Short read

Artificial Intelligence and School Education

Erica Southgate

Karen Blackmore

Stephanie Pieschl

Susan Grimes

Jessey McGuire

Kate Smithers

August 2019

Commissioned by the Australian Government Department of Education

To cite this document in APA referencing style:

Southgate, E., Blackmore, K., Pieschl, S., Grimes, S., McGuire, J. & Smithers, K. (2018). Short read: *Artificial intelligence and school education.* Newcastle: University of Newcastle, Australia.

This document is based on the research report below. Please refer to the main report for research and references:

Southgate, E., Blackmore, K., Pieschl, S., Grimes, S., McGuire, J. & Smithers, K. (2018). *Artificial intelligence and emerging technologies in schools: A research report.* Newcastle: University of Newcastle, Australia.

Artificial Intelligence and School Education

Defining artificial intelligence

AI is an umbrella term that refers to a machine or computer program that uses human-like thinking to complete a task. AI can take many forms. This includes robots that mimic human interaction, robotic manufacturing systems and ubiquitous computing systems. AI can be invisibly infused into the everyday computing application we use, such as internet search engines or the facial recognition tagging technology of social media applications.

The history of AI

The term artificial intelligence was first coined in the 1950s. Since the early 2000, advancement in AI has been driven by the expansion of the internet, availability of 'Big Data', and more powerful computing and algorithms.

Different types of AI

There are various ways to classify AI. One of the most common and popular ways to classify is by the categories, Narrow AI, General AI and Superintelligence (Super AI). At this time the type of AI we have is Narrow AI.

Types of AI Types of Artificial Intelligence (Adapted from Hintz, 2016)[1].

Superintelligence Example: None. Does not exist.		AI that exceeds human intelligence in every field.
General AI	**Self Awareness** Example: None. Does not exist.	AI at this level would extend the 'theory of mind' to predict the internal states of others. Having achieved human-like consciousness, it might choose to exhibit non-human abilities.
General AI	**Theory of Mind** Example: None. Does not exist.	This type of AI would have an updatable representation of the world that includes an understanding that other entities in the world also have their own internal states.
Narrow AI	**Limited Memory AI** Example: Virtual assistants, self-driving cars.	This type of AI receives current input, and adds pieces of this input to its programmed representation of the world. This can change the way the AI makes current or new decisions.
Narrow AI	**Reactive AI** Example: AI chess player.	Designed for a specific task, this AI receives input, and acts on this input. They cannot be applied to different tasks, and past experiences do not affect current decisions.

Understanding and misunderstanding AI

Children and adults often overestimate the intelligence of embodied and disembodied AI. Embodied AI refers to the way the technology can be embedded in some robots for example, while disembodied AI refers to the way AI can be infused into computing applications such as Internet search engines. Children and adults need to have personal experience, guidance and realistic information about the capability of AI and how AI works so that they do not develop misconceptions about the technology. Some common misconceptions of AI are:

- *Misconception 1 - AI is more intelligent than humans*: Current AI has not yet mastered developmental milestones that average human children master around age four. However, current AI can outperform humans in specific domains and on certain narrow or focused tasks.
- *Misconception 2 - AI is synonymous with robots:* There is overlap between AI and robots, but most AI tools are not embodied in robots and many robots are not powered by AI.
- Misconception 3 – *AI and telerobots are the same thing*. Telepresence is the use of remote controlled technology for apparent participation in distant events. For example, a student in a remote location could use a telerobot to attend regular classes. Telepresence can be achieved without AI. The user can provide the 'intelligence' necessary to control the technology.
- *Misconception 4 – AI is too hard to understand*: Even young children can learn about AI (see Milford's book for young learners[2]). It is up to educators, across disciplines, to come together to make learning about and with AI, accessible, interesting and relevant.

A brief overview of machine learning

Machine learning is a subfield of AI. Machine learning involves getting computers to learn over time in an autonomous fashion by giving them data from the real world. Machine learning can offer powerful ways to detect patterns in 'big data', sets and be useful in helping humans make decisions. There are a number of concerns about machine learning. These include its need to learn from 'big data' which may have been collected without consent or used in ways that people did not anticipate. Another concern relates to biased data sets used to train AI, and that allow AI to make decisions that discriminate against certain groups. The 'opaque' nature of 'deep' machine learning (e.g. artificial neural networks) also makes it difficult to understand how and why an AI makes its decisions. This 'black box' type of machine learning, where it is difficult to see the exact logics of the machine decision-making, is an ethical concern for human fields such as education where being able to explain why decisions were made is a central tenet of good teaching and governance.

