42

ARTIFICIAL INTELLIGENCE IN EDUCATION

Wayne Holmes/ Maya Bialik/ Charles Fadel[545]

42.1 AIED Can Offer More

Much of the *Artificial Intelligence in Education (AIED)* involves the application of AI techniques to mainstream learning approaches, and tends to reflect (or automate) existing educational assumptions and practices. In addition, much AIED has been designed (whether intentionally or not) to supplant teachers or to reduce them to a functional role[546] and

[545] The article is an excerpt from Wayne Holmes/ Maya Bialik/ Charles Fadel, Artificial Intelligence in Education, The Center for Curriculum Redesign, Boston, 2019, 151-180. With permission of the publisher.

[546] Worryingly, one of the developers we have mentioned has suggested that the sophistication of their AIED means that teachers only need to play an auxiliary role, working like fast-food chefs ("KFC-like") to strictly regulated scripts.

not to assist them to teach more effectively. This approach, while potentially useful in contexts where teachers are few and far between, clearly undervalues teachers' unique skills and experiences, as well as learners' needs for social learning and guidance. However, instead of just automating the teaching of students sat at computers, conceivably AI might help open up teaching and learning possibilities that are otherwise difficult to achieve, that challenge existing pedagogies, or that help teachers to be more effective. Here we will speculate on some possibilities, some of which have been foreshadowed by the AIED tool, while others are both novel and complex to achieve, and all of which raise interesting social questions. We begin with AI to support collaborative learning, then AI-driven student forum monitoring, AI to support continuous assessment, AI learning companions for students, and AI teaching assistants for teachers, before concluding with AIED as a research tool to further the learning sciences (i.e. in order to help us better understand learning).[547]

42.1.1 Collaborative Learning

Collaborative learning, where students work together to solve problems, is well known to be able to lead to better learning outcomes, but effective collaboration between learners can be difficult to achieve.[548] AIED offers various possibilities. To begin with, an AIED tool could automatically suggest groups of students best suited for particular collaborative tasks, drawing on and making intelligent connections between individual student models (each of which comprises knowledge about

---

[547] One intriguing use of AI in education that we will not consider in detail, because its efficacy has not yet been demonstrated, but that should still be acknowledged is the automatic generation of quiz questions (https://mt.clevere.st and https://learningtools.donjohnston.com/product/quizbot).

[548] Luckin, R., et al. 2017. Solved! Making the Case for Collaborative Problem-Solving. Nesta. https://www.nesta.org.uk/report/solved-making-the-case-for-collaborative-problem-solving/

the student's previous learning experiences and achievements, what the student is learning in other classrooms, their personalities, and more).[549] Having elicited the teachers' requirements, the tool might also suggest groups of mixed or similar-ability students, or groups designed to give particular students opportunities to take on leadership roles, or groups that avoid personality or temperament clashes, and so on, all the while enabling the teacher to quickly and easily override any of the tool's suggestions (which the AI will learn from, for next time). An AIED tool might also take on the role of expert facilitator or moderator, monitoring student collaborative activities, recognizing when students are having trouble understanding shared concepts, and then providing targeted support. Alternatively, the AIED might involve a virtual agent that actively contributes to the group discussions (acting as a virtual peer or a teachable agent), or that makes dynamic connections (either with discussions being held by other groups in the same classroom, or with relevant materials drawn from the semantic web). In fact, some research into AI to support collaborative learning has been undertaken,[550] but there are many technical issues to overcome before it becomes possible in real classrooms.

---

[549] The Universitat Politècnica de València have been researching just such a system: Alberola, J.M., del Val, E., Sanchez-Anguix, V., Palomares, A., and Teruel, M.D. 2016. "An artificial intelligence tool for heterogeneous team formation in the classroom." Knowledge-Based Systems 101: 1–14. https://doi.org/10.1016/j.knosys.2016.02.010

[550] E.g., Diziol, D., et al. 2010. "Using intelligent tutor technology to implement adaptive support for student collaboration." Educational Psychology Review 22 (1): 89–102. https://doi.org/10.1007/s10648-009-9116-9 and Spikol, D., et al. (2016). "Exploring the interplay between human and machine annotated multimodal learning analytics in hands-on stem activities." In Proceedings of the Sixth International Conference on Learning Analytics & Knowledge. 522–523.

42.1.2 Student Forum Monitoring

Increasingly, students of all ages are participating in online education, which usually involves the use of discussion forums. Students might post to forums in response to given tasks or to engage in collaborative learning opportunities, or they might want to contact their tutors to clarify course requirements or to ask about course materials. Accordingly, especially when there are large cohorts of students (as can be typical of some distance universities and MOOCs), these online forums can generate massive numbers of forum posts, all of which must be monitored, moderated, and addressed. However, as the number of forum posts increases, this becomes at best an inefficient use of a tutor's time (dealing with repetitive and minor practical issues) and at worst an increasingly impossible task. It also makes it difficult for students to keep up to date with other student posts that might connect to their interests.

AIED might help in a number of ways (again, some research has already been conducted in this area)[551] — in particular by helping the teachers/tutors to be better able to support their students. First, an AIED tool might triage the forum posts, identifying those that can be dealt with automatically (perhaps practical questions around course dates, such as "When do I need to submit…?"), and those that require a response from a human tutor (such as those discussing more in-depth core subject issues). The simple posts, the ones that the AIED is capable of handling, would receive immediate automatic responses, relieving the human tutors of much repetitive work while enabling the students to move on quickly to more substantive work. Other posts would automatically be referred up to a human tutor, to ensure that students receive high quality, appropriate responses whatever the nature of their posting.

[551] Goel, A.K., and Joyner, D.A. 2017. "Using AI to teach AI: Lessons from an online AI class." AI Magazine 38(2): 48. https://doi.org/10.1609/ aimag.v38i2.2732

Taking this a step further, the more demanding posts (of which there still might be many) would be further analysed, the aim being to identify and aggregate similar posts or posts that raise overlapping issues (in a course with a thousand students, it is unlikely that there will be a thousand unique responses to a single course activity, but rather a much smaller number of closely related posts). A human tutor would then write a response to the much smaller number of aggregated posts, which in turn would be issued to all of the original posters. Although this is unlikely to be as good as replying to each individual student, it would clearly be better than the students receiving no responses at all—which, in a large online course, can all too often be the case. Another approach that might also help in student forums is for the AIED to interpret and make dynamic connections between posts, informing tutors when particular issues have been raised (e.g., known and unknown misconceptions), for them to address, or informing students about other posts that they might find interesting.

