




# Wayne Holmes

# AI and Education: A Critical Studies Perspective

# Utilizing AI in Developing Education Systems



Artificial Intelligence (AI) is among us. Although we might not notice it, increasingly every day we engage with AI in almost every aspect of our daily lives: whether we are using our mobile phones, reading an article in a newspaper, buying something online, checking the weather forecast, or playing a computer game. AI is also increasingly becoming embedded in education.

In fact, it is easy to find many enthusiasts (including many international organisations and well-funded corporations) who will tell you all about how AI will revolutionise education. (e.g. OECD, 2021)

“Many of the tasks that we have allowed ourselves to believe can be done uniquely by the teacher, can in fact be as well, or better carried out by AI machines.... AI is an altogether new way of spreading quality education across the world.” (Seldon & Abidoye, 2018, p. 4)

It is true that there are some extraordinary possibilities. However, there are also very many challenges, over-promises, and hyperbole, which together require the critical studies perspective adopted in this paper.

## ARTIFICIAL INTELLIGENCE

In order to facilitate a full understanding of what is meant by AI and education and its implications, we first need to understand what we mean by AI itself. UNICEF, building on a definition agreed by OECD member states, offer the following helpful summary:

“AI refers to machine-based systems that can, given a set of human-defined objectives, make predictions, recommendations, or decisions that influence real or virtual environments. AI systems interact with us and act on our environment, either directly or indirectly. Often, they appear to operate autonomously, and can adapt their behaviour by learning about the context.” (UNICEF, 2021)

‘Artificial Intelligence’ in one form or another has in fact been with us since the mid1950s, with the name being coined at a now infamous workshop at Dartmouth College (McCarthy et al., 1956). Ever since, AI has experienced periods of heightened interest and exciting predictions, punctuated by periods known as AI winters, when the predictions failed and the funding all but dried up.

Today's dominant version of AI is Machine Learning (ML), which uses algorithms to analyse data and make predictions. It has come about mainly because of the arrival of fast computer processors and the availability of huge amounts of data. Today, it has been estimated, the world creates around 3,000,000,000,000,000 bits of data every second,<sup>1</sup> thanks to the profusion of electronic devices (from mobile phones to health trackers) and the ubiquity of the Internet, while even today's pocket computers (i.e. mobile phones) are more powerful than supercomputers were only 40 years ago.<sup>2</sup> ML has led to the innumerable dramatic achievements of AI in recent years (such as automatic translation between languages<sup>3</sup> and figuring out what shapes proteins fold into).<sup>4</sup> Nonetheless, interestingly, some researchers now argue that ML is soon to hit its own development ceiling, such that significant further progress will only happen if there is a new paradigm (Marcus, 2020).

Although the achievements of ML are impressive, like many cutting-edge technologies before it, AI suffers from over-selling, hyperbole, myths, and fundamental limitations. For example, ML-driven AI systems can be brittle (e.g. a small change to a road sign can prevent an AI system recognising it, Heaven, 2019). They can also be biased (because the data on which they are trained, or the algorithms that they use, are biased, Buolamwini & Gebru, 2018). In any case, only in very specific and narrow applications can AI be said to be "better than humans".<sup>5</sup> Even the recently launched impressive AI language models, such as GPT-3,<sup>6</sup> which can write text in response to a prompt (Romero, 2021), are not all they appear: because they don't understand the meaning of the text that they are processing, they often write nonsense (Marcus & Davis, 2020). As a final example of AI's limitations, AI approaches made disappointingly little impact on addressing the COVID-19 pandemic (Benaich, 2020; Roberts et al., 2021; Walleiser, 2021).

