

## Research for CULT Committee - The use of Artificial Intelligence (AI) in education

### Concomitant expertise for INI report

#### KEY FINDINGS

There are two different types of AI in wide use today. Recent developments have focused on data-driven machine learning, but in the last decades, most AI applications in education (AIEd) have been based on representational / knowledge-based AI.

Data-driven AI uses a programming paradigm that is new to most computing professionals. It requires competences which are different from traditional programming and computational thinking. It opens up new ways to use computing and digital devices. But the development of state-of-the-art AI is now starting to exceed the computational capacity of the largest AI developers. The recent rapid developments in data-driven AI may not be sustainable.

The impact of AI in education will depend on how learning and competence needs change, as AI will be widely used in the society and economy. AIEd should be used to help schools and educational institutions in transforming learning for the future.

Many AIEd systems have been developed over the years, but few of these have shown clear scientific impact on learning. Evidence is lacking partly because the contexts of teaching and learning vary across classrooms, schools, educational systems, and countries. Local knowledge and capacity is critical for effective adoption and shaping of AIEd, and new scaling models are needed. Co-design of AIEd with teachers is a possible way to advance new scaling models.

AI has a great potential in compensating learning difficulties and supporting teachers. The Union/ the EU needs a “clearing house” that helps teachers and policy-makers make sense of the fast developments in this area.



## The big picture

In her 2019 Mobile Learning Week opening keynote, Director-General of UNESCO, Audrey Azoulay stated that AI was the biggest innovation in the human history since the paleolithic time. This may well be the case—if AI someday is invented.

Despite the common error of misplaced concreteness, AI is not a thing. It is a domain of research with many sub-disciplines, each with their own histories, domains of expertise, and developmental dynamics. This is important to understand, when assessing the potential impact of AI in education and policy.

There are three essentially different approaches to develop AI systems. **Symbolic computation** was a key driver in the emergence of AI research in the 1950s. The key idea was that computers are logical machines that can process bits of knowledge instead of only calculating numbers. This led to highly optimistic declarations that as soon as the generic logic of human reasoning could be programmed, computers would gain the essential qualities of human intelligence.

It was soon realised, however, that intelligent action requires extensive amounts of domain-specific knowledge. Towards the 1980s, this led to rapid growth in AI systems that relied on the manipulation of knowledge representations. In particular, representational AI—now often called “good-old-fashioned-AI” or GOFAI—focused on how the cognitive structures of expert decision-makers could be automatically processed. Since the 1980s, many such “**expert systems**” have been developed and deployed in large companies. Many programming techniques developed in the GOFAI research are now routinely used in all software development. This has led to the adage “When it works, it is not AI anymore.”

In pedagogic uses, the representational approach to AI has been dominant since the 1980s. **Intelligent tutoring systems** (ITS) typically contain representations of student’s current knowledge, a domain model that describes the knowledge to be learned, and a pedagogic model that steers the learner towards the learning objectives.

The recent interest in AI has its roots in a third approach: **artificial neural networks**. The first mathematical models of biological neural networks were developed in the 1930s. They became highly influential when it was shown that “universal logical machines” could be constructed from the simplest possible models of neurons as digital on-off elements. For many scientists influenced by logical positivism, this suggested that all rational thinking could be modelled with such networks.

Since the 1950s, many different models of artificial neural networks have been created. With inspiration from studies on biological neural networks, learning in these networks has typically been modelled as the strengthening of connections between simultaneously active neurons. This is known as **Hebbian learning**. For several decades, a major challenge in these network models, however, was the difficulty of using this simple associative learning rule in larger networks that contained many layers of artificial neurons.

The current AI revolution, to a large extent, results from the fact that it has now become possible to program computers with this simple learning rule. **This represents a new paradigm for using computers**. This approach is commonly called **data-driven AI**.

## Where do the AI breakthroughs come from?

Data-driven AI has generated major breakthroughs in the last nine years. To put these in a context, it is important to understand the root causes of progress.