The field of Artificial Intelligence in Education

Since the 1970s, the specialised field of Artificial Intelligence in Education (AIED) has grown to encompass the application of the technology to learning and instruction. The goal of AIED is to develop systems that enable personalised flexible and engaging learning and to automate mundane teaching tasks such as assessment and feedback. AIED is interested in developing AI-powered intelligent tutoring systems, virtual pedagogical agents that act as a peer or instructor, embodied AI robots, and 'smart' classrooms. The following diagram provides an overview of where AI may be included in educational applications.

Overview of typical AIED applications and their relationships.

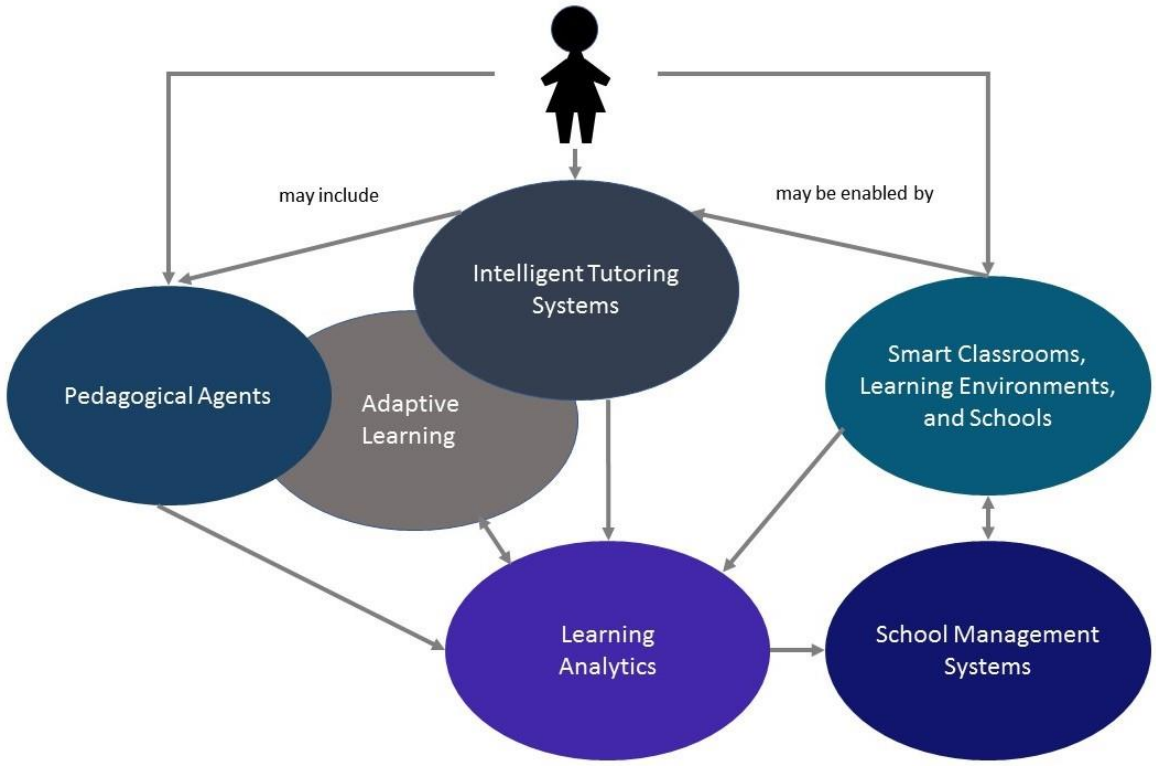

Intelligent tutoring systems

Intelligent Tutoring Systems (ITS) simulate (mimic) one-on-one human tutoring. Existing ITS implementation in schools is usually in well-defined areas and uses very simple pedagogy and student modelling. To date, the most technically sophisticated ITS have not been taken up by schools. On most educational levels, ITS has similar positive effects on learning as human tutors; however, it may not be suitable for all learners.

