European Commission

JRC SCIENCE FOR POLICY REPORT

The Impact of Artificial Intelligence on Learning, Teaching, and Education

Policies for the future

Author: **Tuomi, Ilkka**

Editors: Cabrera, Marcelino; Vuorikari, Riina; Punie, Yves

Joint Research Centre

EUR 29442 EN

Contact information
Address: Edificio Expo. C/ Inca Garcilaso 3, E-41092 Seville (Spain)
Email: marcelino.cabrera@ec.europa.eu
Tel.: +34 9544 88246
EU Science Hub
https://ec.europa.eu/jrc

JRC113226

EUR 29442 EN

PDF	ISBN 978-92-79-97257-7	ISSN 1831-9424	doi:10.2760/12297

Luxembourg: Publications Office of the European Union, 2018

How to cite this report: Tuomi, I. *The Impact of Artificial Intelligence on Learning, Teaching, and Education. Policies for the future,* Eds. Cabrera, M., Vuorikari, R & Punie, Y., EUR 29442 EN, Publications Office of the European Union, Luxembourg, 2018, ISBN 978-92-79-97257-7, doi:10.2760/12297, JRC113226.

Title: The Impact of Artificial Intelligence on Learning, Teaching, and Education

Abstract

This report describes the current state of the art in artificial intelligence (AI) and its potential impact for learning, teaching, and education. It provides conceptual foundations for well-informed policy-oriented work, research, and forward-looking activities that address the opportunities and challenges created by recent developments in AI. The report is aimed for policy developers, but it also makes contributions that are of interest for AI technology developers and researchers studying the impact of AI on economy, society, and the future of education and learning.

Contents

Preface

Artificial Intelligence (AI) is currently high on the political and research agendas around the world. With the emergence of every new technology, there is always both a lot of hype and scepticism around its implications for society and the economy. Although acknowledging that the foundations for AI have been already around for several decades, recent technological breakthroughs are accelerating what AI could do. This study looks at what this could mean for learning, teaching, and education. It aims to provide a critical review and prospective angle on relevant AI developments as a basis for well-informed policy-oriented discussions about the future of these domains.

This report is a contribution to the Digital Education Action Plan[1] which foresees policy research and guidance on the impact and potential of digital technologies in education. It is done on behalf of the Directorate-General for Education, Youth, Sport and Culture, authored by Ilkka Tuomi and edited by the JRC. Another report, appraising AI from different perspectives, entitled "Artificial Intelligence: A European perspective", will be released soon under the label of JRC flagship reports, providing an overall assessment of opportunities and challenges of AI from a European outlook, and supporting the development of European action in the global AI context.

The JRC has carried out research on Learning and Skills for the Digital Era since 2005. It aims to provide evidence-based policy support to the European Commission and its Member States on how to harness the potential of digital technologies to encourage innovation in education and training practices; improve access to lifelong learning; and impart the new (digital) skills and competences needed for employment, personal development and social inclusion. More than 20 major studies have been undertaken on these issues, resulting in more than 120 different publications.

Recent work has focused on the development of digital competence frameworks for citizens (DigComp), educators (DigCompEdu), educational organisations (DigCompOrg) and consumers (DigCompConsumers). A framework for opening up higher education institutions (OpenEdu) was also published in 2016, along with a competence framework for entrepreneurship (EntreComp). Some of these frameworks are accompanied by (self-) assessment instruments. The JRC is also entrusted to develop a future framework for personal and social development, including learning to learn. Additional research has been undertaken on Learning Analytics, MOOCs (MOOCKnowledge, MOOCs4inclusion), Computational thinking (Computhink) and policies for the integration and innovative use of digital technologies in education (DigEduPol).

More information on all our studies can be found on the JRC Science hub: https://ec.europa.eu/jrc/en/research-topic/learning-and-skills.

[1] Communication from the Commission to the European Parliament, the Council, the European Economic and Social Committee and the Committee of the Regions on the Digital Education Action Plan (COM(2018) 237 final).

Executive summary

At the November 2017 Gothenburg Summit, the Commission presented the Communication 'Strengthening European Identity through Education and Culture', that set out a vision for a European Education Area and announced a dedicated Digital Education Action Plan[2], which aims to foster digital skills and competences for all citizens. The Action Plan focuses on implementation and the need to stimulate, support and scale up purposeful use of digital and innovative education practices. It has three priorities: making better use of digital technology for teaching and learning; developing relevant digital competences and skills for the digital transformation; and improving education through better data analysis and foresight. Artificial Intelligence (AI) will have an impact on all these, and in the last priority the Communication specifically invites to explore its impact in education and training through pilots. **This policy foresight report suggests that in the next years AI will change learning, teaching, and education**. The speed of technological change will be very fast, and it will create high pressure to transform educational practices, institutions, and policies. It is therefore important to understand the potential impact of AI on learning, teaching, and education, as well as on policy development.

AI is currently high on the political agendas around the world. Several EU Member States have declared it as a political priority. Influential studies now suggest that perhaps one in two occupations in the industrialized countries is likely to become automated using already existing AI technologies. Policy makers at the European Parliament have highlighted the importance of the issue, and the European Commission, in its 2018 annual work programme, sets its wish to make the most of AI, which will increasingly play a role in our economies and societies[3]. **AI is now often called "the next electricity."** The transformative impact of general purpose technologies, like AI, however, becomes visible only gradually, when societies and economies reinvent themselves as users of new technologies. Technological change brings social and cultural change that is reflected in lifestyles, norms, policies, social institutions, skills, and the content and forms of education.

Wide availability of cheap processing power and vast amounts of data in recent years have enabled impressive breakthroughs in machine learning and created extraordinary commercial and research interest in artificial neural networks, i.e. computational models based on the structure and functions of biological neural networks. Neural AI, and machine learning methods associated with it, are now used for real-time language processing and translation, image analysis, driverless cars and autonomous vehicles, automated customer service, fraud detection, process control, synthetic art, service robots, and in many other applications. **Although some of this excitement may be based on unrealistic expectations and limited knowledge of the complexities of the underpinning technologies, it is reasonable to expect that the recent advances in AI and machine learning will have profound impacts on future labour markets, competence requirements, as well as in learning and teaching practices.** As educational systems tend to adapt to the requirements of the industrial age, AI could make some functions of education obsolete and emphasize others. **It may also enable new ways of teaching and learning.**

2 Communication from the Commission to the European Parliament, the Council, the European Economic and Social Committee and the Committee of the Regions on the Digital Education Action Plan (COM(2018) 237 final).

3 Communication from the Commission to the European Parliament, the Council, the European Economic and Social Committee and the Committee of the Regions Commission Work Programme 2018 - An agenda for a more united, stronger and more democratic Europe (COM(2017) 650 final).