Three key technical developments underpin recent advances in data-driven AI. First, in the last 15 years, the rapid expansion of social media, internet use, and smart phones have generated vast amounts of data, text, voice, and images. Second, data-driven AI uses very simple computations that can be done using hardware that was originally developed for graphics processing in computer games. Using these specialised processor architectures, very high computational power can be achieved at low cost for the types of processing that is needed for developing and training data-driven AI models. Third, **the Internet has enabled the low-cost distribution of human work at a massive scale.** Many of the advances in data-driven AI are based on the availability of data collections that have been processed and labelled by humans.

In the last decade, these three technical trends have converged in a very special innovation dynamic. Internet platform firms, who have access to data and real-time connectivity, have become the dominant users and developers of data-driven AI and major developers of AI research knowledge, software platforms, and processor hardware. As the effective use of real-time big data is impossible without automatic data processing, machine learning and data-driven AI have become a necessity for these companies. At the same time, hundreds of millions of end-users on these platforms constantly classify and categorise data, making separate labelling and categorising redundant. This has led to what can perhaps now be called google-sized natural monopolies of the Internet.

**This dynamic is not necessarily a sustainable one.** Extrapolations from the extraordinary developments of the last decade may have little predictive power. The number of computations required to generate state-of-the-art models has doubled every 3.4 months since the deep-learning breakthrough in 2012. In June 2019, Jérôme Presenti, Vice-President of AI at Facebook, said that Google and Facebook were now quickly running out of compute power. **Data-driven AI can solve some difficult practical problems, but it is probably the most wasteful computational approach invented in the human history.** If the brute-force approach to data-driven AI continues, AI may, indeed, rapidly become an important source of global warming.

Digitisation is often considered to be immaterial. It is therefore important to note that, currently, the mining of one Bitcoin requires fossil energy equivalent of 750 tons of concrete, or 60 barrels of oil. Effective policy development for AI in education therefore requires understanding also the technical drivers of AI, as well as the future of education in a world where AI technologies are widely used.

## Skills and competences in an AI-enabled world

To understand the potential impact of AI, it is useful to reconsider the EU competence frameworks. In these frameworks, “competence” is understood as a combination of expertise and attitude. Expertise, in turn, is viewed as a combination of knowledge, skill and experience.

A practical interpretation of competence is that it is a capability to get things done. This requires **epistemic components**, such as domain knowledge, accumulated experience, and skill. Epistemic components of competence are, however, not enough. This is what the EU competence frameworks conceptually capture with “attitude.” More broadly, however, epistemic components of competence need to be complemented with non-epistemic components, such as creative problem-solving, meta-cognitive capabilities, including self-reflection and emotional control, and the capacity to mobilise social resources and knowledge.

Skill, understood as an epistemic component of competence, is commonly associated with specific tools and techniques. In this sense, a car creates a car-mechanic, a computer creates a software programmer, and an anvil and a forge create a blacksmith. “Skill,” therefore, is conceptually a mirror image of current technology. When technology changes, skills become obsolete.

In the current digital transformation, technologies and tools used for work are rapidly changing. Epistemic components of competence are rapidly becoming obsolete. Education, therefore, is shifting its emphasis from epistemic content-related components of competence that were central in the last two centuries towards generic technology-independent “soft skills.” Social skills and capabilities to mobilise networked resources are becoming increasingly important as the Internet enables new forms of access and collaboration.

It is in this “post-Kondratiev” innovation dynamic, where the long-term impact of AI can best be understood. When the technical context of competence is losing its stability, experience, skill and domain-specific knowledge become less important. **Generic non-epistemic components of competence, such as creative problem-solving and meta-cognitive learning skills, in turn, become increasingly important.** It becomes less relevant what you know, and more important whose knowledge you can mobilise and what you can learn. Social and cultural skills that are necessary to effectively operate in the global networks of production and communication, become increasingly important.

