in Classrooms: A Review of Field-Based Studies. *Educational Research Review*, *33*, 100388. https://doi.org/10.1016/j.edurev.2021.100388

Woolf, B. P., Lane, H. C., Chaudhri, V. K., & Kolodner, J. L. (2013). AI Grand Challenges for Education. *AI Magazine*, *34*(4), 9.
Xu, Z., Wijekumar, K. (Kay), Ramirez, G., Hu, X., & Irey, R. (2019). The Effectiveness of Intelligent Tutoring Systems on K-12 Students' Reading Comprehension: A Meta-Analysis. *British Journal of Educational Technology*, *50*(6), 3119–3137. https://doi.org/10.1111/bjet.12758
Zawacki-Richter, O., Marín, V. I., Bond, M., & Gouverneur, F. (2019). Systematic Review of Research on Artificial Intelligence Applications in Higher Education – Where are the Educators? *International Journal of Educational Technology in Higher Education*, *16*(1), 39. https://doi.org/10.1186/s41239-019-0171-0
Zhang, K., & Aslan, A. B. (2021). AI Technologies for Education: Recent Research & Future Directions. *Computers and Education: Artificial Intelligence*, *2*, 100025. https://doi.org/10.1016/j.caeai.2021.100025