Finally, the AIED might also use sentiment analysis AI techniques to identify posts that reveal negative or non-productive student emotional states (perhaps a student is overly challenged, or likely to drop out of the course, or possibly suffering from mental health issues), posts that are unacceptable (perhaps because they include racist, misogynist or gratuitously aggressive comments), or posts that suggest topic drift (the tendency for forum posts to drift from the original intent). Any such posts (which, because of the overall number of posts, can be easy for humans to miss) would be referred up to a human tutor, so that the tutor can respond quickly, appropriately and effectively (perhaps by calling the student by phone, rather than depending on a digital intervention). Together, these various techniques might also enable tutors to be kept well informed of student opinions, collective worries, or recurrent themes that emerge from the forums.

42.1.3 Continuous Assessment

Psychologists and educators know that it is wrong to make decisions based upon a single test score and that decisions should reflect a balanced, complete understanding of each child. Numbers and scores can be very misleading if we don't consider the whole picture, something that means using both a qualitative and quantitative approach.[552]

Although there is little evidence for their validity, reliability or accuracy, high-stakes examinations are core to educational systems around the world.[553] Perhaps this is because that is how it has always been, perhaps because they efficiently rank students, perhaps because no practical, cost effective at scale, alternative has ever been devised, or perhaps because those who run the systems are typically those who were most successful at exams (and do not emotionally resonate with the need for change). Whatever the reason, with high-stakes examinations in place, schools and universities all too often end up teaching to the test, prioritising routine cognitive skills and knowledge acquisition over in-depth understanding and authentic application. In other words, the examinations, rather than the needs of students or wider society, determine what is taught and learned. Meanwhile, ironically, AI technologies are automating exactly the type of knowledge that examinations mostly assess: "There's lots of elements of human intelligence that cannot be automated but the bit that we've tended to value, that relates to academic exam success, is one of the bits that we've managed to automate."[554] In any case, stop-and-test examinations (standardised, unseen tests that are at set points in the learning schedule, thus potentially interrupting the learning) are not able to rigorously evaluate a student's understanding of all that has been learned—at best they can only provide a snapshot of frag-

---

[552] Gunzelmann, B.G. 2005. "Toxic testing: It's time to reflect upon our current testing practices." Educational Horizons 83 (3): 214.

[553] Evolving Assessments for a 21st Century Education.

[554] Rose Luckin quoted in https://www.jisc.ac.uk/news/the-ai-revolution-is-here-17-aug-2018

ments of what has been studied over the duration of a course. Last, but not least, students of all ages can sometimes suffer from serious exam anxiety, which can easily negatively impact on the student's success in a typical three-hour end of course examination (further clouding their accuracy and trustworthiness).

Nonetheless, most AIED research in this area has been unambitious. It has focused on improving existing examination systems (developing AI-driven techniques to authenticate the identity of students taking exams online),[555] rather than challenging the underlying principles. However, as we have seen, typical ITS and other AIED tools already and constantly monitor student progress to provide targeted feedback and to assess whether the student has achieved mastery of the topic in question. Similar information could be captured by AIED tools designed to support collaborative learning, while intelligent essay assessment tools can also make inferences about a student's understanding. All of this information and more might be collated throughout a student's time in formal educational settings (the learning sciences have long understood the value of students engaging with constructive assessment activities), together with information about the student's engagement with non-formal learning (such as learning a musical instrument, or a craft or other skills) and informal learning (such as language learning or enculturation by means of learning from experience or active participation), to help create a picture of the whole learner. In other words, the AI-driven assessment would happen in the background, all of the time—making it next to impossible for students to cheat or subvert the system's intention (as can be the case when wealthier students employ personal tutors),[556] or take the test as many times as necessary until they achieve a good-enough score.

---

[555] For example, http://tesla-project.eu

[556] Luckin, R. 2017. "Towards artificial intelligence-based assessment systems." Nature Human Behaviour 1. https://doi.org/10.1038/s41562-016-0028

This more detailed and nuanced information about an individual student might then be represented (and perhaps visualised in dynamic graphics) in an AI-driven e-portfolio,[557] an intelligent personal resumé (in fact, an extended open student model). This e-portfolio could perhaps be underwritten and authenticated by blockchain technologies[558] as used by virtual currencies such as Bitcoin (essentially open, distributed ledgers, hosted simultaneously by millions of computers across the internet and linked using cryptography, that can share data in a verifiable, incorruptible, and accessible way). In this way, students would have a robust, accredited, in-depth record of all their learning experiences and achievements, far more detailed and useful than a collection of certificates. Parts or all of this smart resumé they might share when applying for admission to another course or for a new job, while retaining full control of their academic persona and data. From a learner's perspective, an additional benefit is that continuous assessment can act as a moving average, a fluid-like shock absorber that evens out the blips of bad days and disadvantageous personal situations (it simply does not sense that a young person's academic outcomes and future life can be determined by difficulties at home that coincide with the day of an important exam).

In short, although the constant monitoring of student behaviours and achievements raises significant and far-reaching ethical questions that must first be properly investigated and addressed, it is conceivable that stop-and-test examinations could soon be entirely removed from our educational systems and relegated to a more primitive past.

---

[557] Per one of the authors' US patent numbers 9,262,640 and 9,582,567, which also protect privacy and security.

[558] Sharples, M. and Domingue, J. 2016. "The blockchain and kudos: A distributed system for educational record, reputation and reward." In European Conference on Technology Enhanced Learning. Springer. 490–496.

42.1.4 AI Learning Companions

The smart resumés that we have just proposed could also play a role in a much larger AIED possibility: AI-driven lifelong students' learning companions.[559] As we have seen, the desire for every student to have their own personalised tutor is what first inspired the development of ITS, but what about taking this to its logical conclusion? AI has the potential to provide every student with their very own personalised learning companion, operating sometimes as a learning partner, other times as a guide through the mass of available learning opportunities, and sometimes as an instructor, all the time recording the student's interests and progress in their blockchain-protected, smart resumé. The arrival and rapid developments of Siri, Cortana, Google Home and Alexa, suggest that this possibility is tantalisingly close.[560] In many countries, smartphones with extraordinary processing power and always-on internet access are more than common. It would not necessarily be a big technical step to leverage these capabilities, to create an AI-driven smartphone learning companion that could accompany and support individual learners throughout their studies, from kindergarten to old age.

Such a learning companion brings many possibilities. Once the student has decided on a particular topic of interest, it might provide some instructional activities, monitor the student's progress, remind them when a task needs to be completed, and offer targeted feedback and guidance—all on their speech-driven smartphone (and available on all their other devices). In other words, it might function as what we have called an ITS+.