In this setting, AI becomes a general-purpose technology that can perform tasks that previously required human knowledge and skill. Data-driven AI becomes necessary when the world becomes connected in real time, and when constant adaptation is needed to optimise activities in complex global networks that link actors across time and space. This also drives rapid change in skill and knowledge demand. As productive activities become increasingly automated in these real-time networks, human intervention, however, can become difficult. The future of work and demand for skills and education, therefore, cannot properly be understood just by focusing on AI systems themselves. It is necessary to understand the broader drivers that make AI systems economically interesting and socially important. These same drivers also generate important tensions in current educational systems. As a result, AI is often viewed as a way to reduce tensions between the institutions of the past and the needs of the present. For example, AI is commonly viewed as a tool that can provide individual personalised teaching of course material or as a way to automate repetitive teacher tasks. Another way of perceiving the potential of AI in education is to see it as a tool that could allow the transformation of education towards the needs of the future.

## AI and digital skills

The rapidly rising visibility of AI has led many educational institutions to expand the provision of AI-related content. The ‘Elements of AI’ online course, developed by the University of Helsinki and Reaktor, has been a very successful effort to provide introductory-level knowledge about AI for broad audiences. Its main objective has been to “demystify AI,” and now over 350,000 people from 170 countries have signed up to this free online course.

A particularly interesting aspect of ‘Elements of AI’ is that about 40 per cent of the learners have been women. This is more than double the average for computer science courses. In general, women now represent less than one fifth of AI researchers. **If job growth in the future increasingly occurs in tasks that require AI-related skills, as many experts claim, this may generate important gender imbalances in the labour market.**

The very high media visibility of AI and the stories about AI experts being hired with seven-figure annual salaries have created what could be called an AI gold rush. This has important implications for AI skill development. In particular, it raises the question whether AI skill gaps will be filled without policy interventions.

To understand the dynamics of AI competence development, it is useful to distinguish use, modification, development, and creation skills. The ‘Elements of AI’ course addresses basic usage

skills, creating awareness and helping the learners to make sense of the essential AI-related concepts and claims. This level of learning is necessary for effective and appropriate use of existing AI-systems. To achieve this level of competence requires less than a week of effort.

The majority of current AI systems are now created by relatively novice developers who rely on tools, frameworks, code, and learning material openly distributed by large companies, such as Google, Facebook and Microsoft. To be able to take into use existing tools and to modify them for specific purposes, some programming skills are needed.

On a very basic level, secondary school students can now build basic chatbots and machine learning systems in a few hours using this approach. There are also development interfaces specifically intended for children that reduce the need to know programming languages. For example, Machine Learning for Kids (<https://machinelearningforkids.co.uk/>) is now used in many schools and coding clubs. It provides simple programming interfaces to the IBM Watson services and allows children to develop programs in Scratch, Python and APP Inventor. A prominent initiative in this area has been AI4k12.org that has a very active mailing list for teachers who implement AI-projects in the classroom.

More generally, many systems used in service and manufacturing industries are modifications of freely available AI systems. For a relatively competent computer programmer, it takes some months to learn state-of-the-art AI development platforms and to modify existing code for business purposes.

The development of new AI systems and skilful modification of existing computational architectures using state-of-the-art development approaches requires substantially more effort. In general, graduate-level theoretical knowledge, support from competent peers or experts, and access to open source tools and commercial hardware platforms provided by the leading AI firms is necessary. Many universities have expanded their educational offers at this level of competence. Due to the high visibility and economic attractiveness of AI, the number of competent people at this level is increasing very rapidly. In particular, existing education in statistics, mathematics, computer science and physics can relatively easily be converted to AI-specific skills at this level.

The creation of new state-of-the-art AI models requires advanced theoretical knowledge and practical skills. Until recently, there have been only very few researchers with competences required to create new breakthroughs in AI and machine learning. Most of these have been employed by universities. As AI has become a strategic issue for many large corporations, this top-level talent has been in high demand. The largest data-driven AI conference, NeurIPS had some 13,500 attendees in 2019. This represents about 8-fold increase from 2012 and 41 per cent increase from 2018. Based on these numbers, one might roughly estimate that the number of highly competent AI researchers, able to create new AI architectures, is now perhaps 20,000 globally. These are the people that move the current technological frontier.