It has been suggested that AI's hyperbole and myths can be traced back to the language used (Rehak, 2021). For example, the early decision to call the field Artificial 'Intelligence' effectively presupposes that the creation of a non-human intelligence is even possible. This anthropomorphism of AI also inevitably suggests agency (for example, 'learning' requires someone or something that 'learns'). Accordingly, perhaps a first step towards demythologising AI would be to stop using human descriptors to label mechanical processes. For example,

<sup>1</sup> <https://www.statista.com/statistics/871513/worldwide-data-created>

<sup>2</sup> <https://www.theclever.com/-15huge-supercomputers-that-were-less-powerful-than-your-smartphone>

<sup>3</sup> <https://ai.googleblog.com/09/2016/a-neural-network-for-machine.html>

<sup>4</sup> <https://deepmind.com/blog/article/alphafold-a-solution-to-a-50-year-old-grand-challenge-in-biology>

<sup>5</sup> In this example (<https://www.theguardian.com/global/2015/may/13/baidu-minwa-supercomputer-better-than-humans-recognising-images>) the systems, unlike humans, only had to get the answer correct once in five attempts.

<sup>6</sup> <https://openai.com/blog/openai-api/>

instead of AI ‘learning’, perhaps we should talk of the system’s configuration being improved, instead of ‘decision’ we should say ‘calculation’, and instead of ‘intelligent’ we should say ‘automated’. However, it is unlikely that we will see such changes anytime soon.

In any case, it is important to note that AI should not be thought of in purely technical terms. Instead, AI is a complex sociotechnical artefact that needs to be understood as something that is constructed through complex social processes (Eynon & Young, 2021). In other words, when we consider artificial intelligence, we must consider both the human dimensions and technical dimensions in tandem (Holmes et al., 2019).

## ARTIFICIAL INTELLIGENCE

Just as what constitutes AI and its implications are complex, the connections between AI and education are more involved than many acknowledge. There are also a lot of misunderstandings out there, and little robust research (Miao & Holmes, 2021). In order to help navigate this complexity, it can be useful to think of AI and education in terms of three “buckets” – the first bucket being “learning with AI”, the second “learning about AI”, and the third “preparing for AI.” As with any categorisation system, these buckets overlap in many ways – nonetheless, they still provide a helpful framing.

### LEARNING WITH AI

“Learning with AI”, using AI tools to support or facilitate teaching and learning, may be further categorised as institution-supporting, student-supporting, and teacher-supporting. Institution-supporting AI includes AI tools designed to help with student recruitment, timetabling, resource allocation, finances, and all the other unglamorous back-end administrative tasks that schools and universities all need to do (Zawacki-Richter et al., 2019). Accordingly, these tools actually have more in common with standard business applications of AI, than with teaching or learning.

Student-supporting AI, on the other hand, is where most of the excitement and money is. This has been researched for more than forty years and has now “escaped” from the lab to be offered by large numbers of million dollar-funded corporations around the world – such as Alef,<sup>7</sup> Assistments,<sup>8</sup> Byjus,<sup>9</sup> Domoscio,

<sup>7</sup> <https://alefeducation.com>

<sup>8</sup> <https://new.assistments.org>

<sup>9</sup> <https://byjus.com>

<sup>10</sup> to name just a few (from the beginning of the alphabet). These are all examples of the most prominent student-supporting AI, the so-called intelligent tutoring systems (ITS). With ITS, the student engages with an online system that delivers some information, an activity, and possibly a quiz. The student's individual responses (where they click and what they answer) then determines the next piece of information, activity, and quiz that they are given. In this way, each student follows their own adapted pathway through the material to be learned. Other student-supporting AI includes dialogue-based tutoring systems, which use a dialogic Socratic-approach to teaching and learning. Examples of this are AutoTutor (Graesser et al., 1999) and Watson Tutor.<sup>11</sup> Currently, almost all AI tutoring systems focus on well-defined and factual subjects such as mathematics, or other non-interpretative subjects like physics or computer science.