But a learning companion would also operate at a higher and more strategic level. Building on the student's individual interests and life goals, it could also help them decide what to learn, as well as where and

---

[559] The University of Southern California have been researching just such an application over many years: http://ict.usc.edu/prototypes/personal-assistant-for-life-long-learning-pal3

[560] Alexa, Should We Trust You?

how to do the learning (the companion might identify and connect with the learning opportunities that are available, both formal and informal, both on and off-line). It could then also guide the student along overarching long-term individualised learning pathways designed to help the student address their emerging personal life-goals, connecting their learning interests and achievements, while reminding them of and encouraging them to reflect on and perhaps develop their long-term learning aims. The learning companion[561] might suggest learning opportunities that focus on some so-called 21st Century Skills,[562] and social-emotional learning.[563] It could also potentially connect learners, in the same classroom or from opposite sides of the world, depending on their shared interests and goals, helping them develop and work together in projects that prioritise both individual and collective achievements (and, in turn, helping to promote other critical skills in collaboration, teamwork, and intercultural awareness).

42.1.5 AI Teaching Assistant

As we have noted several times, most AIED technologies are designed with the aim of relieving teachers of the grunt work of teaching (most often by automating time-consuming activities such as the marking of classroom or homework assignments). However, despite these best of intentions, many AIED technologies in effect take over teaching (they deliver personalised and adapted learning activities better than teachers), or at least they reduce teachers to a functional role (perhaps their job is to work to strictly regulated scripts, or to ensure that the technology is ready for the student to use). Nonetheless, as we and col-

---

[561] World Economic Forum. 2015. New Vision for Education: Unlocking the Potential of Technology. World Economic Forum.

[562] Trilling, B. and Fadel, C. 2012. 21st Century Skills: Learning for Life in Our Times. John Wiley & Sons.

[563] Fadel, C., Bialik, M., and Trilling, B. 2015. Four-Dimensional Education: The Competencies Learners Need to Succeed. Centre for Curriculum Redesign.

leagues have written previously: Crucially we do not see a future in which AIED replaces teachers. What we do see is a future in which the role of the teacher continues to evolve and is eventually transformed; one where their time is used more effectively and efficiently, and where their expertise is better deployed, leveraged, and augmented.[564]

This might be more of an emotional plea than a coherent argument—but it assumes that teaching involves more than delivering knowledge, and that it is a fundamentally social process. From this perspective, a key role for AI is supporting teachers to teach and support students.

One way in which this might be achieved is by augmenting teachers' expertise and skills with an AI teaching assistant, to complement and work with the students' AI learning companion, that goes far beyond the useful but by comparison somewhat primitive teacher dashboards featured in so much education technology. This would be a key way that AIED can support teachers to support students. Just such a possibility has been explored in the short narrative "A.I. is the New T.A in the Classroom,"[565] which describes a possible classroom of the future in which the teacher is supported by a dedicated and personalised AI teaching assistant (AI TA).

Many of the ideas we have suggested could play a role in this possible scenario (such as automatically setting up collaborative groupings of students, replacing stop-and-test examinations with AI-supported continuous assessment, and managing peer-marking and undertaking some automated marking). The AI TA could also automatically provide teaching and professional development resources (texts, images, videos, augmented-reality animations, links, network connections) that the teacher

---

564 Luckin. R., et al. Intelligence Unleashed, 11. https://www.pearson.com/content/dam/one-dot-com/one-dot-com/global/Files/about-pearson/innovation/Intelligence-Unleashed-Publication.pdf.

565 Luckin, R., and Holmes, W. 2017. "A.I. is the new T.A. in the classroom." How We Get To Next. https://howwegettonext.com/a-i-is-the-new-t-a-in-the-classroom-dedbe5b99e9e

might choose to call upon to support their teaching. It could also monitor the students' performance as they engage in their classroom activities, continuously updating their learner models, making connections with the domain models of topics being taught, and tracking progress over time. All of this information (together with data about each student from additional sources: assessments from other classes, informal learning achievements, and relevant medical or family information) could be readily available to the teacher, whenever the AI TA computes it might be useful or whenever the teacher calls for it. In this possible future, what and how to teach the students, and how best to support them, would remain the responsibility and prerogative of the teacher. The AI TA's role would simply be to make the teacher's job easier and more effective.

42.1.6 AIED: a Research Tool to Further the Learning Sciences

As has probably been noticed, each of AIED's possible future uses are firmly rooted in existing AIED research and approaches. This is no less true of our final example, the use of AIED as a research tool to further the learning sciences. Implementing an educational practice in any technology means that the practice has to be both better understood and then systemized. As a consequence, the technology acts much like a virtual spotlight, highlighting issues that have existed for years but that have been hidden or overlooked (for example, around the most effective approaches to teaching). This is particularly true of the introduction of AI to education, which is beginning to throw an extraordinarily bright spotlight onto many learning sciences issues. However, while there have been notable developments in this area of AIED research, mostly it has been at a relatively theoretical level, such that their potential and implications remain somewhat unclear.

In fact, AIED as a learning sciences research tool is often linked to a pair of other independent but overlapping academic fields that use statistical techniques drawn from big data research:[566] learning analytics and educational data mining.[567] While learning analytics involves "the measurement, collection, analysis and reporting of data about learners and their contexts, for purposes of understanding and optimising learning and the environments in which it occurs;"[568] educational data mining "is concerned with gathering and analysing data so as to understand, support and improve students' learning."[569] One example, that avoids this distinction, and that has been shown to be effective, is The Open University's OU Analyse[570] tool, which draws on data from across the university (such as student access of online learning materials, submission of assessments, and outcomes) to identify students who might be at risk of dropping out from their studies—to enable tutors and student-support staff to provide appropriate pro-active remedial support. In fact, with the fields continually informing and cross-fertilizing each other, the distinctions between learning analytics, educational data mining, and AIED as

---

566 Mayer-Schonberger, V. and Cukier, K. 2013. Big Data: A Revolution That Will Transform How We Live, Work and Think. John Murray.

567 Readers who would like to learn more about the similarities and differences between learning analytics and educational data mining might be interested to read Benedict du Boulay and others, "What does the research say about how artificial intelligence and big data can close the achievement gap?" in Luckin, R. (ed.) 2018. Enhancing Learning and Teaching with Technology. Institute of Education Press, 316–27; or Siemens, G., and Baker, R.S.J.d.. 2012. "Learning analytics and educational data mining: Towards communication and collaboration." In Proceedings of the 2nd International Conference on Learning Analytics and Knowledge, 252–254. http://dl.acm.org/citation.cfm?id=2330661

568 Siemens, G. 2011. "1st International conference on learning analytics and knowledge 2011: Connecting the technical, pedagogical, and social dimensions of learning analytics. https://tekri.athabascau.ca/analytics/about

569 Du Boulay, et al., "What does the research say about how artificial intelligence and big data can close the achievement gap?" 270.