In the European framework programmes for research and technological development, AI technologies have been studied and applied in educational contexts in many projects focusing on technology-enabled learning. These projects have used technologies that have deep ties with AI research, including natural language processing, pattern recognition, intelligent tutoring, probabilistic AI planning, intelligent agents, AI game engines, and adaptive user models in personalized learning environments (PLE). **The impact of these technologies in practical educational settings has been relatively modest until recently. Technical developments over the recent years, however, suggest that the situation may be changing rapidly.**

The main intent of the present report is to help educators and policymakers to make sense of these potentially very important technical developments. To understand the impact of AI, we need to understand what AI is and what it can do. In the current "AI avalanche" this is not always easy. Deep expertise in AI technology is scarce, and many educators and policymakers now struggle to get up to date with basic knowledge in this area. In the midst of self-driving cars, speaking robots, and the flood of "AI miracles", **it may be easy to think that AI is rapidly becoming super intelligent, and gain all the good and evil powers awarded to it in popular culture. This, of course, is not the case. The current AI systems are severely limited, and there are technical, social, scientific, and conceptual limits to what they can do.** Perhaps surprisingly, well-established research on human learning provides important tools and concepts that help us understand the state-of-the-art and future of AI. Many current AI systems use rather simplified models of learning and biological intelligence, and learning theories thus help us gain better understanding of the capabilities of current AI systems.

There will be great economic incentives to use AI to address problems that are currently perceived as important by educational decision- and policy-makers. This creates policy challenges. For educational technology vendors it is easy to sell products that solve existing problems, but it is very difficult to sell products that require changes in institutions, organizations and current practices. To avoid hard-wiring the past, it would be important to put AI in the context of the future of learning. Policy may be needed to orient development in AI towards socially useful directions that address the challenges, opportunities, and needs of the future. **As AI scales up, it can effectively routinize old institutional structures and practices that may not be relevant for the future.** Future-oriented work, therefore, is needed to understand the potential impact of AI technologies. How this potential is realized depends on how we understand learning, teaching and education in the emerging knowledge society and how we implement this understanding in practice. **Future-oriented policy experimentation, as suggested by the Digital Education Action Plan, may, therefore, be an effective way to address this challenge.** .

Recent AI breakthroughs are based on supervised machine learning. A critical success factor of these systems is the availability of huge amounts of pre-categorized training data. In contrast to logic- and knowledge-based approaches to AI, we therefore characterize these as "data-based" AI systems in this report. Many of these "deep-learning" neural AI systems may well be characterized as "datavores." **At present, the most important technical bottleneck of AI, therefore, is the availability of data**. This is a qualitatively new development in the history of computing and information processing. Without access to vast training datasets, it is very difficult to develop successful AI systems. In this report, we put forward an argument that EU policies could create data platforms that could redefine the competitive landscape for learning- and education-oriented AI systems.

As these supervised AI learning algorithms are based on historical data, they can only see the world as a repetition of the past. This has deep ethical implications. When, for example, students and their achievements are assessed using such AI systems, the assessment is necessarily based on criteria that reflect cultural biases and historically salient measures of success. Supervised learning algorithms create unavoidable biases, and these are currently extensively debated. From a more fundamental ethical point of view, however, the expression of human agency requires capability to make authentic choices that do not only repeat the past. Although there are already AI systems that deal with creative activities, AI systems will have great difficulties in dealing with people who are creative, innovative, and not only average representations of vast collections of historical examples.

It is often assumed that AI systems enable new levels of personalisation and diversity for information systems; much of this, however, results from fine-grained categorization that puts users into pre-defined classes. Although these systems may be able to efficiently simulate personalisation, they do not necessarily support deeper levels of diversity. At present we can say that the use AI systems in educational settings will shape the development of human cognition and self-efficacy, but we don't know how. It is therefore important to continuously evaluate, for example, how the use of AI in educational contexts constrains and enables human possibilities for responsible and ethical action. AI systems can be excellent predictive machines, but this strength may be an important weakness in domains where learning and development are important. A contribution of this report is to show that different types of AI and machine learning systems operate on different layers of human behaviour[4]. **Most importantly, the level of meaningful activity**—which in socio-cultural theories of learning underpins advanced forms of human intelligence and learning—**remains beyond the current state of the AI art.**

One of the most successful application areas in AI has been video processing. There will be strong economic interests in using video-connected AI systems in classrooms and to complement the collected data with data from social media and Internet of things (IoT) platforms. As it becomes technically possible to monitor student emotions and attention in real time and use such data to help teachers and students, AI privacy and security become important topics also in education. Similarly, AI systems are well suited for collecting informal evidence of skills, experience, and competence from open data sources, including social media, learner portfolios, and open badges. This creates both ethical and regulatory challenges.

Several high-profile econometric studies on the future of work have shown that many occupations can be automated with current AI technologies. These studies have relied on task- and skill-biased models of technical change. In this report, we argue that a data-biased model is more appropriate for current AI systems. We also explore a similar methodology to see how the future of the teaching profession might look like. The results suggest that many currently defined high-priority teacher tasks might be automated. However, this is based on the assumption that the role of teachers is rather mechanical and purely instructional with summative assessment playing a central role, reflecting deep beliefs about the functions of education and the social institutions around it. In educational systems that emphasize development and, for example, social competences, formative assessment might be higher on the list. As a result, **there is a risk that AI might be used to scale up bad pedagogical practices.** If AI is the new electricity, it will have a broad impact in society, economy, and education, but it needs to be treated with care.

[4] Readers may also be interested in "HUMAINT", an interdisciplinary JRC project aiming to understand the impact of machine intelligence on human behaviour, with a focus on cognitive and socio-emotional capabilities and decision making (see https://ec.europa.eu/jrc/communities/community/humaint).

1 Introduction

All human actions are based on anticipated futures. We cannot know the future because it does not exist yet, but we can use our current knowledge to imagine futures and make them happen. The better we understand the present and the history that has created it, the better we can understand the possibilities of the future. To appreciate the opportunities and challenges that artificial intelligence (AI) creates, we need both good understanding of what AI is today and what the future may bring when AI is widely used in the society. AI can enable new ways of learning, teaching and education, and it may also change the society in ways that pose new challenges for educational institutions. It may amplify skill differences and polarize jobs, or it may equalize opportunities for learning. The use of AI in education may generate insights on how learning happens, and it can change the way learning is assessed. It may re-organize classrooms or make them obsolete, it can increase the efficiency of teaching, or it may force students to adapt to the requirements of technology, depriving humans from the powers of agency and possibilities for responsible action. All this is possible. Now is a good time to start thinking about what AI could mean for learning, teaching, and education. There is a lot of hype, and the topic is not an easy one. It is, however, both important, interesting, and worth the effort.