Other student-supporting “learning with AI” include automatic writing evaluation tools, examples of which (such as e-Rater<sup>12</sup>) are increasingly being used in standardised summative assessments, especially in the USA. However, because of the limitations of AI and the high-stakes nature of summative assessment, AI for formative assessment probably is likely to be more useful (Rolfe, 2011). Two final examples of student supporting “learning with AI” tools are network orchestrators and chatbots. Network orchestrators use AI technologies to establish connections between people (such as Smart Learning Partner from Beijing, that connects students safely with human tutors, so that they can receive support on something that they want to learn). Chatbots, on the other hand use natural language processing to provide on-demand responses to student queries (examples are Ada<sup>13</sup> and Genie<sup>14</sup>).

Finally, there's the third type of “learning with AI”, teacher-supporting AI, AI tools that do not replace teacher functions but aim to help teachers directly – think of an AI-driven virtual exoskeleton for teachers, helping them to do what they want to do but more easily and effectively. However, given that dashboards do not count, there are almost no examples of genuine teacher-supporting AI available today. The one example that does exist is automatic resource-curation: AI that scans the Internet in response to a teacher's query to identify quality resources that they can use in their teaching (e.g. X5GON<sup>15</sup> and Teacher Advisor<sup>16</sup>). However, these tools still have a way to go before they are really useful.

<sup>10</sup> <https://domoscio.com/solutions/adaptive-learning>

<sup>11</sup> <https://www.pearson.com/ped-blogs/blogs/11/2018/ai-based-tutoring-new-kind-personalized-learning.html>

<sup>12</sup> <https://www.ets.org/erater/about>

<sup>13</sup> <https://www.youtube.com/watch?v=xXQ2bxQrKuQ>

<sup>14</sup> <https://www.youtube.com/watch?v=ml0gdSCjGQ8>

<sup>15</sup> <https://platform.x5gon.org>

<sup>16</sup> <https://teacheradvisor.org>



## LEARNING ABOUT AI

The second bucket, “learning about AI”, is all about learning how AI works, the techniques and technologies, and how to create it. For years, universities across the world have offered degrees in a range of AI subjects (including topics such as data analytics, machine learning, neural networks, and so on), enabling many graduates to take up extremely well-paid jobs as data scientists or AI engineers. But what about schools? In recent months, UNESCO has been mapping AI school curricula from around the world, which revealed that the teaching of AI in schools remains relatively uncommon, although growing. However, it is possible to access an enormous variety of AI courses and tutorials via the Internet that anyone might choose to take, most often for free. One that is extremely popular, that introduces the technologies of AI and is available to all learners, is Elements of AI,<sup>17</sup> from the University of Helsinki. Another one, from France, is ClassCode<sup>18</sup>.

## PREPARING FOR AI

The second bucket, “learning about AI”, is all about learning how AI works, the techniques and technologies, and how to create it. For years, universities across the world have offered degrees in a range of AI subjects (including topics such as data analytics, machine learning, neural networks, and so on), enabling many graduates to take up extremely well-paid jobs as data scientists or AI engineers. But what about schools? In recent months, UNESCO has been mapping AI school curricula from around the world, which revealed that the teaching of AI in schools remains relatively uncommon, although growing. However, it is possible to access an enormous variety of AI courses and tutorials via the Internet that anyone might choose to take, most often for free. One that is extremely popular, that introduces the technologies of AI and is available to all learners, is Elements of AI, from the University of Helsinki. Another one, from France, is ClassCode.

To be fair, “preparing for AI” is not always ignored. The EU is currently developing the latest version of its DigComp Digital Competence Framework,<sup>19</sup> which specifically looks at the impact of AI on humans and the competences that all citizens should have to enable them to deal with the growing issues (this is due to be published in 2022). The Council of Europe is also investigating the impact of using AI in education on human rights, democracy, and the rule of law (which again is due to be published soon). In addition, the EU has established an expert group that aims to apply the EU’s Ethics Guidelines for Trustworthy AI<sup>20</sup> (which

<sup>17</sup> <https://www.elementsofai.com>

<sup>18</sup> <https://pixees.fr/classcode-v2/iai>

<sup>19</sup> <https://ec.europa.eu/jrc/en/digcomp/digital-competence-framework>

<sup>20</sup> <https://op.europa.eu/en/publication-detail/-/publication/d11-0434-3988569ea8-c1f01-aa75ed71a1>

address issues such as respect for human autonomy, prevention of harm, fairness, and explicability) to education and training.