570 See Herodotou, C., et al. 2017. "Predictive modelling for addressing students' attrition in higher education: The case of OU analyse." http://oro.open.ac.uk/49470/and https://analyse.kmi.open.ac.uk

a learning sciences research tool are becoming increasingly blurry. Often, it simply comes down to the communities who are involved in the research and the terminology that they use. Here, as we are writing about AIED, we will continue to use AIED terminology.

One prominent example of AIED as a learning sciences research tool has recently been published by the Medical Research Council Cognition and Brain Sciences Unit at the University of Cambridge.[571] The traditional grouping of students with learning difficulties in broad categories such as ADHD, dyslexia, and autism has long been known to be insufficiently helpful, when educators try to improve learning outcomes for individuals. For this reason, the Cambridge researchers are investigating the use of machine learning to categorise struggling students at a more granular level (based on measures of listening skills, spatial reasoning, problem solving, vocabulary, and memory). By analysing data from more than 500 children, the machine learning revealed four clusters of learning difficulties (which had not previously been so clearly delineated): difficulties with working memory skills, difficulties with processing sounds in words, broad cognitive difficulties in many areas, and typical cognitive test results for the student's age. The researchers found that diagnosing struggling learners in terms of these four clusters was both more accurate and more useful, helping educators address individual learning difficulties, than the traditional diagnostic labels.

We will conclude our brief discussion of AIED as a learning sciences research tool with one final example, one that is in the early stages but has important potential: the use of machine learning to improve learning design. Learning design refers to a range of methodologies "for enabling teachers/designers to make more informed decisions in how they go

[571] See Astle, D.E., Bathelt, J. and Holmes, J. 2018. Remapping the cognitive and neural profiles of children who struggle at school." Developmental Science. https://doi.org/10.1111/desc.12747 and, for a short summary, https://www.opencolleges.edu.au/informed/learning-strategies/artificial-intelligence-identifies-students-struggle-school

about designing learning activities and interventions."[572] These methods are intended to inform decisions about pedagogy (teaching and learning) and about ways to support student learning experiences, and can also be used to provide core data for learning analytics or educational data mining. Most approaches in use in universities[573] draw on teachers' professional knowledge of teaching and learning (knowledge that is often tacit and thus has had to be elicited from them, which is a non-trivial task and can lead to fuzziness and inconsistencies). Instead, the approach currently being researched at the Open University involves machine learning from thousands of existing module activities to identify categories of activities at a highly granular level. Once these learning design activity categories are identified, and have been robustly authenticated, it should then be possible to correlate the actual learning designs of course modules with student outcomes, to help us better understand how students learn. In turn, this might inform teachers and learning designers about which learning designs (depending on, for example, domain, specific subject, duration and level of study) are most effective in practice.

42.2 AI in Education—A Tentative Summary

In the previous sections we have discussed a wide variety of existing and potential AIED technologies. One way to access this variety is to consider the technologies in terms of whether they are mainly student teaching (they take a mainly instructionist approach), or student support-

---

[572] Conole, G. 2012. Designing for Learning in an Open World (v. 4). Springer Science & Business Media.

[573] E.g., Cross, S., et al. 2012. "OULDI-JISC project evaluation report: The impact of new curriculum design tools and approaches on institutional process and design cultures." http://oro.open.ac.uk/34140/; Laurillard, D., et al. 2013. "A constructionist learning environment for teachers to model learning designs." Journal of Computer Assisted Learning 29 (1): 15–30; Dalziel, J. (ed.), Learning Design. Routledge.

ing (they take a mainly constructivist approach), or teacher supporting (they mainly help teachers do what they already do but more quickly or with less effort). A summary representation of this is shown in the following table. A cursory examination of this table will reveal that the categorization provides only a high-level overview, while many of the AIED approaches overlap, and most of the technologies could easily appear in another place in the table. It is also likely that over time different AIED technologies will merge into multi-capable systems, perhaps incorporating sequenced (ITS), Socratic (DBTS), and self-directed (ELE) learning in one technology.[574]

	Student Teaching (mainly instructionist)	**Student Supporting (mainly constructivist)**	**Teacher Supporting**
AIED Applications	• ITS • DBTS • Language learning apps	• ELEs • Automatic writing evaluation (formative) • Learning network orchestrators • Language learning apps • AI Collaborative learning • AI Continuous assessment • AI Learning companions	• ITS+ • Automatic writing evaluation (summative) • Student forum monitoring • AI Teaching Assistants • AI as a research tool to further the learning sciences
AIED Technologies and Approaches		• Chatbots • AR and VR • Natural Language Processing • Adaptivity	

Student teaching, student supporting, and teacher supporting AIED.

This summary is given more flesh in the following table, Characteristics of AIED Technologies.

574 Early examples of this include Holmes, W. 2013. "Level up! A design-based investigation of a prototype digital game for children who are low-attaining in mathematics." (Unpublished PhD thesis, University of Oxford) and Rummel, N., et al. 2016. "Transforming learning, empowering learners." The International Conference of the Learning Sciences 1.

Type of AIED	Characteristics	Determined by	Target
Intelligent Tutoring Systems	• Step-by-step sequence of instruction and tasks. • Individualized pathways. • System-determined content and pathways. • Students working with computers (or mobile devices). • Individualized feedback. • Real-time adaptivity.	System	For students
Dialogue-based Tutoring Systems	• Step-by-step dialogue-based instruction and tasks. • Individualized conversations. • System-determined content and pathways. • Students working with computers (or mobile devices). • Individualized feedback. • Real-time adaptivity.	System	For students
Exploratory Learning Environments	• Exploratory tasks. • Individualized pathways. • System-determined content and pathways, with student choice within tasks. • Students working with computers (or mobile devices). • Individualized feedback. • Real-time adaptivity.	System and learner	For students
Automatic Feedback and Scoring of Essays	• Essays (and other assignments) uploaded and analyzed by the system. • Some provide individualized formative feedback (to help students improve their writing), some only summative assessment (to score/grade the essay).	System	For students (formative) For teachers (summative
ITS+	• Depends on the ITS+.		