Radical breakthroughs in AI may, however, require broad and trans-disciplinary skills and knowledge. As was pointed out above, the current deep-learning gold rush is at least partly inspired by historical trends that may be running out of steam and electricity. **It is possible that qualitatively new types of processing and compute architectures will be needed to create environmentally viable AI systems.** Whereas continuous training of state-of-the-art machine learning systems now requires megawatts of electricity, the human brain works well with about 20 watts. Neuromorphic computing and new non-digital hardware architectures may become increasingly interesting in the future, and many of the hard-learned skills and knowledge of current-day AI experts may have only limited lifetime. **It is therefore not clear that formal education in AI-specific knowledge and skills will be able to generate competences that will be relevant in the future.** At present, most estimates of AI competence needs are based on rather straightforward

extrapolations of the past, and there do not seem to be ongoing attempts to develop more informed scenarios about AI-related competence needs and their potential development paths.

A specific characteristic of AI skill development is also that state-of-the-art knowledge and tools can only be accessed through the Internet. A similar dynamic of competence creation characterises open source communities. For example, in the 1990s, the Linux development community was able to create very rapidly high-level software and computer architecture competences outside the formal systems of education. Internet-enabled peer-to-peer learning is also important in the AI domain, and is now explicitly supported by some of the largest global companies. **As very high economic incentives now drain top talent from universities, the universities may face considerable challenges to be able to provide state-of-the-art knowledge in meaningful ways.** These challenges are exacerbated by the fact that at present the development of breakthrough AI systems requires access to big data and compute platforms that are, to a large extent, controlled by commercial actors.

## The use of AI in education

One common classification of AI in education is based on the main user of the system. A recent NESTA report distinguishes student-, teacher- and system-facing AI. Holmes, Bialik and Fadel (2019) further divide the student-facing AI systems in systems that aim at teaching students, usually based on instructivist pedagogy, and systems that aim at supporting learning, often building on more constructivist pedagogic approaches. The following table shows examples of such systems.

Different types of current AIEd systems (modified from Holmes et al. 2019, p. 165)

Student teaching	Student supporting	Teacher supporting	System supporting
<ul style="list-style-type: none"> <li>Intelligent tutoring systems (including automatic question generators)</li> <li>Dialogue-based tutoring systems</li> <li>Language learning applications (including pronunciation detection)</li> </ul>	<ul style="list-style-type: none"> <li>Exploratory learning environments</li> <li>Formative writing evaluation</li> <li>Learning network orchestrators</li> <li>Language learning applications</li> <li>AI Collaborative learning</li> <li>AI Continuous assessment</li> <li>AI Learning companions</li> <li>Course recommendation</li> <li>Self-reflection support (learning analytics, meta-cognitive dashboards)</li> <li>Learning by teaching chatbots</li> </ul>	<ul style="list-style-type: none"> <li>ITS+learning diagnostics</li> <li>Summative writing evaluation, essay scoring</li> <li>Student forum monitoring</li> <li>AI teaching assistants</li> <li>Automatic test generation</li> <li>Automatic test scoring</li> <li>Open Education Resources (OER) content recommendation</li> <li>Plagiarism detection</li> <li>Student attention and emotion detection</li> </ul>	<ul style="list-style-type: none"> <li>Educational data mining for resource allocation</li> <li>Diagnosing learning difficulties (e.g. dyslexia)</li> <li>Synthetic teachers</li> <li>AI as a learning research tool</li> </ul>

A recent review of peer-reviewed academic AIEd articles found that extant research has covered four main areas of AI in higher education:

- adaptive systems and personalisation
- assessment and evaluation
- profiling and prediction
- intelligent tutoring systems

The different uses of AI can also be categorised based on the student life cycle. In the above dataset, 63 percent of the academic articles described systems for academic support services, 33 percent described administrative and institutional services, and 4 percent covered both. Academic support

services included systems for teaching and learning (e.g. assessment, feedback, tutoring). Administrative and institutional systems included systems such as admission, counselling, and library services. A further coding of the articles generated four main areas of AI application, shown in Table 2.