These initiatives are important, but they are not well-known, and they tend to be limited to Europe, and citizens need to be made aware of them. Everyone also need opportunities to gain the AI knowledge, skills and attitudes recommended by DigComp 2.2, and to become familiar with the issues identified by the EU's expert group. In other words, preparing for AI means all citizens achieving some level of AI literacy. Just as education seeks to ensure that all citizens are literate in reading or mathematics, the argument is that we now need to ensure that all citizens understand the core benefits and challenges brought by AI. To date, teaching about AI has largely been the preserve of computer scientists. The resultant inevitable focus on the technology dimension of AI has tended to forget the human dimension – issues such as the impact on human well-being, respect for autonomy, protection of privacy, solidarity, democratic participation, equity, diversity, prudence, responsibility, and sustainable development (The Montréal Declaration for Responsible Development of Artificial Intelligence, 2018).

A key set of issues in which all citizens should be encouraged to engage is the ethics of AI. However, the ethics are complex, and so this is not so easy to achieve. In fact, the ethics of artificial intelligence in general has received a great deal of attention, by researchers (e.g. Boddington, 2017) and more widely (e.g., the House of Lords,<sup>21</sup> UNESCO,<sup>22</sup> World Economic Forum<sup>23</sup>), with numerous institutes for AI ethics being set up (e.g., the Ada Lovelace Institute,<sup>24</sup> the AI Ethics Initiative<sup>25</sup>, AI Now<sup>26</sup>, and DeepMind Ethics and Society<sup>27</sup>, to name just a few). In 2019, Jobin and colleagues identified 84 published sets of ethical principles for AI, which they concluded converged on five areas: transparency, justice and fairness, non-maleficence, responsibility, and privacy (Jobin et al., 2019). While what each of these mean and include, and how they may be applied both to the development or use of AI, remain subject to ongoing debates, they should surely be part of an AI curriculum.

<sup>21</sup> <https://publications.parliament.uk/pa/ld201719/ldselect/ldai/100/100.pdf>

<sup>22</sup> <https://unesdoc.unesco.org/ark:/48223/pf0000376713>

<sup>23</sup> [https://www3.weforum.org/docs/WEF\\_A\\_Holistic\\_Guide\\_to\\_Approaching\\_AI\\_Fairness\\_Education\\_in\\_Organizations\\_2021.pdf](https://www3.weforum.org/docs/WEF_A_Holistic_Guide_to_Approaching_AI_Fairness_Education_in_Organizations_2021.pdf)

<sup>24</sup> <https://www.adalovelaceinstitute.org>

<sup>25</sup> <https://aiethicsinitiative.org>

<sup>26</sup> <https://ainowinstitute.org>

<sup>27</sup> <https://deepmind.com/about/ethics-and-society>

## BEYOND THE HYPE

Building on this brief overview of the connections between AI and education, the intention now is to return to “student-supporting AI”, the focus of most of the excitement and the hype, to consider the claims made by the growing numbers of often multi-million-dollar-funded AI and education companies.<sup>28</sup>

To begin with, many of these companies claim that their AI education tools are intelligent (after all the word “intelligent” is right there in the title). But they're not. The reality is that no AI system today comes anywhere close to human intelligence: no AI truly understands anything, including as mentioned earlier the superficially impressive GPT-3, and no AI education tool is anywhere near as intelligent as a human teacher. In fact, as we have seen, AI systems in education are extremely limited in what they cover and can achieve: they are limited in the subjects that they address, what they can adapt to, and what explanations they can give. Sometimes, these AI in education tools might APPEAR intelligent, but that's a long way from actually BEING intelligent. It's all smoke and mirrors, and we mustn't allow ourselves to be fooled by the marketing.