	• Whole-school wraparound ITS. • Student data visible to teacher; superimposed above each student via augmented-reality glasses. • Back-end ITS functionality (AIED as a service) for other providers of EdTech products.	n/a	For students and teachers
Language Learning Apps	• Step-by-step sequence of instruction and tasks. • System-determined content and pathways. • Students working with computers (or mobile devices). • Individualized feedback.	System	For students
Chatbots	• Mostly providing information.	Student (i.e., responds to student questions)	For students
Augmented and Virtual Reality	• Mostly providing access to otherwise unavailable environments.	Mixed	For students
Learning Network Orchestrators	• Mostly providing access to learning opportunities.	Mixed (i.e., sometimes responds to student requests)	For students
Collaborative Learning	• Facilitating the organization of collaborative learning. • Facilitating collaborative learning.	System	For students
Student Forum Monitoring	• Providing automatic feedback to forum posts, perhaps making connections between posts and sentiment analysis.	n/a	For students, and for teachers
Continuous Assessment	• Assessing student competencies on an ongoing basis (e.g., during talk), rather than using tests or exams.	System	For students

Finally, we might compare all of the AIED technologies with the SAMR model discussed in the context section of this book. This highlights how most of the near and medium-term advantages of AIED are in the augmentation and modification of present-day activities, while the long term might see a substantial Holy Grail benefit in redefinition.

	EdTech at Large (using the SAMR model)	**AIED in Particular**
Redefinition	Technology allows for the creation of new tasks, previously inconceivable.	• AI removing the need for stop-and-test examinations (i.e. by providing continuous highly adaptive assessments).
Modification	Technology allows for significant task redesign.	• AR and VR learning experiences • AI Learning Companions • AI Teaching Assistants • AI as a Learning Sciences research tool.
Augmentation	Technology acts as a direct tool substitute, with functional improvement.	• ITS • DBTS • Exploratory Learning Environments • Automatic Writing Evaluation • ITS+ • Language Learning • Chatbots • Collaborative Learning support • Student forum monitoring
Substitution	Technology acts as a direct tool substitute, no functional change.	Not applicable (as of this writing)

AIED and the SAMR model.

42.3 The Social Consequences of AI in Education

As we have seen, the application of AI in educational contexts is growing rapidly. In this book, we have explored the various AI techniques being used, the applications that have been in development for almost fifty years, and the futuristic possibilities that are becoming ever more likely (whatever our personal values).

Clearly, AIED has achieved some notable successes, while the conceivable applications are at the least intriguing. However, AIED's potential impact on students, teachers and wider society is yet to be fully worked out. This is true of issues as broad as accuracy, choice, predictions, privacy, teachers' jobs, and what we should be teaching school and university students.[575] But it is especially true for AIED's emerging ethical questions: "Around the world, virtually no research has been undertaken, no guidelines have been provided, no policies have been developed, and no regulations have been enacted to address the specific ethical issues raised by the use of artificial intelligence in education."[576]

In any case, one wonders why, if AIED is so effective, has it not yet been widely adopted by schools, universities and training companies? In fact, it is not yet even clear whether the AI technologies being imported into education are actually up to the task. For many years, non-AI technologies in educational settings have been critiqued. The question is whether AIED is destined to become the latest computer technology to

---

[575] E.g., "Machine learning: universities ready students for AI revolution," https://www-timeshighereducation-com.libezproxy.open.ac.uk/news/broader-four-year-degrees-offered-in-response-to-ai-revolution and "The most important skills for the 4th industrial revolution? Try ethics and philosophy." https://www.edsurge.com/news/2018-10-06-the-most-important-skills-for-the-4th-industrial-revolution-try-ethics-and-philosophy

[576] Holmes, W., et al. 2018. "Ethics in AIED: Who cares?" In: Artificial Intelligence in Education (ed. Rosé, C.P., et al.). 19th International Conference Proceedings, Part II. https://doi.org/10.1007/978-3-319-93846-2

be oversold yet underused in classrooms.[577, 578] We also need to consider what might happen, what might be the impact on individual learners, if ineffective AI techniques (or biased data sets) are used in classrooms (for example, what might happen if the face recognition technology that achieved 95% false positives for the UK's Metropolitan Police[579] was used in classroom monitoring)? Meanwhile, there are few examples of cumulative or replicable AIED research: the field is developing so rapidly while AIED data sets and algorithms tend to be (jealousy?) guarded. There is also little available robust evidence of the efficacy at scale of the rapidly increasing numbers of AIED tools. Even those, such as Mathia and Assistments, that do have some evidence, have typically been compared with business as usual rather than with another technology that has at least some level of comparability.[580] The purported effectiveness of many other tools may be due to their novelty in classrooms,[581] rather than anything to do with the AI employed—we simply do not have the evidence to say one way or another.

---

[577] Cuban, L. 2001. Oversold and Underused: Computers in the Classroom. Harvard University Press.

[578] "Pretty much all edtech sucks. And machine learning is not going to improve edtech."—Al Essa, McGraw-Hill Education; and "I don't see a child sitting in front of an Alexa and being taught, because there is a whole other set of cues they need to learn. I don't see machine learning reaching that point."—Janel Grant. Both quoted in Johnson, S. 2018. "What can machine learning really predict in education?" EdSurge. https://www.edsurge.com/news/2018-09-26-what-can-machine-learning-really-predict-in-education

[579] The Independent, May 2018. https://ind.pn/2InMfGf

[580] Holmes, W., et al. Technology-Enhanced Personalised Learning, 65 and 68.

[581] Schomaker, J. and Meeter, M. 2015. "Short- and long-lasting consequences of novelty, deviance and surprise on brain and cognition." Neuroscience & Biobehavioural Reviews. https://doi.org/10.1016/j.neubiorev.2015.05.002

42.3.1 The Implications of AIED Technologies for Classrooms

We began our AIED journey with intelligent tutoring systems, which as we saw are the most common of AIED applications, and which we will now use to scaffold and highlight some social consequences of AI applied to education that deserve more detailed attention. It has long been recognized that AI by design amplifies hidden features of its initial data and effectively reinforces its underlying assumptions. In particular, if the algorithms "are trained on data which contains human bias then of course the algorithms will learn it, but furthermore they are likely to amplify it. This is a huge problem, especially if people assume that algorithms are impartial."[582] In this respect, both rule-based and machine learning ITSs are no different. Their very design, their implementation of step-by-step instructionist methods focused on a knowledge curriculum while ignoring contextual and social factors, amplifies existing yet contested assumptions about effective approaches to teaching, and even to what it means to learn.[583]