Since 2013, when Frey and Osborne[5] estimated that almost half of U.S. jobs were at a high risk of becoming automated, AI has been on top of policymakers' agendas. Many studies have replicated and refined this study, and the general consensus now is that AI will generate major transformations in the labour market.[6] Many skills that were important in the past are becoming automated, and many jobs and occupations will become obsolete or transformed when AI will be increasingly used. At the same time, there has been a tremendous demand for people with skills in AI development, leading to seven figure salaries and sign-up fees. China has announced that it aims to become the world leader in AI and grow a 150 billion AI ecosystem by 2030. The U.S. Department of Defense invested about 2.5 billion USD in AI in 2017, and the total private investment in the U.S. is now probably over 20 billion USD per year. In 2017, there were about 1200 AI start-ups in Europe,[7] and the European Commission aims to increase the total public and private investment in AI in the EU to be at least 20 billion euros by the end of 2020.[8]

In limited tasks, AI already exceeds human capabilities. Last year, with just about one month of system development, researchers at Stanford were able to use AI to diagnose 14 types of medical conditions using frontal-view X-ray images, exceeding the human diagnostic accuracy for pneumonia.[9] In 2017, given no domain knowledge except the game rules, an artificial neural network system, AlphaZero, achieved within 24 hours a superhuman level of play in the games of chess, shogi, and Go.[10] In May 2018, Google CEO Sundar Pichai caused a firestorm when he demonstrated in his keynote an AI system, Duplex, that can autonomously schedule appointments on the phone, fooling people to think they are discussing with another human. In the midst of self-driving cars, speaking robots, and the flood of AI miracles, it may be easy to think that AI is rapidly becoming superintelligent, and gain all the good and evil powers awarded to it in popular culture. This, of course, is not the case. The current AI systems are severely limited, and there are technical, social, scientific, and conceptual limits to what they can do. As one

[5] Frey and Osborne (2013, 2017).

[6] E.g., European Political Strategy Centre (EPSC 2018), United States Government Accountability Office (GAO 2018), Finnish Steering Group of Artificial Intelligence Programme (2017), and UK House of Lords (2018).

[7] Data from the U.K. House of Lords Select Committee on Artificial Intelligence report (House of Lords 2018, 48).

[8] Artificial Intelligence for Europe (EC 2018b).

[9] Rajpurkar et al. (2017).

[10] Silver et al. (2017).

recent author noted, AI may be riding a one-trick pony as almost all AI advances reported in the media are based on ideas that are more than three decades old.[11] A particular challenge of the currently dominant learning models used in AI is that they can only see the world as a repetition of the past. The available categories and success criteria that are used for their training are supplied by humans. Personal and cultural biases, thus, are an inherent element in AI systems. A three-level model of human action presented in the next section suggests that norms and values are often tacit and expressed through unarticulated emotional reactions. Perhaps surprisingly, the recent successes in AI also represent the oldest approach to AI and one where almost all the intelligence comes from humans.

Instead of a beginning of an AI revolution, we could be at the end of one. This, of course, depends on what we mean by revolution. Electricity did not revolutionize the world when Volta found a way to store it in 1800 or when Edison General Electric Company was incorporated in 1889. The transformative impact of general purpose technologies becomes visible only gradually, when societies and economies reinvent themselves as users of new technologies. Technological change requires cultural change that is reflected in lifestyles, norms, policies, social institutions, skills, and education. Because of this, AI—now often called the "new electricity"—may revolutionize many areas of life when it is taken into use even if it keeps on driving its "one-trick" pony for the foreseeable future. Many interesting things will happen when already existing technologies will be adopted, adapted, and applied for learning, teaching, and education. For example, AI may enable both new learning and teaching practices, and it may generate a new social, cultural, and economic context for education.

Below we ask simple questions that illustrate the relevance of AI for educational policies and practices. Which vocations and occupations will become obsolete in the near future? What are the 21st Century skills in a world where AI is widely used? How should AI be incorporated in the K-12 curriculum? How will AI change teaching? Should real-time monitoring of student emotions be allowed in classrooms? Can AI fairly assess students? Do we need fewer classrooms because of AI? Does AI reduce the impact of dyslexia, dyscalculia, or other learning difficulties? These questions are simple to ask, and relevant for understanding the future of learning, teaching, and education. The answers, of course, are more complex.

The main aim of this report is to put these and other similar questions in a context where they can be meaningfully addressed. We do not aim to provide final answers; instead, we hope to provide background that will facilitate discussion on these and other important questions that need to be asked as AI becomes increasingly visible in the society and economy around us. To do this, we have to first open the "black box" of AI and peek inside. There are several things AI can do well, and many things it cannot do. At present there is an avalanche of reports and newspaper articles on AI, and it is not always easy to distinguish important messages from noise. It is, however, important to understand some key characteristics of current AI to be able to imagine realistic futures. In the next sections, we put AI in the context of learning, teaching, and education, and then focus on the specific form of AI, adaptive artificial neural networks, that have generated the recent interest in AI.

[11] Somers (2017).

2 What is Artificial Intelligence?

Artificial Intelligence has many different definitions. In the headlines of newspaper articles, AI is a machine that thinks, understands languages, solves problems, diagnoses medical conditions, keeps cars on the highways, plays chess, and paints impressionistic imitations of van Gogh paintings. AI is often defined as a computer system with the ability to perform tasks commonly associated with intelligent beings. As this definition somewhat problematically requires us to define intelligence and is inconveniently tautological, artificial intelligence is now commonly defined as a scientific discipline; as the activity that creates machines that can function appropriately and with foresight in their environment.[12] The first explicit definition of artificial intelligence was suggested in a funding proposal to the Rockefeller Foundation in 1955. It was based on the "conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it." This early definition rapidly led to deep controversies. In practice, the early developers of AI interpreted intelligence and thinking as mechanical processing of logical statements, thus, in effect, defining human intelligence as computation of truth values. This interpretation was historically aligned with logical positivism and attempts to formalize mathematics using purely syntactic means, but it also raised important questions about the philosophical foundations of AI.[13]

In the following section, we propose a different way to understand the nature of AI. It will help us locate the different capabilities of different types of AI in the context of learning. Adaptability, learning, and anticipatory action are commonly viewed as key characteristics of AI. We therefore use a theory of human action and learning as a starting point. For this we use a three-level model, along the lines of cultural-historical activity theory and a similar model proposed by Harré, Clarke and Carlo.[14]

2.1 A three-level model of action for analysing AI and its impact

Cultural-historical theory of activity distinguishes three hierarchically linked levels of human behaviour.[15] First, behaviour can be analysed as socially meaningful *activity* directed by culturally and socially constructed motives. Activity is realized through goal-oriented *acts* that essentially are ways of solving problems at hand that need to be solved to accomplish the activity. *Operations*, in turn, implement the acts in the present situation and concrete context, using the tools available. An important aspect of this three-level hierarchy is that the levels cannot be reduced to each other. We can explain the meaning of an *activity* only using social, cultural and historical terms that do not make sense at the level of *acts* or *operations*. For example, we can explain the object and motive of *activity* by saying that we are teaching children so that they become citizens, realize their potential as human beings, and get good jobs. The "content" of this activity—how it is translated into concrete *acts*—depends on social institutions, norms, social division of labour and knowing, the ways in which social production is organized, and many other similar things. Most importantly, we rarely are explicitly aware of all those social factors that shape our activities. Cultural norms, values, expectations, social

12 Nilsson (2009).

13 Since the early 1960s, the rather straightforward epistemological views adopted by the early AI developers were criticized mainly in reference to continental phenomenologists, including Husserl, Heidegger and Merleau-Ponty. See, e.g., Dreyfus (1979), Winograd and Flores (1986), Heinämaa and Tuomi (1989).