Table 2: Number of AI applications across peer-reviewed studies, multiple mentions possible (source: Zawacki-Richter et al., 2019)

AI applications	n	%
Adaptive systems and personalisation (teaching course content; recommending personalized content; supporting teachers and learning design; using academic data to monitor and guide students; representation of knowledge in concept maps)	27	18%
Assessment and evaluation (automated grading; feedback; evaluation of student understanding, engagement and academic integrity; evaluation of teaching)	36	24%
Profiling and prediction (admissions decisions and course scheduling; drop-out and retention; student models and academic achievement)	58	39%
Intelligent tutoring systems (teaching course content; diagnosing strengths and automated feedback; curating learning materials; facilitating collaboration; the teacher's perspective)	29	19%
<b>Total</b>	<b>150</b>	<b>100%</b>

Zawacki-Richter et al. (2019) note that a large majority of the papers that they analysed in detail, were authored by computer scientists and authors from STEM departments. Only nine of the 146 first article authors were from education departments. Almost four fifths of all the articles applied quantitative methods. This, however, probably reflects the selection criteria for journals that were included in the study. In particular, in the U.S. academic setting, publishable articles tend to require quantitative methods, and this often leads to quantitative studies where computer-science researchers use their classrooms as research laboratories. In Europe, much of the AIED research has been published in conferences that were not covered by this study.

**There is relatively scarce evidence about the benefits of AI-based systems in education.** In some specific areas, such as mathematics and physics, intelligent tutoring systems have been shown to improve learning, but it is also clear that learning benefits cannot be achieved simply by introducing new tools in a classroom. A recent review of critiques on ITS by Benedict du Boulay argued that a key to educational impact is teacher training that helps the teacher to orchestrate technology use. **In other words, the learning outcomes do not depend on technology. It depends on how the teachers can use technology in pedagogically meaningful ways. An appropriate approach, therefore, is to co-design the uses of technology with teachers.** This approach has been the starting point in the EU-funded New Era of Learning -project, where the largest Finnish cities have provided opportunities for rapid AIED experiments and co-design with technology developers, teachers and students. A similar approach was also used in the Joint Research Centre (JRC) of the European Commission "AI Handbook with and for Teachers" -pilot project that finished in December 2019. In the project, teachers knowledgeable of AI and AI developers knowledgeable of teaching jointly developed a prototype model for EU-level co-creation network that would produce an AI handbook that would help teachers and education developers to deploy and use AI in appropriate ways. Such an approach aims to move beyond conventional technology-push and demand-pull models of innovation, adopting a middle-up-down diffusion model. As technology is advancing rapidly and many products will become available, teachers and education administrators will need high-quality information that will help them make sense of this rapidly changing landscape. In the U.S., the Department of Education has invested in the "What Works Clearinghouse" that consolidates scientific evidence on educational products and policies, and there have been many similar initiatives in the Member States. There seems to be clear potential for coordinating such initiatives at the EU level.

## The AIEd business landscape

The rapidly increasing interest in AI has resulted in many start-ups and more established firms to re-profile their products as AI products. Global private investment in AI start-ups was about 40 billion US dollars in 2019, and the cumulative investment in the 2014-2019 period has been over 135 billion. According to the Stanford AI 2019 Index, over 3000 AI companies received over 400k USD private funding in 2018. There are now hundreds of businesses in Europe that claim to be in the AI business, and many educational technology firms do this as well. Although it is possible to provide some numbers of firms existing in this domain using business directories and information on venture capital funding, at present, there does not seem to be any systematic studies available. **Many AI firms that develop solutions in the educational domain do not define their core business as education.** For example, HeadAI in Finland has developed AI-based job market and skill demand prediction systems that are also used in career guidance. The official industry classification of the company is IT services.

Data from 2018 Asgard Global AI database shows that **in Europe, the United Kingdom, Finland and Sweden have been the host countries with most AIEd firms.** In the U.K., the EU-funded EDUCATE accelerator, hosted by the University College London (UCL), has generated several education technology start-ups that use AI technologies in their products. Most visible of these has been Century that markets its product as a tool that combines research in AI, neuroscience, and learning science. In Finland, Claned provides a relatively mature cloud-based learning platform that uses AI technologies. In Sweden, Sana Labs provides AI functionality that can be integrated with learning platforms through APIs. The recently launched H2020 IMPACT EdTech incubator/accelerator project also aims at helping 42 digital education innovators to bring their digital learning solutions into the market. It is to be expected that several proposals for the IMPACT EdTech calls will include AIEd.