ITSs also embody a usually unacknowledged paradox, the dependence of personalised approaches to learning on identifying what is collective or average. "[ITS] recommends lessons to users based on how other learners on the system have performed. These systems "learn" each student by presuming them to be similar to others.... We herald an intervention as a success if [an efficacy study shows that] it works on average, discarding the nuances of why it may work for some students more

---

[582] Douglas, L. 2017. "AI is not just learning our biases; it is amplifying them." Medium. https://medium.com/@laurahelendouglas/ai-is-not-just-learning-our-biases-it-is-amplifying-them-4d0dee75931d

[583] Instructionism "is based on cognitive learning theories that centre on teaching as education performed by a teacher. In the view of instructionism, instruction has to be improved in order to achieve better learning results." Seel, N.M., ed. 2012. Encyclopedia of the Sciences of Learning. Springer.

than others, and to what degree. [In summary], the individual struggle of the individual learner is easily lost in the noise."[584]

In other words, focusing on the average to determine an appropriate intervention is inevitably limiting: if a robust study shows that one approach is more effective on average compared with a second approach, the second approach is likely to be fully rejected, despite the fact that it might be more effective for particular individuals or groups.

ITSs by design also can reduce student agency. Although constrained by the curriculum (as decided by local or national policy-makers), it is generally the ITS (its algorithms and student models) and, at a higher level, the ITS designers, that determine what should be learned, in what order and how; while the student is given little choice but to follow the ITS-determined individual pathway (it also in some sense makes the teacher somewhat redundant—it is the system, not the teacher, that decides what is best for a student to learn). For example, most ITS begin with the basics, before guiding the individual student through tasks that take them step-by-step towards mastery targets, minimising failure along the way. However intuitively appealing, the assumptions embodied in this instructionist approach[585] also ignore the value of other approaches researched in the learning sciences (such as collaborative learning, guided discovery learning, blended learning, and productive failure).[586]

ITSs also raise issues centred on the selection of data, raising complex issues centred on trust.[587] For example, it has been argued that there

---

[584] Mubeen, J. 2018. "When 'personalised learning' forgets to be 'personalised."' Medium. https://medium.com/@fjmubeen/when-personalised-learning-forgets-to-be-personalised-48c3558e7425

[585] Gagné, Conditions of Learning and Theory of Instruction.

[586] Dean Jr., D. and Kuhn D. 2007. "Direct instruction vs. discovery: The long view." Science Education 91. https://doi.org/10.1002/sce.20194

[587] E.g., https://www.theatlantic.com/magazine/archive/2018/11/alexa-how-will-you-change-us/570844/

is no such thing as raw data:[588] data used in any analysis has been preselected (it is not possible to include all data generated by a system in its computations), and these choices are inevitably subject to conscious or unconscious, explicit or implicit, selection biases.[589] Similarly, the algorithms chosen or developed raise additional issues, such as those centred on the accuracy and implications of their predictions (if the computation is incorrect, are students being guided away from their best interests, and how do we ensure that mistakes err on the side of failing in the least harmful way?), the increasing focus on inferring and responding to the students' affective states (are a student's innermost feelings not private anymore?)[590] and the usual focus on teaching the type of knowledge that is the easiest to automate and thus potentially the least useful in the long-term for students.[591]

In any case, as we discussed earlier, the efficacy of ITSs in real educational settings remains to be confirmed (although many have been shown to be broadly effective when compared against usual classroom teaching).[592] Indeed, one ITS, Summit Learning,[593] which was developed by engineers from Facebook and is being used in around 400 schools, has been the focus of student protests and boycotts.

---

[588] Gitelman, L., et al. 2013. "Raw *Data" Is an Oxymoron*. MIT Press.

[589] "Data is easily obtained, but it has a lot of bias in it." John Behrens (Pearson), quoted in Johnson, What Can Machine Learning Really Predict in Education? https://www.edsurge.com/news/2018-09-26-what-can-machine-learning-really-predict-in-education

[590] "Tech firms want to detect your emotions and expressions, but people don't like it." https://theconversation.com/tech-firms-want-to-detect-your-emotions-and-expressions-but-people-dont-like-it-80153

[591] Rose Luckin quoted in https://www.jisc.ac.uk/news/the-ai-revolution-is-here-17-aug-2018

[592] Du Boulay, B. "Artificial intelligence as an effective classroom assistant." IEEE Intelligent Systems 31. https://doi.org/10.1109/MIS.2016.93

[593] https://www.summitlearning.org

"Unfortunately we didn't have a good experience using the program, which requires hours of classroom time sitting in front of computers... The assignments are boring, and it's too easy to pass and even cheat on the assessments. Students feel as if they are not learning anything and that the program isn't preparing them for the Regents exams they need to pass to graduate. Most importantly, the entire program eliminates much of the human interaction, teacher support, and discussion and debate with our peers that we need in order to improve our critical thinking. Unlike the claims made in your promotional materials, we students find that we are learning very little to nothing. It's severely damaged our education, and that's why we walked out in protest."[594]

Finally, ITSs typically set themselves up as doing at least some of the job of teachers, increasingly more effectively than teachers, thus questioning the role of teachers in future classrooms.[595] As we have seen, the ambition of many researchers is to relieve teachers of the burdens of teaching (such as monitoring progress and marking assignments), enabling them to focus on the human aspects of teaching (such as social engagement). In fact, "AI cannot create, conceptualise, or manage complex strategic planning; cannot accomplish complex work that requires precise hand-eye coordination; cannot deal with unknown and unstructured spaces, especially ones that it hasn't observed; and cannot, unlike humans, feel or interact with empathy and compassion… tasks that can only be done by a human teacher. As such, there will still be a great need

---

[594] The Chan Zuckerberg Initiative funded the Summit Learning project and disputes these claims. https://www.washingtonpost.com/education/2018/11/17/students-protest-zuckerberg-backed-digital-learning-program-ask-him-what-gives-you-this-right/?noredirect=on&utm_term=.27d5e322ac1c

[595] At least one ITS company appeared to pivot from attempting to sell their product into schools, because teachers were unsure why they should use a technology that did their job instead of them.

for human educators in the future."[596] But, on the other hand, if we (students, educators, and parents) do not critically engage, perhaps AIED might lead to fast-food chef, script-driven classroom managers[597] rather than teachers, while the AI deals with all of the cognitive demands of teaching (a dystopian scenario that is only some short steps away from removing humans from teaching entirely).