Another claim is that these tools will save teacher time, a promise that has been made about educational technologies for almost a hundred years but has never actually materialised (Watters, 2021). Of course, AI, we are told, is different, more sophisticated, and so AI tools will finally save teacher time. Most teachers would love a tool that takes care of their marking, but this ignores the fact that no AI system is capable of the depth of interpretation or accuracy of analysis that a human teacher can give. It also ignores how much a teacher learns about a student when they read what the student has written, giving insights that no dashboard will ever give. So maybe AI will save teacher time, although there's little evidence for that, but at what cost to the quality of teaching and learning?

A third claim, perhaps the most alluring, is that the AI tools offer “personalised” learning, an approach that – so the argument goes – is better than anything a human teacher can provide (if they were not ‘better’, why would we spend money on them?). This is something we hear about all the time. Again, it's an ambition that's been around for almost 100 years, which has re-emerged most recently

<sup>28</sup> For example; Alef (<https://alefeducation.com>), ALEKS (<https://www.aleks.com>), Alta (<https://www.knewton.com>), AmritaCREATE (<http://www.amritacreate.org>), Area9 (<https://area9learning.com>), Assistments (<https://new.assistments.org>), Better Marks (<https://bettermarks.com>), Byjus (<https://byjus.com>), Century (<https://www.century.tech>), CogBooks (<https://www.cogbooks.com>), Cognii (<https://www.cognii.com>), Domoscio (<https://domoscio.com/solutions/adaptive-learning>), Dreambox (<http://www.dreambox.com>), EnLearn (<https://www.enlearn.org>), Inq-ITS (<http://www.inqits.com>), iReady (<https://www.curriculumassociates.com/Products/i-Ready>), Laix (<https://www.liulishuo.com>), Mathia (<https://www.carnegielearning.com>), Qubena (<https://qubena.com>), Realizeit (<http://realizeitlearning.com>), Querium (<http://querium.com>), Smart Learning Partner, Smart Sparrow (<https://www.smartsparrow.com>), Snappet (<https://nl.snappet.org>), Soffos (<https://soffos.ai>), Squirrel AI (<http://squirrelai.com>), Summit Learning (<https://www.summitlearning.org>), Thinkster Math (<https://hellothinkster.co.uk>), and Toppr (<https://www.toppr.com>), to name just a few.



from Silicon Valley: if we can have personalised shopping on Amazon and personalised viewing on Netflix, why can't we have personalised learning? But this completely misses the point. The so-called intelligent tutoring systems might provide each student with their own individual pathway through the materials, but they still take every student to the same fixed learning outcomes as everyone else. Although there are some benefits, this is a weak understanding of personalisation. Real personalisation is not about pathways but about helping each individual student to achieve their own potential, to self-actualise, which is something that no existing AI tool does. It all depends on what you want education to achieve. Homogenised students with lots of exam certificates? Or independent, creative, and critical thinkers, who are able to realise their own potential and ambitions, and are able to contribute constructively to society?

In any case, these student-supporting AI tools mostly adopt an extremely primitive behaviourist or instructionist approach to teaching and learning, which involves spoon feeding information, while avoiding failure, and which ignores more than sixty years of pedagogical developments. And in so doing, they disempower teachers, turning them all too often into mere technology facilitators. Again, this misses the point, ignoring the amazing skills that human teachers have, that no AI tool can replicate. These tools also undermine student agency. The student has no choice but to do what the AI requires, meaning that there is no opportunity for them to develop self-regulation skills or to self-actualise. Again, there might be some benefits, but the general point is that they are nowhere near as good as many suggest.

Finally, even if you reject these concerns, if you think all this is a fuss over nothing, and that AI is on track to solving education's big and intractable problems, we still do not know if any of the AI tools even do what they claim to do. The reality is that almost no AI and education companies' products have been independently evaluated – in real classrooms, do they genuinely improve student outcomes, without increasing the gap between rich and poor students? Instead, it is almost all just marketing. If the AI tools are so good, where is the evidence? Where are the robust, independent, randomised controlled trials? Apart from a couple of examples from the US (e.g. Roschelle et al., 2017), there are almost none.