An example of a well-known commercially successful and extensively tested GOFAI intelligent tutoring system is MATHia, developed by a Carnegie Mellon -spinoff Carnegie Learning. MATHia is intended for blended mathematics learning for US Middle and High School students. A more recent product from Carnegie Learning is MATHiaU, that is intended for college students who need remedial mathematics learning. MATHia and several other commercial AIEd systems have been reviewed by Holmes et al.

AI technologies are now used in a wide variety of everyday technologies, such as smart phone cameras, voice assistants, social media apps, and search engines. AI, therefore, is already embedded in many learning activities. Many students now use language translators, image classification systems, and automatic speech to text subtitling for their daily schoolwork. Instead of separate AI-based systems in and for education, AI is now becoming a technology that enables new educational ecosystems.

One example is Gooru Navigator. It was originally developed as a side-project at Google, in an attempt to create a "Google Navigator for learning." Gooru uses a number of AI technologies to classify learners' current knowledge, skills and mindsets, and to match learning objectives with learning content created by an open community of teachers. Gooru aims to create an ecosystem of content, where Gooru Navigator operates as a personal "search engine" and "route mapper" that suggests relevant content items that help the learner to move from one's current location to the stated aim. The system also uses analytic capabilities to generate insights at different levels, including the progress of a specific class, school or geographic region.

**AI can potentially have a substantial impact on special needs education.** One interesting company in this area is Lexplore. The business was started in Sweden by researchers from Karolinska



Institutet based on the large amount of eye movement data on students with reading and writing difficulties collected in the Kronoberg Project over the years. Their product can rapidly scan the eye movements of a child during reading, and the AI model in the system can detect with high accuracy whether the child suffers from dyslexia.

## Ethics of AI in education

Ethics of AI is an important topic for teachers both from a conceptual and practical point of view. The users of educational AI need to know when AI systems are aligned with the general ethical principles that underpin education. This requires that we make explicit these principles as well as the challenges in implementing them. Similarly, when AI is a topic of study, it becomes almost immediately clear how technical designs rely on often implicit ethical assumptions. Already very basic experiments with programming AI systems allows teachers and students to understand how ethical principles and values influence our current technological world.

The JRC pilot 'AI Handbook with and for Teachers' has developed a workbook on ethics of AI in education. This focuses on three aspects of AI: **impact on human autonomy, fairness, and explainability.**

Many recommendations on the ethical use and development of AI have been published by international organisations, national government bodies, business organisations and academic researchers. A recent review of 84 reports on ethics found great variance in the content and approach of these proposals. **A recurrent theme in these guidelines, however, is the need for transparency in automated decision-making, non-discrimination, accountability and safety.**

One of the most frequently discussed ethical challenges in AI is bias. It should be noted that data-driven AI systems are essentially systems that find biases in data that are used in their training. The question, therefore, is what are socially and culturally acceptable biases, and how they should be taken into account in developing, deploying and using AI systems. The different types of biases are shown in the table below.

Sources of unacceptable bias in data-driven AI systems (source: Tuomi, 2019)

	Source of bias	Example
Model bias	Unbalanced training data Bias in training data Model artefacts	Different sample sizes for different groups produce different estimates Historical data reflects actual biases, e.g., gender bias Unexpected associations are generated in the training process
System bias	Feedback loops System use Data representation and measurement error	Model predictions affect action and further data collection Scoring of people in inappropriate contexts, e.g., control of opinion and social interaction Implicit classification in converting knowledge and observation into computable data structures
Social bias	Cultural bias AI divide	Existing norms, practices, routines, and power structures Accumulated effects of unequal access to knowledge and technology

Ethics of AI, in general, is a complex conceptual, technical and political challenge. There are over 24 different definitions of "fairness" in use by AI researchers. Fairness in public services is based on social contracts that have been reflected in the various human rights declarations, but these are not necessarily appropriate for private service providers or relations between individual AI system users. **Generic principles are difficult to operationalise**, and there is therefore, increasing interest in

regulating the use of AI around the world. A review of a recent study on privacy and ethics informed design in the learning domain noted, for example, that the GDPR does not adequately address the need to process relational network data that may be needed for many social learning models.