Pedagogical agents

Pedagogical Agents (PAs) are virtual characters integrated into learning technologies to facilitate instruction. Often PAs are embodied which means that learners can see pictures of virtual characters or avatars on screen that realistically or abstractly resemble humans, fictional characters, or objects. PAs were created to add a social, emotional, and motivational component to learning technologies and to communicate with learners in natural human-like ways. PAs can: serve as an information source; demonstrate or model learning content; coach or scaffold information processing self-regulation or motivation; and assess learners. Some PAs act as learning companions or peers. Some PAs are teachable which allows the human student to teach their PA and gain knowledge through explanation. Evidence regarding the effectiveness of PAs is mixed and some argue that students may be distracted from their learning by PAs.

Smart classrooms, learning environments and schools

The Internet of Things (IoT) describes the increasing capacity for everyday items to connect to the Internet and interact with other devices. Data from the IoT provides powerful input to AI systems. The IoT will be essential for smart classrooms and schools of the future. A smart learning

environment is equipped with wireless communication, personal digital devices, sensors, and learning platforms that connect with each other to provide input into AI systems. The AI then make decisions about regulating physical aspects of the environment (e.g. climate control) or learning systems. Most smart classroom applications are only at the stage of pilot testing and feasibility studies.

Adaptive learning and learning analytics

While ITS and PAs can provide a user interface, AI can power adaptive, personalised learning system. Adaptive systems can be used to gradually withdraw scaffolding until the learner reaches competence. Big data is drawn from ITS or learning management systems and can offer insights on how individual learning unfolds over time, and on which instructional approaches are effective under which conditions. The analysis of this data is only possible via learning analytics, educational data mining approaches, and with certain types of AI systems, and techniques such as evolutionary computation and artificial neural networks. One of the most important envisioned benefits is more adaptive personalised learning and instruction. However, the ethical issues around the use of AI powered systems in education and elsewhere are far from resolved.

Ethics, artificial intelligence and education

There are decades of research on the ethics of autonomous and intelligent systems. Many countries are engaged in thinking through ethical, regulatory and legislative approaches to ensure AI is used for the good of humanity. The way school systems will engage with AI will be largely determined by broader public education that can build trust in the technology. In schools, this trust must be founded on the ability of AI-powered systems to promote worthwhile and fair learning opportunities and the well-being of students, and their school community as whole.

One of the main ethical issues in the use of AI in education is who is ultimately accountable if something goes wrong. It is in our ethical interest to clarify the norms, values and assumptions reflected in AI systems so that these *enhance* human potential, creativity and well-being rather than foreclose or homogenise it. Algorithmic 'nudging' embedded in AI systems for education, and affective computing applications that attempt to influence a person's emotional state, raise concerns about respect for the right of humans to make their own choices based on sufficient information. The proliferation of machine learning generated 'deep fakes' – media which is a very realistic but untrue simulation of people or events – presents challenges for teaching digital literacy and critical thinking. The effects of using AI-powered robots for caring and teaching purposes, especially with young children, needs further research and regulatory frameworks. The *Education, Ethics and AI (EEAI) framework (diagram below)* developed for this project provides a comprehensive way of thinking through ethical issues for teachers, school leaders and policy makers.