## ALL A BIT NEGATIVE?

By now, you might think this paper has been unreasonably negative about the potential of AI for education. However, it is clear that what some AI tools can achieve is impressive, and the researchers have overcome some huge technical challenges. But still the tools are very limited, and frequently they are highly questionable (Holmes et al., 2019, 2021). So, how might we ensure that we do take advantage of AI in education to genuinely enhance teaching and learning?

At the moment, most AI and education developers start with the AI, and look around to see what problems they might fix – but this needs to change. Instead, we need to start with the problems that teachers and their students know best, and then ask AI engineers to use their extraordinary technologies to help solve those specific problems. In any case, rather than automating poor pedagogic practices, AI developers need to leverage AI's amazing capabilities to innovate teaching and learning, without compromising pedagogy or human values, to empower students and teachers, and to support inclusion and equity. For example, rather than using AI to facilitate exams, as the hugely controversial e-proctoring systems claim to do (Nigam et al., 2021), how might we use AI to assess and accredit learning while empowering teachers and enhancing student agency?

In conclusion, despite the potential and the hype, AI is yet to make much difference in education. Nonetheless, whether it is welcomed or not, its use in classrooms is growing quickly, and there are many tools available that teachers might employ. However, most of these tools simply automate poor pedagogic practices (e.g. e-proctoring), undervalue teacher skills (student-supporting AI is typically teacher replacing), reduce student agency (e.g. behaviourism), or address problems that are not anything that a teacher would recognise as being especially important. So, what needs to happen? In my opinion, we need teachers and AI developers to work together. In particular, we need to help the teachers see how AI could – if done right – transform pedagogy for the better, and we need to help the AI engineers understand the real education issues that they should focus on addressing. In addition, governments and curriculum developers need to focus on enabling all citizens to achieve an appropriate level of AI literacy involving both the human and technology dimensions of AI. Finally, governments need to ensure new funding, to enable learning scientists, philosophers, social scientists, and computer scientists to work together, to develop and

independently evaluate AI applications that address real education problems, that empower teachers and enable students to achieve their individual potential – all for the common good.

## REFERENCES

- Benaich, N. (2020, September 20). AI has disappointed on Covid. *Financial Times*.  
<https://www.ft.com/content/0aafc2de-f46d-4646-acfd-4ed7a7f6feaa>
- Boddington, P. (2017). *Towards a code of ethics for artificial intelligence research*. Springer, Berlin Heidelberg.
- Buolamwini, J., & Gebru, T. (2018). *Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification*. 15.
- Eynon, R., & Young, E. (2021). Methodology, Legend, and Rhetoric: The Constructions of , by Academia, Industry, and Policy Groups for Lifelong Learning. *Science, Technology, & Human Values*, 46(1), 166–191. <https://doi.org/0162243920906475/10.1177>
- Graesser, A. C., Wiemer-Hastings, K., Wiemer-Hastings, P., Kreuz, R., Group, T. R., & others. (1999). AutoTutor: A simulation of a human tutor. *Cognitive Systems Research*, 1(1), 3. 51.
- Heaven, W. D. (2019). Why deep-learning AIs are so easy to fool. *Nature*, 574(7777), 163–166. <https://doi.org/10.1038/d41586-019-03013-5>
- Holmes, W., Bialik, M., & Fadel, C. (2019). *Artificial Intelligence in Education. Promises and Implications for Teaching and Learning*. Center for Curriculum Redesign.
- Holmes, W., Porayska-Pomsta, K., Holstein, Ken, Sutherland, E., Baker, T., Buckingham Shur S., Santos, O. C., Rodrigo, M. M. T., Cukorova, M., Bittencourt, I. I., & Koedinger, K. (2021). Ethics of AI in Education: Towards a Community-wide Framework. *International Journal of Artificial Intelligence in Education*.
- Jobin, A., Ienca, M., & Vayena, E. (2019). Artificial Intelligence: The global landscape of ethics guidelines. *Nature Machine Intelligence*, 1(9), 389–399.  
<https://doi.org/10.1038/s42256-019-0088-2>
- Marcus, G. (2020). The Next Decade in AI: Four Steps Towards Robust Artificial Intelligence. *ArXiv:2002.06177 [Cs]*. <http://arxiv.org/abs/2002.06177>
- Marcus, G., & Davis, E. (2020). GPT-3, Bloviator: OpenAI's language generator has no idea what it's talking about. *MIT Technology Review*.  
<https://www.technologyreview.com/1007539/22/08/2020/gpt3-openai-language-generator-artificial-intelligence-ai-opinion/>
- McCarthy, J., Minsky, M. L., Rochester, N., & Shannon, C. E. (1956). *A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence*. <http://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html>
- Miao, F., & Holmes, W. (2021). *AI and education: Guidance for policy-makers*. UNESCO.  
<https://unesdoc.unesco.org/ark:/48223/pf0000376709>