Naturally, there are many examples of ITSs that challenge at least some of these issues (such as Mathia, whose developers recommend it is delivered in a blended context). We have also looked at alternative approaches, such as DBTSs (that prioritise a Socratic, albeit step-by-step, approach rather than an instructionist approach to learning) and AI-driven ELEs (that prioritise a guided-discovery approach to learning). And we have considered alternative ways in which AI is being or might be used in innovative ways, that have the potential to step outside dominant educational practices: for example, relatively simple AI that enables students to connect to their choice of human tutors (to get support on what they want to learn), and complex AI that provides a lifetime learning companion dedicated to their needs. Yet even these approaches depend on huge amounts of personal data and efficient algorithms, raising privacy and ethical issues that have yet to be fully considered.

---

[596] https://www.linkedin.com/pulse/10-jobs-safe-ai-world-kai-fu-lee. Also see, "Intelligent machines will replace teachers within 10 years, leading public school head teacher predicts." https://www.independent.co.uk/news/education/education-news/intelligent-machines-replace-teachers-classroom-10-years-ai-robots-sir-anthony-sheldon-wellington-a7939931.html; "Could artificial intelligence replace our teachers?" https://www.educationworld.com/could-artificial-intelligence-replace-our-teachers; and "Why artificial intelligence will never replace teachers," https://www.thetechedvocate.org/artificial-intelligence-will-never-replace-teachers

[597] As we mentioned earlier, one ITS developer has suggested that the sophistication of their AIED means that teachers only need to play an auxiliary role, working like fast-food chefs ("KFC-like") to strictly regulated scripts.

42.3.2 The Ethics of AIED

Indeed, the ethics of AI applied in education, although left to last in this book, requires urgent attention. For example, one school has installed facial recognition technology to monitor how attentive students are in class. Every movement of pupils ... is watched by three cameras positioned above the blackboard.... Some students are already changing their behaviour due to the increased monitoring....

"I don't dare be distracted after the cameras were installed in the classrooms. It's like a pair of mystery eyes are constantly watching me." The system works by identifying different facial expressions from the students, and that information is then fed into a computer which assesses if they are enjoying lessons or if their minds are wandering.... The computer will pick up seven different emotions, including neutral, happy, sad, disappointed, angry, scared and surprised. If it concludes that the student is distracted with other thoughts during the class, it will send a notification to the teacher to take action."[598]

This example of AI being used to maximise student attention is from China. However, before we dismiss it as a culturally-specific phenomenon, we should remember that ALT Schools[180] also uses AI-driven classroom cameras to monitor student behaviour (while in the UK, "tens of thousands of pupils aged as young as five are at risk of being spied on through their webcams..., often without students or their parents ever knowing").[599] This is not to say that the use of AI to analyse classroom video feeds is by definition unethical. For example, researchers at the University of Pittsburgh are using AI and classroom videos to help better understand how the quality of classroom talk, the liveliness of discus-

---

[598] Connor, N. 2018. Chinese school uses facial recognition to monitor student attention in class. Telegraph, https://www.telegraph.co.uk/news/2018/05/17/chinese-school-uses-facial-recognition-monitor-student-attention/

[599] 2018. Children Young 5 Risk Spied Webcams Using School Software, Telegraph, https://telegraph.co.uk/technology/2018/12/15/children-young-5-risk-spied-webcams-using-school-software

sion, and the level of student engagement contributes to effective learning, to inform better approaches to teaching.[600]

On the other hand, there are examples of AI companies[601] collecting huge amounts of student interaction data, in order to use machine-learning techniques to "search for patterns." The aim, naturally, is to "improve student learning by teaching the software to pinpoint when children are feeling happy, bored, or engaged."[602] Nonetheless, this approach is controversial, with such data collection being characterized as "borderline mental-health assessments..., [that] encourage a view of children as potential patients in need of treatment."[603]

The reality is that, while the range of AI techniques and technologies researched in classrooms and discussed at conferences are extensive and growing, the ethical consequences are rarely fully considered (at least, there is very little published work considering the ethics). In fact, most AIED research, development, and deployment has taken place in what is essentially a moral vacuum (for example, what happens if a child is unknowingly subjected to a biased set of algorithms that impact negatively and incorrectly on their school progress?). In particular, AIED researchers are working without any fully worked out moral groundings.

In fact, as we have seen, AIED techniques raise an indeterminate number of self-evident but as yet unanswered ethical questions. To begin with, as with mainstream AI, concerns exist about the large volumes of data collected to support AIED—albeit data that is collected with the

---

[600] Kelly, S., Olney, A.M., Donnelly, P., Nystrand, M., and D'Mello. S.K. (2018). "Automatically measuring question authenticity in real-world classrooms." Educational Researcher 47. https://doi.org/10.3102/0013189X18785613

[601] E.g., https://www.algebranation.com

[602] "How (and why) ed-tech companies are tracking students' feelings." https://mobile.edweek.org/c.jsp?cid=25919761&bcid=25919761&rssid=25919751&item=http%3A%2F%2Fapi.edweek.org%2Fv1%2Few%2Findex.html%3Fuid=C08929D8-6E6F-11E8-BE8B-7F0EB4743667

[603] Jane Robbins, American Principles Project Foundation, quoted in preceding note, "How (and why) ed-tech companies are tracking students' feelings."

best of intentions (such as the recording of student competencies, emotions, strategies, misconceptions, and screen usage,[604] to better help students learn). Who owns and who is able to access this data, what are the privacy concerns, how should the data be analysed, interpreted, and shared, and who should be considered responsible if something goes wrong? In a parallel domain, healthcare, the use of personal data can be contentious and is frequently challenged[605] —but this has yet to happen noticeably in education.

However, while data raises major ethical concerns for the field of AIED, AIED ethics cannot be reduced to questions about data. Other major ethical concerns include the potential for bias[606] (conscious or unconscious) incorporated into AI algorithms (i.e., how the data is analysed)[607] and into AIED models (what aspects of a domain are assumed

---

[604] "FaceMetrics lands $2 million to gamify kids' screen time and track immersion with AI." https://venturebeat.com/2018/06/13/facemetrics-lands-2-million-to-gamify-kids-screen-time-and-track-immersion-with-ai

[605] For example, https://www.bbc.co.uk/news/technology-46206677: "A controversial health app developed by artificial intelligence firm DeepMind will be taken over by Google …" Lawyer and privacy expert Julia Powles [said]: "DeepMind repeatedly, unconditionally promised to 'never connect people's intimate, identifiable health data to Google.' Now it's announced... exactly that. This isn't transparency, it's trust demolition."