14 Socio-cultural activity theory, or more accurately cultural-historical activity theory, was inspired by the pedagogic studies of Vygotsky and his colleagues in the 1920s and 1930s. It became an important approach to study pedagogic methods and psychological theory in the Soviet Union in the subsequent decades. We use here the activity-theoretic model as described in Leont'ev (1978) and reinterpret its three-level structure using terminology from Harré et al. (1985).

15 We follow here the terminology from Leont'ev (1978).

institutions, and other essentially contextual factors shape our activities and provide a tacit normative, emotional, and anticipatory background that allows the ongoing stream of activity to go on. This is also the level that provides the foundation for ethics of action.

The relation between *acts* and *activity* is, thus, similar to the relation between words and utterances. We need words to express utterances, and *acts* to express *activity*. It is, however, impossible to understand the meaning of an utterance by adding up definitions of words. On the contrary, the sense of the word depends on its role in the context of an utterance. A written sentence needs words, and words need letters, but the meaning of a sentence cannot be found by studying letters or words. This, in effect, says that it is not possible to build models of human activity from bottom up, simply combining some elementary behavioural components.[16] *Activity*, properly understood, requires social and inter-generational learning, and the level of human activity cannot be accessed simply by empirical observation of human behaviour. The level of *acts*, in contrast, consists of externally and internally observable behaviour. Whereas the level of *activity* answers a socially, culturally, and historically meaningful question "why", the level of *acts* answers the question "what". This is also the level where we think with concepts, plan, and solve problems. If we call the level of *activity* a "cultural" level, the level of *acts* could perhaps be called "cognitive." A description of teaching at this level could be, for example, that "I am authoring course material for the class." The third level of *operations* addresses the question "how." It implements *acts* in concrete settings. For example, there are many ways to assess student skills, many kinds of homework, and many ways to deliver homework to students. This is the level where technology operates as a tool, and where behaviour can be best understood as routine and habit. A description of teaching activity at this level could be, for example, that "I'm inserting a picture on a slide."

Psychologists and learning theorists have focused on different levels of this three-level hierarchy during the last century. Behaviouristic and associationist theories of learning have addressed mainly the level of operations. Cognitivist and constructivist theorists have mainly addressed the cognitive level, with constructionists also emphasizing the material, affective, and social context. Socio-cultural theorists, in turn, have often focused on the social, cultural and materially embedded dimensions of knowing and learning. Figure 1 depicts these three levels and maps some well-known learning theorists to these levels.[17] Human learning occurs on all three levels of the activity hierarchy. When habit and routine hits an obstacle, we become aware of it, operation ceases, and action replaces it. We start to interpret the problem, and try to find a solution.[18] At this level, learning consists of problem solving, creative reframing, and formation of new anticipatory models. New ways of doing and thinking emerge, can be internalized, and can become the basis for new habits and routines. Lev Vygotsky, the founder of cultural-historical theory, however, also pointed to the importance of the social and cultural level of activities that shape human thinking and learning. Advanced forms of thought are made possible because they rely on culturally and historically developed stocks of knowing.[19] Cognitive level acts, thus, use resources from both the top level of activity and the bottom level of operations. Whereas Vygotsky emphasized

[16] This also means that any straightforward attempt to build artificial intelligence by combining elementary logical components into more complicated networks fails. For example, in an influential early contribution to AI, John von Neumann (1951) argued that it is possible to describe the human brain by interpreting neurons as logical switches and the brain as a complex network of such logical elements. Although von Neumann noted that we may need radically new forms of logic to do this, he also believed that the bottom-up approach is enough.

[17] Such a description is, of course, a simplification. In particular, Papert (1980; 1991) emphasized the affective and material dimensions of learning, and Piaget also wrote extensively about the social factors that underpin cognitive development, see, e.g. Cole and Wetsch (1996).

[18] This is known as Claparède's law of conscious awareness. It has informed many theories of learning from Dewey (1991) and Vygotsky (1986) to more recent ones, such as action research and action learning in organizational development (Lewin 1946).

[19] See, e.g., Vygotsky (1986), Vygotsky and Luria (1992), van der Veer & Valsiner (1994).

the influence of social and cultural factors in cognitive development, critical pedagogists such as Paulo Freire and newer activity theorists such as Yrjö Engeström have emphasized the role of learning in changing existing social practices.[20] Engeström, in particular, has highlighted the role of learning in the creation of new educational practices.[21]

Figure 1. Three levels of human and machine learning

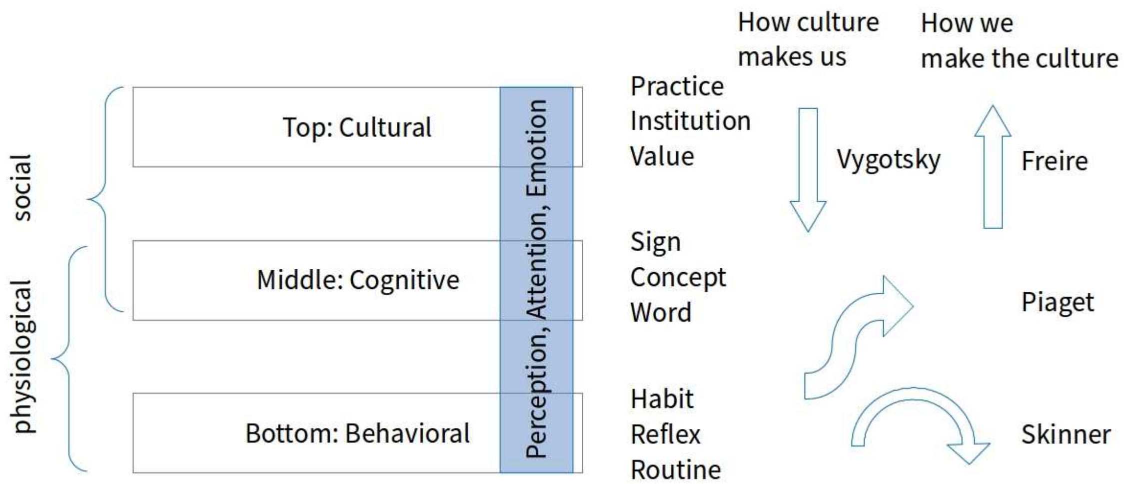

Source: Author's elaboration.

In this conceptual frame, learning at the level of activity can be understood as innovation and realization of imagined futures.[22] Possibilities that have been figured out at the level of cognition can start to change social practices and systems of activities, eventually leading to new motives and reasons that start to organize the society. Much of this activity-level development, however, is also emergent and unintended.[23] Social structures, practices and institutions get their shape as a result of complex ongoing social interaction and highly diversified interests and interpretations, and to a large extent remain unobservable for the members of society.

This three-level model provides a useful entry point for understanding artificial intelligence and its potential impact on human activities. When AI enters social practices at the level of *operations*, it augments and complements them, increasing the efficiency and effectiveness of current ways of doing things. When it enters at the level of *acts*, it replaces, substitutes, and automates acts that were previously done by humans. When it

20 See, e.g., Freire (1972) and Engeström (1996).

21 Engeström (1987). It should perhaps be noted that the "cognitive" level is in cultural-historical approaches understood as inherently social and materially embedded. Psychology has commonly viewed cognition from an individualistic point of view. To highlight the inadequacy of such an individualistic construct of cognition, terms such as "socially shared cognition," "situated cognition," "distributed cognition," and "extended cognition" are now commonly used. See, e.g., (Brown, Collins, and Duguid 1989; Cole 1986; Hutchins 1995; Mace 1977; Norman 1993; Suchman 1987; Salomon 1993).