In education, a special challenge is that data-driven AI systems encode patterns of historical data. **These systems therefore are inherently unable to make sense of acts that do not follow precedents.** When individuals become something that they were not before, i.e., when they learn, data-driven AI systems have difficulties in understanding them. **To the extent that education aims at the realisation of human potential, creativity and authentic social change, the wide use of data-driven AI systems may therefore also limit human and social development.**

## Bibliography

- Baker, Toby, and Laurie Smith. 2019. "Educ-AI-Tion Rebooted? Exploring the Future of Artificial Intelligence in Schools and Colleges." London: NESTA.
- Barocas, Solon, Moritz Hardt, and Arvind Narayanan. 2019. *Fairness and Machine Learning*. fairmlbook.org. <http://www.fairmlbook.org>.
- Benfatto, Mattias Nilsson, Gustaf Öqvist Seimyr, Jan Ygge, Tony Pansell, Agneta Rydberg, and Christer Jacobson. 2016. "Screening for Dyslexia Using Eye Tracking during Reading." *PLOS ONE* 11 (12): e0165508. <https://doi.org/10.1371/journal.pone.0165508>.
- Boulay, Benedict du. 2019. "Escape from the Skinner Box: The Case for Contemporary Intelligent Learning Environments." *British Journal of Educational Technology* 50 (6): 2902–19. <https://doi.org/10.1111/bjet.12860>.
- Hagendorff, Thilo. 2019. "The Ethics of AI Ethics -- An Evaluation of Guidelines." *ArXiv:1903.03425 [Cs, Stat]*, October. <http://arxiv.org/abs/1903.03425>.
- Holmes, Wayne, Stamatina Anastopoulou, Heike Schaumburg, and Manolis Mavrikis. 2018. "Technology-Enhanced Personalised Learning: Untangling the Evidence." Stuttgart: Robert Bosch Stiftung. [https://www.bosch-stiftung.de/sites/default/files/publications/pdf/2018-08/Study\\_Technology-enhanced%20Personalised%20Learning.pdf](https://www.bosch-stiftung.de/sites/default/files/publications/pdf/2018-08/Study_Technology-enhanced%20Personalised%20Learning.pdf).
- Holmes, Wayne, Maya Bialik, and Charles Fadel. 2019. *Artificial Intelligence in Education: Promises and Implications for Teaching & Learning*. Boston, MA: The Center for Curriculum Redesign.
- Jobin, Anna, Marcello Lenca, and Effy Vayena. 2019. "The Global Landscape of AI Ethics Guidelines." *Nature Machine Intelligence* 1 (9): 389–99. <https://doi.org/10.1038/s42256-019-0088-2>.
- Johnson, Khari. 2019. "Facebook VP: AI Has a Compute Dependency Problem." *Venture Beat*, July 11, 2019. <https://venturebeat.com/2019/07/11/facebook-vp-ai-has-a-compute-dependency-problem/>.
- Luckin, Rosemary. 2018. *Machine Learning and Human Intelligence: The Future of Education for the 21st Century*. London: UCL Institute of Education Press.
- Luckin, Rosemary, and Mutlu Cukurova. 2019. "Designing Educational Technologies in the Age of AI: A Learning Sciences-Driven Approach." *British Journal of Educational Technology* 50 (6): 2824–38. <https://doi.org/10.1111/bjet.12861>.
- Metz, Cade. 2018. "A.I. Researchers Are Making More Than \$1 Million, Even at a Nonprofit." *The New York Times*, May 4, 2018, sec. Technology. <https://www.nytimes.com/2018/04/19/technology/artificial-intelligence-salaries-openai.html>.
- Mittelstadt, Brent, Chris Russell, and Sandra Wachter. 2019. "Explaining Explanations in AI." In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 279–288. FAT\* '19. New York, NY, USA: ACM. <https://doi.org/10.1145/3287560.3287574>.