Figure 2: The Education, Ethics and AI (EEAI) framework

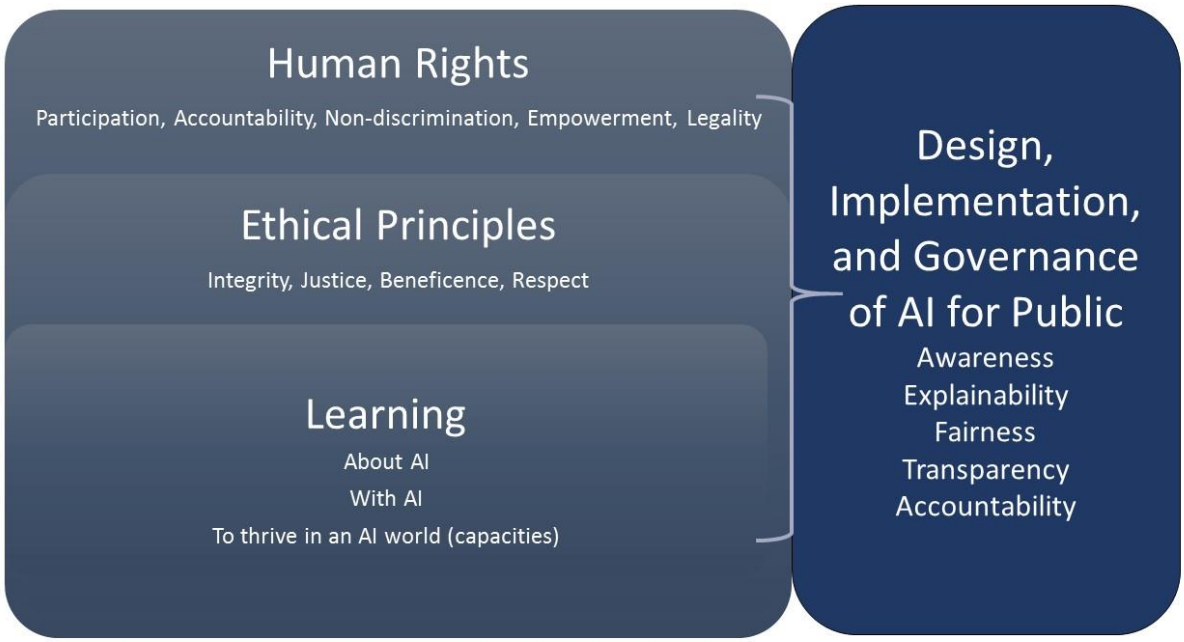

- **Human Rights:** The use of AI should accord with human rights which entail adherence to the principles of participation, accountability, non-discrimination and legal frameworks to support these.
- **Ethical Principles:** The contemporary ethical principles of integrity, justice, beneficence and respect are key to using AI-powered systems ethically and safely.
- **Ethical design, implementation and governance of AI systems in education:** There are many complex issues related to the design, implementation and governance of AI-powered systems. The five pillars of ethical AI - awareness, explainability, fairness, transparency and accountability – allow us to ask critical questions about AI-powered systems to ensure the technology is used to enhance learning, increase equity, and foster the well-being of all stakeholders in a school community.

Conclusion

AI could potentially offer benefits to teachers and students in the form of personalised learning and pedagogical agents designed to deliver appropriate and sequenced content and feedback to learners. However, AI is still in a relatively early stage of development for education and there is much work to be done around the ethical and legal frameworks that can ensure that the technology is used for good and not harm, and that transparent processes are in place to ensure accountability at classroom, school community, and school systems levels. Teachers, school leaders and policy-makers should begin to engage with developments in AI for education and society, in order to empower their students in the present and for future change.

Advice to teachers

A. Start your lifelong AI learning journey now:

Whether you are a pre-service teacher, a seasoned classroom educator, or a school leader, now is the time to begin your professional learning journey on the potential of AI for education. Even with a

good foundational knowledge, professional learning in this fast-moving area will need to be refreshed annually. School communities can begin to identify places across the curriculum (inclusive of but beyond STEM) where learning about AI can be integrated so that we can begin to equip students with knowledge about how the technology works, and how it is infused into and influences our interactions in everyday life.

B. Engage in the ethics of AI so that your school community can make informed decisions:

There are numerous ethical issues regarding the design, implementation and governance of AI-powered systems in education that need to be addressed, in an ongoing manner, by all stakeholders in school communities. It is vital that there is open, informed dialogue and transparency about the ethical quandaries of AI and education if trust is to be developed in the technology. The teaching profession has a long history of leading public discussion and providing accessible explanations on complex issues which affect students and their families and of grappling with issues of fairness, ethics, duty of care, and accountability in schooling. This makes the teaching profession well-equipped to both use AI technology for good and to ask critical questions regarding when and how machines should guide student learning and decision processes within educational settings, and whose values should be imbued into AI-powered systems.

C. Provide opportunities for students to learn with and about AI across curriculum learning areas and in the general capabilities area.

Free resources on learning with and about AI can be found at the Digital Technologies Hub https://www.digitaltechnologieshub.edu.au/footer/about-dth and CSER https://csermoocs.adelaide.edu.au/available-moocs

[1] Hintze, A. (2016, Nov 14). Understanding the four types of AI, from reactive robots to self-aware beings. The Conversation. Retrieved 13 December 2018, https://theconversation.com/understanding-the-four-types-of-ai-from-reactive-robots-to-self-aware-beings-67616

[2] Milford, M. (2018). The complete guide to artificial intelligence for kids. Available from https://www.kickstarter.com/projects/1740916372/the-complete-guide-to-artificial-intelligence-for/community