22 In contrast to many common interpretations, innovation is here defined as creation of new technologically mediated social practice, see (Tuomi 2002a).

23 This observation underpins both Engels' (1966, chap. 5) description of the development of human cognition and Hayek's (1945) views on the impossibility to design policies that, in general, would produce better outcomes than free markets.

enters social practice at the level of *activity*, it transforms the system of motives, making current activities and specializations redundant and obsolete. For example, technical and routine skills emphasize the level of *operations*. Vocational education has traditionally focused on this level, teaching students how to use tools and domain-specific knowledge. The recent calls for competence-based education, in turn, emphasize problem solving, critical thinking, decision-making and analytical skills, focusing on the *cognitive* level. Entrepreneurial and innovation competences, highlighted in frameworks for key competences and 21st century skills, mainly address the opportunities for social and cultural change at the level of *activities*.

Consequently, learning at the level of *operations* requires data on the current concrete environment. This data can be generated using perception and physical interaction. Learning at the level of socially motivated *activity*, in contrast, requires knowledge about social systems of meaning. To gain such knowledge, communication, language, and dialogue become necessary. An important indicator of the current change in the dynamics of development is that whereas technology in the industrial age focused on tools for automating and supporting operations, the focus is now increasingly on technologies for social change. The three levels of activity have complex dependencies. In the course of historical development, what originally was a means may become an end in itself. "Zooming in" to modern social life, therefore, we may see a rather fractal structure or activities and acts. Using this three-level model of activity, it becomes, however, clear that different types of artificial intelligence and machine learning systems operate on different layers of this hierarchy. Most importantly, the level of meaningful activity, which according to socio-cultural theories of learning underpins advanced forms of human intelligence and learning, remains beyond the current state of the AI art. This paradigm is currently being explored in the field of Child-Robot Interaction and social robotics[24]. In the next section, we briefly outline the main characteristics of three different types of AI to locate their capabilities in this hierarchy, and discuss their potential impact.

2.2 Three types of AI

The history of AI can relatively cleanly be categorized into three alternative approaches: *data-based, logic-based*, and *knowledge-based.* The first of these is now also called artificial neural networks and machine learning. Perhaps surprisingly, the recent successes in AI also represent the oldest approach to AI.

2.2.1 Data-based neural AI

Mathematical models of neural networks were first developed by Nicolas Rashevsky in the early 1930s,[25] and they became famous when his student Walter Pitts interpreted biological neural networks in 1942 as networks of logical switches. The publication of these ideas by Warren McCulloch and Pitts[26] occurred at a time when Alan Turing had shown that formal logic can be mechanized and the first digital computers were being developed. It was therefore quickly recognized that all formal logical operations could be simulated by such neural networks. Brain started to look like a computer, and the computer became known as the electronic brain. This two-way metaphor has since then become widely influential. It underpins cognitive science and research in organizational

[24] See, for instance Vouloutsi, V. et al. 2016. Towards a synthetic tutor assistant: the EASEL project and its architecture. In Conference on Biomimetic and Biohybrid Systems (pp. 353-364). Springer, Cham.

[25] Early work on neural network models is reviewed in Rashevsky (1960). Rashevsky's work is little known among AI researchers, but his indirect impact is considerable. A collection of classic articles up to late 1980s is Anderson and Rosenfeld (1988).

[26] McCulloch and Pitts (1943).

information processing, and now influences economics, connectivist models of learning, and many areas of scientific and popular thinking.[27]

The present neural AI is to a large extent based on neural network models that were informed by neurobiology. An important early contribution was made by Frank Rosenblatt in 1958, when he—inspired by neuropsychologist Donald Hebb's idea that learning occurs in neural networks through synaptic modifications and economist Friedrich Hayek's work on distributed learning—suggested that learning in biological neural networks could be modelled as gradual change in network connections.[28] The multi-layer photo-perceptron described by Rosenblatt is in many ways identical to current state-of-the-art image processing neural networks.[29] Its main difference with today's neural AI systems is that modern systems have very many "neural layers," and "deep learning" in such multi-layer networks is done using machines that are about trillion times faster than the IBM 704 computer that Rosenblatt used for his experiments.

Figure 2: Organization of a perceptron

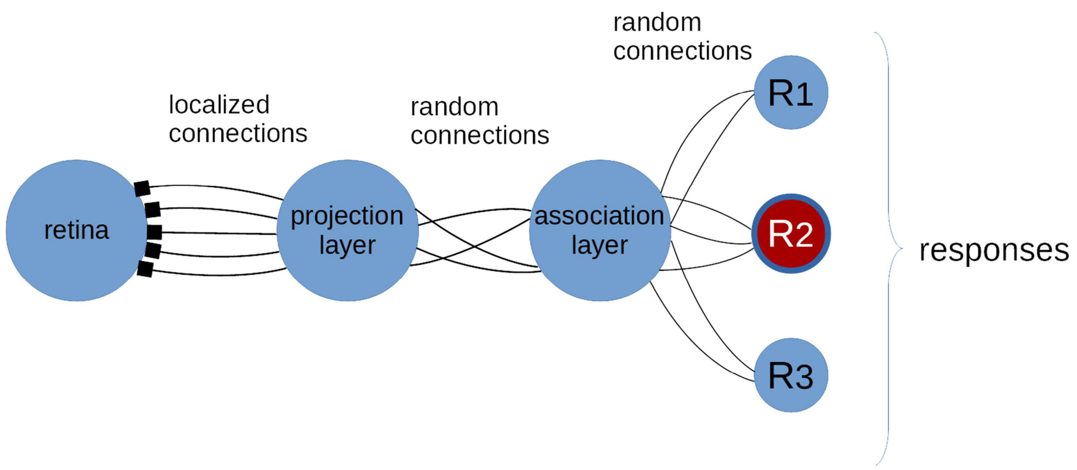

Source: Adapted from Rosenblatt, 1958.

The distinctive characteristics of most neural AI systems are their simple behaviouristic learning models, very high computational needs during learning, and their need for data. For these systems, the availability of data is the most critical success factor. Using

27 Many excellent histories of AI and cognitive science exist that describe the interdependent development of computers, cognitive psychology, and artificial intelligence. See, e.g., McCorduck (1979), Gardner (1987) and Boden (2016).

28 Rosenblatt (1958). Hebb was, in turn, influenced by Rashevsky's work on neural networks. Hayek's connectionist model of learning is described in Hayek (1952).