- Pedró, Francesc, Miguel Subosa, Axel Rivas, and Paula Valverde. 2019. "Artificial Intelligence in Education: Challenges and Opportunities for Sustainable Development." Working Papers on Education Policy. Paris: UNESCO.
- Perrault, Raymond, et al. 2019. "The AI Index 2019 Annual Report." Stanford, CA: Stanford University. [https://hai.stanford.edu/sites/default/files/ai\\_index\\_2019\\_report.pdf](https://hai.stanford.edu/sites/default/files/ai_index_2019_report.pdf).
- Salvo, B. De. 2018. "Brain-Inspired Technologies: Towards Chips That Think?" In *2018 IEEE International Solid - State Circuits Conference - (ISSCC)*, 12–18. <https://doi.org/10.1109/ISSCC.2018.8310165>.
- Shum, Simon J. Buckingham, and Rosemary Luckin. 2019. "Learning Analytics and AI: Politics, Pedagogy and Practices." *British Journal of Educational Technology* 50 (6): 2785–93. <https://doi.org/10.1111/bjet.12880>.
- Songer, Nancy B., Michelle R. Newstadt, Kathleen Lucchesi, and Prasad Ram. 2020. "Navigated Learning: An Approach for Differentiated Classroom Instruction Built on Learning Science and Data Science Foundations." *Human Behavior and Emerging Technologies* 2 (1): 93–105. <https://doi.org/10.1002/hbe2.169>.
- Strubell, Emma, Ananya Ganesh, and Andrew McCallum. 2019. "Energy and Policy Considerations for Deep Learning in NLP." *ArXiv:1906.02243 [Cs]*, June. <http://arxiv.org/abs/1906.02243>.
- Tuomi, Ilkka. 2007. "Learning in the Age of Networked Intelligence." *European Journal of Education* 42 (2): 235–54. <https://doi.org/10.1111/j.1465-3435.2007.00297.x>.
- ———. 2018. "The Impact of Artificial Intelligence on Learning, Teaching, and Education: Policies for the Future." JRC Science for Policy Reports. Luxembourg: Publications Office of the European Union. <https://doi.org/10.2760/12297>.
- ———. 2019a. "AI Energy Consumption Is Becoming a Global Problem." *Helsingin Sanomat*, August 31, 2019. Translated from the original op-ed in Finnish: <https://www.linkedin.com/pulse/ai-energy-consumption-becoming-global-problem-ilkka-tuomi>.
- ———. 2019b. "Ethics of AI in Education: Workbook for the JRC 'AI Handbook with and for Teachers' Project." unpublished deliverable, December 2019.
- What Works Clearinghouse. 2020. Studies on Distance Learning. Institute of Education Sciences. <https://ies.ed.gov/ncee/wwc/DistanceLearningStudy>
- Zawacki-Richter, Olaf, Victoria I. Marín, Melissa Bond, and Franziska Gouverneur. 2019. "Systematic Review of Research on Artificial Intelligence Applications in Higher Education – Where Are the Educators?" *International Journal of Educational Technology in Higher Education* 16 (1): 39. <https://doi.org/10.1186/s41239-019-0171-0>.

## Further information

This briefing is available at: <https://bit.ly/3lCMotK>

More information on Policy Department research for CULT:  
<https://research4committees.blog/cult/>



**Disclaimer and copyright.** The opinions expressed in this document are the sole responsibility of the authors and do not necessarily represent the official position of the European Parliament. Reproduction and translation for non-commercial purposes are authorised, provided the source is acknowledged and the European Parliament is given prior notice and sent a copy. © European Union, 2020.

Research manager: Pierre HERIARD      Editorial assistant: Lyna PÄRT

Contact: [Poldep-cohesion@ep.europa.eu](mailto:Poldep-cohesion@ep.europa.eu)

This document is available on the Internet at: [www.europarl.europa.eu/supporting-analyses](http://www.europarl.europa.eu/supporting-analyses)

IP/B/CULT/IC/2020-022

Print      ISBN 978-92-846-6703-1 | doi:10.2861/618143 | QA-03-20-297-EN-C  
PDF      ISBN 978-92-846-6704-8 | doi:10.2861/439239 | QA-03-20-297-EN-N