## REFERENCES

- Benaich, N. (2020, September 20). AI has disappointed on Covid. *Financial Times*.  
<https://www.ft.com/content/0aafc2de-f46d-4646-acfd-4ed7a7f6feaa>
- Boddington, P. (2017). *Towards a code of ethics for artificial intelligence research*. Springer Berlin Heidelberg.
- Buolamwini, J., & Gebru, T. (2018). *Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification*. 15.
- Eynon, R., & Young, E. (2021). Methodology, Legend, and Rhetoric: The Constructions of , by Academia, Industry, and Policy Groups for Lifelong Learning. *Science, Technology & Human Values*, 46(1), 166–191. <https://doi.org/0162243920906475/10.1177>
- Graesser, A. C., Wiemer-Hastings, K., Wiemer-Hastings, P., Kreuz, R., Group, T. R., & others. (1999). AutoTutor: A simulation of a human tutor. *Cognitive Systems Research*, 1(1), 3. 51.
- Heaven, W. D. (2019). Why deep-learning AIs are so easy to fool. *Nature*, 574(7777), 163–166. <https://doi.org/10.1038/d41586-019-03013-5>
- Holmes, W., Bialik, M., & Fadel, C. (2019). *Artificial Intelligence in Education. Promises and Implications for Teaching and Learning*. Center for Curriculum Redesign.
- Holmes, W., Porayska-Pomsta, K., Holstein, Ken, Sutherland, E., Baker, T., Buckingham Shur S., Santos, O. C., Rodrigo, M. M. T., Cukorova, M., Bittencourt, I. I., & Koedinger, K. (2021). Ethics of AI in Education: Towards a Community-wide Framework. *International Journal of Artificial Intelligence in Education*.
- Jobin, A., Ienca, M., & Vayena, E. (2019). Artificial Intelligence: The global landscape of ethics guidelines. *Nature Machine Intelligence*, 1(9), 389–399.  
<https://doi.org/10.1038/s42256-019-0088-2>
- Marcus, G. (2020). The Next Decade in AI: Four Steps Towards Robust Artificial Intelligence. *ArXiv:2002.06177 [Cs]*. <http://arxiv.org/abs/2002.06177>
- Marcus, G., & Davis, E. (2020). GPT-3, Bloviator: OpenAI's language generator has no idea what it's talking about. *MIT Technology Review*.  
<https://www.technologyreview.com/1007539/22/08/2020/gpt3-openai-language-generator-artificial-intelligence-ai-opinion/>
- McCarthy, J., Minsky, M. L., Rochester, N., & Shannon, C. E. (1956). *A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence*. <http://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html>
- Miao, F., & Holmes, W. (2021). *AI and education: Guidance for policy-makers*. UNESCO.  
<https://unesdoc.unesco.org/ark:/48223/pf0000376709>