29 Current deep-learning architectures use computational "backpropagation" of output error during learning to adjust network weights. In contrast, Rosenblatt's perceptron used feedback connections from its output layer for learning how to separate different input patterns. Although deep-learning networks are essentially perceptrons, Rosenblatt used in 1958 a vacuum tube computer that able to do about 12,000 mathematical additions or multiplications per second (12kFLOPS). Google's newest tensor processing compute "pods," announced in May 2018, can run more than hundred petaflops when training a machine learning system. That is 100,000,000,000,000,000 multiplications per second. This, in itself, is superhuman: If every person on Earth would make one multiplication every second, about ten million planet Earths would be needed to achieve the same computational capability.

biological terminology, they could be called “datavores.” Because of this, we call this the “data-based” approach to AI.[30]

2.2.2 Logic- and knowledge-based AI

Neural network models were popular in the 1950s and 1960s. They were also a key area of study—among learning, language, creativity, and abstraction—in the Dartmouth summer research project in 1956 that established the term Artificial Intelligence. Although work continued on neural networks, research on AI soon moved to “symbolic processing.” As mathematicians and logic-oriented philosophers had since Hilbert and Russell believed that logical truths could be derived by formal manipulation of sentences, it was apparent that computers could do all those inferences that are logical. A pioneering effort in this line of AI was the Logic Theorist, developed by Allen Newell, John Shaw, and Herbert Simon over the Christmas break in 1955. It was able to manipulate logical statements and derive proofs for logical theorems, and its creators were certain that they had produced a machine that thinks. The Logic Theorist was soon followed by the General Problem Solver that was supposed to be able to solve any logically well-defined problem that had a solution. This logic-oriented approach to AI was the dominant one from the late 1950s to early 1970s.[31]

By the 1970s it was generally acknowledged that human thinking cannot be simulated just by formal manipulation of logical statements. As a result, domain-specific knowledge and different ways of representing knowledge became the central focus of AI research. This led to what is now known as "expert systems" or, more broadly, knowledge-based systems. Early examples of these include the SHRDLU natural language understanding program and the MYCIN medical diagnostic system that recommended antibiotics and their dosage based on the symptoms and the patient. Knowledge-based systems typically consisted of a relatively general "inference engine" and a domain-specific "knowledge base" that was used to make inferences based on human input. In particular, in expert systems, domain knowledge tried to imitate knowledge structures used by human experts. Expert systems were very popular in the 1980s, with two thirds of Fortune 500 companies using them in the daily activities. Since then they have been widely used in various sectors of economy, for example in the financial sector, logistics, semiconductor chip design, manufacturing planning, and business process automation. Many expert systems have also been developed for learning and education since the early 1980s.

The interest in knowledge-based AI waned towards the end of 1980s as it became clear that the development of domain-specific knowledge bases required specialized knowledge engineers, and also because the spread of computer networking and the Internet shifted the interests towards system integration and automation of routine business processes. Many ideas from stand-alone expert systems are now widely used in standard programming environments. As the boom of knowledge-based AI decayed at the end of the 1980s, neural AI research became again popular for a few years. Difficulties associated with parallel programming and system integration, however, kept most neural AI systems in university laboratories, and attention moved to new areas such as mobile computing and the World-Wide Web.

---

[30] It should perhaps be noted that the currently popular neural AI models require huge amounts of data because they use learning models that can easily be implemented using digital computers and algorithms. More effective neural models can be implemented using analog computation and measurement-type computers (Tuomi, 1988). The ”third wave” DARPA AI Next campaign, announced in September 2018, and many neural chip initiatives aim to address this challenge.

[31] C.f. McCorduck (1979).

From a practical point of view, both logic-based and knowledge-based approaches in AI focus on the cognitive level of activity hierarchy. They also interpreted cognition in a purely individualistic way. Logic-based AI tried to develop general algorithms for thinking that manipulate symbols, arguing that this is what also humans do. Whereas logic-based systems focused on general problem-solving processes, knowledge-based approaches used simple models of inference and more elaborate representations of domain-specific knowledge, arguing that effective decision-making requires more knowledge than logic. In contrast, machine learning and artificial neural networks typically use learning models that can be characterized as behaviouristic. These systems are typically provided with vast amounts of data and pre-defined criteria for optimal response. In these systems, the algorithms do not try to imitate human intelligence; instead, they define strategies for adapting system output to expected output using extensive amounts of what is called "training data". In some applications, such as games, this training data can be automatically generated; in most currently important neural AI systems the data are provided by humans. For example, the development of state-of-the-art image recognition AI systems now, to a large extent, relies on the publicly available ImageNet database that consists of 14 million images. The labelling of objects in these images was done in 2007-2010 using the Amazon's Mechanical Turk crowdsourcing platform by 48,940 people in 167 countries.

2.3 Recent and future developments in AI

The recent interest in AI results from three parallel developments. First, increasingly realistic computer games have required specialized graphics processors. When the PC graphics card manufacturer *Nvidia* published the *CUDA* programming interface to its graphics accelerator cards in 2007, fast parallel programming became possible at low cost. This allowed researchers to build neural network models that had many connected layers of artificial neurons and large numbers of parameters that the network could learn. Second, huge amounts of data have become available as computers and computer users have been networked. The digitalization of images, videos, voice and text has created an environment where machine learning can thrive. This has allowed AI researchers to revisit old artificial neural network models, training them with very large datasets.

Somewhat surprisingly, these huge data sources have proven to be enough for some of the hard problems of AI, including object recognition from digital images and machine translation. Whereas it was earlier believed that computers need to understand language and its structures before they can translate text and speech from one language to another, for many practical uses it is enough to process millions of sentences to find out the contexts where words appear. By mapping words into high-dimensional representational spaces, enough of this contextual information is retained so that translation can be done without linguistic knowledge. A common approach is to use the publicly available *GloVe* word representations that have been developed using text corpora that contains up to 840 billion word-like tokens found on documents and content on the Internet, subsequently translated to a vocabulary of over 2 million words.[32] Using this dataset and machine learning algorithms, the words have been mapped into points in a 300-dimensional vector space.[33] The location and geometric relations between words in this space capture many elements of word use, and can be also used as a basis for translation from one language to another. Although such a purely statistical and data-

[32] See Pennington et al. (2014)

[33] There exist several versions of *GloVe* vectors. Pre-trained *GloVe* vectors, trained using different corpora, can be downloaded from https://nlp.stanford.edu/projects/glove/

based approach is not able to comprehend new or creative uses of language, it works surprisingly well in practice.

Third, specialized open source machine learning programming environments have become available that make the creation and testing of neural networks easy. In most current neural AI models, learning occurs by the gradual adjustment of network weights, based on whether the network makes right predictions with the training data. A central task in such learning is to propagate information about how important each neuron's activity is to right and wrong predictions made by the network. When an active neuron is associated with a wrong prediction, the activity of the neuron is decreased by decreasing the weights of its incoming connections. As there can be very many layers of neurons and many connections between neurons, this is a task that is difficult even for powerful traditional computers. The influence of each neuron to the prediction can, however, be computed using the chain rule of calculus, propagating the information from the output layer of the network layer-by-layer towards the input layer. This is known as "backpropagation" of error.[34] Although the computation of network weights using this method may involve hundreds of millions of computations in state-of-the-art networks, current neural AI development environments can do this with a couple of lines of program code.