[606] "[A]s algorithms play an increasingly widespread role in society, automating—or at least influencing —decisions that impact whether someone gets a job or how someone perceives her identity, some researchers and product developers are raising alarms that data-powered products are not nearly as neutral as scientific rhetoric leads us to believe." Kathryn Hume, integrate.ai, quoted in "AI needs debate about potential bias," by Carole Piovesan, https://www.lawtimesnews.com/article/ai-needs-debate-about-potential-bias-15180. Also see, The Fairness Toolkit, https://unbias.wp.horizon.ac.uk/fairness-toolkit

[607] A recent survey by The Pew Research centre found that "the public is frequently sceptical of [algorithms] when used in various real-life situations. … [with] 58% of Americans feel[ing] that computer programs will always reflect

worth learning, what approaches to pedagogy are assumed to be most effective, and what student information is assumed to be the most pertinent?). On the other hand, if a computer's decisions are indistinguishable from that of a human, or at least from a panel of human experts (because humans are well known to sometimes disagree, for example when marking essays),[608] perhaps those decisions should be accepted.[609] Nonetheless, each decision that goes into constructing these algorithms and models might impact negatively on the human rights of individual students (in terms of gender, age, race, socio-economic status, income inequality, and so on)—at present we just do not know whether or not they will.

But these particular AI ethical concerns, centred on data and bias, are the "known unknowns," and are the subject of much research and discussion in mainstream AI research.[610] What about the "unknown un-

---

some level of human bias." http://www.pewinternet.org/2018/11/16/public-attitudes-toward-computer-algorithms/

[608] To give an anecdotal example, a Master's thesis written by one of the authors at a prestigious university was marked as a distinction by one professor and a fail by another.

[609] From another perspective, the UCLA law professor Eugene Volokh argues that "a computer should be accepted if a panel of humans thinks the opinions it writes are on par with or better than those written by a human judge…" (https://www.axios.com/artificial-intelligence-judges-0ca9d45f-f7d3-43cd-bf03-8bf2486cff36.html)

[610] E.g., Ada Lovelace Institute (https://www.adalovelaceinstitute.org), AI Ethics Initiative (https://aiethicsinitiative.org), AI Ethics Lab (http://www.aiethicslab.com), AI Now (https://ainowinstitute.org), DeepMind Ethics and Society (https://deepmind.com/applied/deepmind-ethics-society), and the Oxford Internet Institute (https://www.oii.ox.ac.uk/blog/can-we-teach-morality-to-machines-three-perspectives-on-ethics-for-artificial-intelligence).
Also see Winfield, Alan F. T., and Jirotka, M. (2018). "Ethical governance is essential to building trust in robotics and artificial intelligence systems." Phil. Trans. R. Soc. A 376. https://doi.org/10.1098/rsta.2018.0085 And see "Top 9 ethical issues in artificial intelligence." https://www.weforum.org/agenda/2016/10/top-10-ethical-issues-in-artificial-intelligence "Establishing an AI code

knowns," those ethical issues raised by the application of AI in education that have yet to be even identified?

AIED ethical questions include (there are many more):

- What are the criteria for ethically acceptable AIED?
- How does the transient nature of student goals, interests and emotions impact on the ethics of AIED?
- What are the AIED ethical obligations of private organisations (developers of AIED products) and public authorities (schools and universities involved in AIED research)?
- How might schools, students, and teachers opt out from, or challenge, how they are represented in large datasets?
- What are the ethical implications of not being able to easily interrogate how

AIED deep decisions (using multi-level neural networks) are made?

Strategies are also needed for ameliorating risk, since AI algorithms are vulnerable to hacking and manipulation (as the Facebook–Cambridge Analytica data scandal showed was more than possible): "It's impossible to have personal privacy and control at scale, so it is critical that the uses to which data will be put are ethical – and that the ethical guidelines are clearly understood."[611] Where AIED interventions target behavioural change (such as by nudging individuals towards a particular behaviour or course of action), the entire sequence of AIED-enhanced pedagogical activity also needs to be ethically warranted. Finally, it is important to recognize another perspective on AIED ethical questions: in

---

of ethics will be harder than people think." https://www.technologyreview.com/s/612318/establishing-an-ai-code-of-ethics-will-be-harder-than-people-think, and Willson, M. (2018). "Raising the ideal child? Algorithms, quantification and prediction." Media, Culture & Society, 5. https://doi.org/10.1177/0163443718798901

[611] Tarran, B. (2018). "What can we learn from the Facebook–Cambridge Analytica scandal?" Significance 15 (3): 4–5.

each instance, the ethical cost of inaction and failure to innovate must be balanced against the potential for AIED innovation to result in real benefits for learners, educators, and educational institutions.

42.3.3 In short, the ethics of AIED is complicated

As is likely already clear, the authors of this book are excited by what AI has to offer teaching and learning ... but we are also very cautious. We have seen an extraordinary range of AIED approaches (from Mathia, AutoTutor and Betty's Brain, to the Ada chatbot, OpenEssayist, and Lumilo, and more) and some amazing future AIED possibilities (from the end of exams, to lifelong learning companions, and AI teaching assistants). However, we have also identified a range of critical issues that need to be addressed before AI becomes an acceptable integral part of everyday learning.

Most importantly, the ethics of AIED need to be fully worked out—a non-trivial task that requires the involvement of a wide range of stakeholders (from students to philosophers, teachers to policymakers, and parents to developers). We (teachers, policymakers, and learning scientists) need to understand the key issues raised by the collection of data (such as the choice of what data to collect and what data to ignore, the ownership of data, and data privacy). We also need to understand the computational approaches being applied (what decisions are being made, what biases are creeping in, and how do we ensure that decisions are 'correct and transparent?).[612] This much is self-evident, which is why so many initiatives to both determine and govern the ethics of AI have been established around the world.

However, we also need to have a thorough understanding of the ethics of education, of teaching and learning (the ethics of particular approaches, curriculum choices, focusing on averages, the allocation of

---

[612] See, Miller, T. (2019). "Explanation in artificial intelligence: Insights from the social sciences." Artificial Intelligence 267. https://doi.org/10.1016/j.artint.2018.07.007

available funds, and much more besides), another non-trivial task. For without that, how will we know what might happen when these three areas (data, computation, and education) collide?

This returns us to our introduction, and is hopefully our main takeaway. Whether we welcome it or not, AI is increasingly being used widely across education and learning contexts. We can either leave it to others—the computer scientists, AI engineers[613] and big tech companies—to decide how artificial intelligence in education unfolds, or we can engage in productive dialogue. It is up to each of us to decide whether we acquiesce, take what we are given, or whether we adopt a critical stance, to help ensure that the introduction of AI into education reaches its potential and has positive outcomes for all.

[613] "You and AI–machine learning, bias and implications for inequality." https://royalsociety.org/science-events-and-lectures/2018/07/you-and-ai-equality.