These three trends started to come together around 2012. In that year, a multilayer network trained using *Nvidia*'s graphics processor cards showed outstanding performance in an image recognition competition. The competition was based on the ImageNet database that contains about 14 million human-annotated digital images. The ImageNet Large Scale Visual Recognition Challenge (ILSVRC) is now one of the main benchmarks for progress in AI. Its object detection and classification challenge uses 1.2 million images for training, with 1,000 different types of objects. In 2017, the best neural network architectures were able to guess the correct object category with 97,7 per cent "top-5" accuracy, meaning that the correct object class was among the five most probable classes as estimated by the network. The rapid improvement in object recognition can be seen in Figure 3 that gives the top-5 error rates of the winners over the years.

Figure 3: Error rates in the ImageNet ILSRC object recognition competition

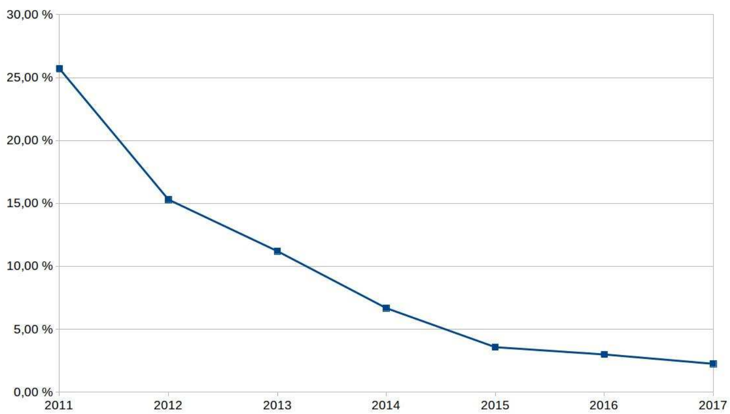

Source: Data compiled from imagenet.org

[34] This method was first explicitly described by Seppo Linnainmaa in 1970 in his master's thesis at the University of Helsinki, but it became widely known in the mid-80s, as part of the parallel distributed processing approach to AI (Rumelhart and McClelland 1986). The difficulty of propagating prediction error signals in complex multilayer neural models limited the use of this methodology until graphics processors started to be used for "deep learning."

The resurrection of neural AI has partly been caused by the availability of data, such as digital images, electronic texts, Internet search patterns, and social network content and linkages. Recent developments, however, have also been driven by the fact that these huge datasets are difficult to analyse and utilize with traditional computing. Machine learning both requires big data but it also makes large quantities of data usable and valuable. There are therefore large commercial incentives in using machine-learned models for processing data that cannot practically be processed using more traditional approaches.

2.3.1 Models of learning in data-based AI

Almost all current neural AI systems rely on what is called a supervised model of learning. Such "supervised learning" is based on training data that has been labelled, usually by humans, so that the network weights can be adjusted when the labels for training data are wrongly predicted. After a sufficient number of examples are provided, the error can in most cases be reduced to a level where the predictions of the network become useful for practical purposes. For example, if an image detection program tries to differentiate between cats and dogs, during the training process someone needs to tell the system whether a picture contains a cat or a dog.

A practically important variant of supervised learning is called "transfer learning." A complex neural network can be trained with large amounts of data, so that it learns to discern important features of the data. The trained network can then be re-used for different pattern recognition tasks, when the underpinning features are similar enough. For example, a network can be trained to label human faces with millions of images. When the network has learned to recognize the faces that have been used for its training, its deep layers become optimized for face recognition. The top levels of the network can then relatively easily be trained to detect new faces that the system has not seen before. This drastically reduces the computational and data requirements. In effect, AI developers can buy pre-trained networks from specialized vendors, or even get many state-of-the-art pre-trained networks for free and adapt them to the problem at hand. For example, the GloVe vectors, available from Stanford University, are commonly used as a starting point for natural language processing, and Google's pre-trained Inception image processing networks are often used for object recognition and similar image processing tasks.

Supervised learning systems can produce statistical guesses of which of possible pre-given class a specific given input data pattern belongs. Supervised learning, thus, assumes that we already know what categories input patterns can represent. This is the most frequently used learning model in AI today because for practical purposes it is often enough to classify patterns into a set of pre-defined classes. For example, a self-driving car needs to know whether an object is a cyclist, truck, a train, or a child. Technically, supervised learning creates machines that map input patterns into a collection of output classes. Their intelligence, thus, is similar to simplest living beings that can associate environmental conditions with learned behaviours. In psychology, these learning models underpin the Pavlovian theory of reflexes and, for example, Skinnerian reinforcement learning. As Vygotsky pointed out in the 1920s, this type of learning represents the developmentally simplest model of learning, and both pigeons and humans are well capable of it.[35]

[35] Tuomi (2018).

A particular challenge of supervised learning models is that they can only see the world as a repetition of the past. The available categories and success criteria that are used for their training are supplied by humans. Personal and cultural biases, thus, are an inherent element in AI systems that use supervised learning. The three-level model presented above suggests that norms and values are often tacit and expressed through unarticulated emotional reactions. It is, therefore, to be expected that supervised learning models materialise and hardwire cultural beliefs that often remain otherwise unexplored. In somewhat provocative terms, supervised learning creates machines that are only able to perceive worlds where humans are put in pre-defined boxes. From ethical and pedagogic points of view this is problematic as it implies that in interactions with such machines, humans are deprived of agency powers that allow them to become something new and take responsibility of their choices.

Many unsupervised or partially supervised neural learning models have been developed since the 1960s, some of which are also currently being developed and applied. Increasing computational power has also allowed researchers to use simple pattern-matching networks as components in higher-level architectures. For example, Google's *AlphaZero* game AI uses "reinforcement learning" where the system generates game simulations and adjusts network weights based on success in these games. Inspired by Skinnerian models of operant conditioning, reinforcement learning amplifies behaviour that leads to outcomes that are defined as positive. A variant of reinforcement learning is known as generative adversarial networks, or GANs, where one network tries to fool another to believe that the data it generates actually comes from the training data set. This approach has been used, for example, to create synthetic images of artworks and human faces that an image recognition system cannot distinguish from real images[36]. It is also commercially used for product design, for example in the fashion industry. A variation of GAN is called "Turing learning," where the system that learns is allowed to actively interact with the world in trying to guess whether the data comes from the real environment or from a machine.[37]

2.3.2 Towards the future

As some economists, philosophers, and scientists have made high-profile statements about the forthcoming emergence of super-intelligent AI systems that eventually may replace humans in many areas of human life, it is perhaps useful to note that most current AI learning models represent cognitive capabilities that most closely resemble biological instincts. Many predictions about the future of AI have been based on extrapolations of historical technical development, and in particular estimates of the continuation of "Moore's Law" in computing, with little concern about differences between advanced forms of human learning and the more elementary capabilities of association. Human learning requires many meta-level competences. In particular, for humans it is important to know what counts as knowledge, how to go on in acquiring, creating, and learning knowledge, how to regulate cognition, attention and emotion in learning

[36] https://www.nytimes.com/interactive/2018/01/02/technology/ai-generated-photos.html and https://www.hs.fi/tiede/art-2000005734015.html

[37] This approach is based on a simplified version of the imitation game suggested by Turing in 1950. Turing argued that if a machine is able to fool a human in this game, the question whether machines can think becomes redundant. This is now known as the "Turing test." The original imitation game, however, is more sophisticated than its popular versions and the model used in Turing learning. The game tries to distinguish a man and a woman, and tries to see if, based on answers to interrogator's questions, a man makes as many errors in detecting a man who imitates a woman than he makes detecting a machine who imitates a woman. Turing's test, thus, measures whether two obviously different humans (a man and a woman) are no more different than a machine and a human when they can be observed only using teletype messages. The philosophical foundation for the test is logical positivism, which essentially claims that if something walks and talks like a duck, it has to be a duck. In the imitation game, the duck is in a closed room with a teletype printer, and the types of ducks that are allowed in the game are strongly constrained (Heinämaa and Tuomi 1989).

processes, and what the social and practical motivation for learning is. As Luckin has recently well pointed out, at present AI lacks most of these meta-cognitive and regulatory capabilities.[38]

It is important to note that the future of the current AI boom will to an important extent be determined by developments in chip design. For almost fifty years, developments in processor and memory chips were driven by rapid continuous improvements in miniaturization of component features on semiconductor chips. During the last ten years it has become increasingly accepted that this development is about to end, and new approaches are needed to keep the semiconductor industry growing. Neural AI addresses this "post-Moore" era by shifting development towards new computing models, including analog computing. This represents a major discontinuity in the technological foundations of knowledge society.[39]

In practice, most AI experts work with "narrow AI," in contrast with "general AI" that would have capabilities similar to humans. In setting up the first Dartmouth summer project on artificial intelligence, the leading researchers believed that computers will soon be intelligent. Such expectations seem to be unrealistic also today. Although it might be possible to develop AI systems that have capabilities that more closely resemble human intelligence, current AI systems use rather simplified models of learning and biological intelligence. Most current AI systems rely on essentially reflexological and behaviouristic models of learning, popularized by Pavlov and Thorndike at the beginning of the 20th century. They could perhaps therefore better be described as mechanical instincts, instead of artificial intelligence.[40] Despite these limitations, the potential of AI in education has been widely recognized during the last three decades. Although the impact on classrooms has been relatively minor, the recent developments suggest that the situation may change. In particular, AI-based systems can become widely used as systems that support teachers and learners. AI can also rapidly change the economy and job market, creating new requirements for education and educational systems.

2.4 AI impact on skill and competence demand

One of the key roles of modern educational system is that it creates competences that allow people to participate in the economic sphere of life. The history of educational systems is closely linked with the development of the industrial society, and wage labour is still a central organizing principle in industrial societies and their everyday life. In high-level policy discussions, education is therefore often understood as a source of employment. Education, in this interpretation, is a key driver of economic productivity and competitiveness, and educational policies are framed in the context of economic growth. It is therefore important to ask also in the context of educational policies how AI will transform work and employment. For economists, a central question has been whether automation and computerization increases unemployment. As machines increase

[38] Luckin (2018).

[39] The claims of rapidly approaching ”singularity” and ”superintelligence,” therefore, are based on somewhat questionable extrapolations of historical trajectories. For more detailed analysis of these developments, see Tuomi (2002b, 2009). In particular, the energy consumption of neural AI systems will be a critical factor for the wide use of AI.

[40] Most current AI researchers are rather agnostic concerning the future of general AI. Historically, many AI researchers have thought that Turing's test is important for AI because it is aligned with the formalist idea that all truths are statements that at least in principle can be typed on a teletype keyboard. From this point of view, it seems irrelevant that the experimenter is prohibited from opening the door and looking inside to check whether there is a human or a machine. It can also be shown that success in the Turing test does not mean that a machine would have similar capabilities for thinking as humans. A finite collection of Google Duplexes do not make a dialogue in mathematical sense. More generally, it can be shown that any finite collection of simulations cannot generate an accurate model of biological systems (Rosen 1985; Louie 2009). This, however, requires the use of mathematical formalism known as category theory.

labour productivity, fewer human workers are needed to maintain production. Unless the demand for products grows enough, unemployment grows.

In reality, this simple model is, of course, too simple. If machines replace some jobs, people may move to other jobs. In general, this is what happened in the last century when agricultural and industrial jobs were automated, and labour moved to services. There are many influential studies that have verified this pattern.[41] Using historical data, they typically conclude that more technology and labour productivity growth have not increased aggregate unemployment. On the other hand, it is well known that an important reason why automation has not generated persistent unemployment is population growth that has continuously increased demand for industrial products and services. Many other factors, such as education, globalization, increased consumption of non-renewable natural resources, as well as developments in science and healthcare have been involved in the 20th century economic growth, and it is, therefore, difficult to make predictions about the future using historical patterns.

Although some influential studies claim that automation has not generated unemployment, it may therefore be useful to recall also the history of industrialization and its social consequences. Industrialization led to social upheavals and revolutions from Prussia to Mexico, Russia, and countries around the world, often with brutal outcomes. Millions of lives were lost. People flocked into cities, and at the turn of the 20th century authors such as Jack London still described in detail the dismal conditions of wage-slaves in the Oakland docks. As the economic system now operates on a global scale, the impact of AI cannot easily be studied on a national scale, where useful econometric data typically is available. Although country-level data can be aggregated, for example, for cross-national comparisons, the global and networked knowledge economy is not just a collection of economically integrated national economies.[42] In considering the social, economic and human impact of AI and its relation to educational policies, a broad view on social change is necessary.

2.4.1 Skills in economic studies of AI impact

Much of the current economic research on the future of work and the impact of AI starts from analysing the impact of computers on skill demand. It is, therefore, important to understand how skills and work tasks have been interpreted in these studies. Below, we put these econometric studies in the context of the three-level model presented above

41 These include, for example, Autor, Levy and Murnane (2003), Acemoglu and Restrepo (2016), and, in a more pessimistic vein, Brynjolfsson and McAfee (2012). Autor, in particular, has argued that the main impact of automation has been in the polarization of labour markets. He also argues that the use of AI will increase the comparative advantage of humans in tasks that require problem-solving skills, adaptability, creativity, flexibility, and common sense (Autor 2015). A recent collection of articles on the economy of AI is available from the US National Bureau of Economics Research (Agrawal, Gans, and Goldfarb 2018). Many of these studies, however, could be put in a somewhat different light by looking time use and hours worked in the economy per capita. For example, in Finland the time used for paid labour has decreased about one fifth per capita in the last forty years.

42 The global and networked character of knowledge economy poses some quite deep methodological challenges here. We have extensive economic data on the national level, and it is therefore natural to assume that we should use those data as a starting point to study the economic impact of computerization and AI. The available data, however, do not necessarily capture the non-local and functional aspects of economy. In biology, the observation that those aspects of living systems that make them "alive" cannot be described using data on their constituent components led in the 1950s to "relational biology." It focuses on the functional organization of biological systems instead of their various material implementations (Rosen 1958, 1991; Rashevsky 1954, 1972). In particular, Robert Rosen argued that dynamic models, such as those used in physics and economics, are not able to capture the essence of biology as systems are alive because of complex networks of interrelated functions. A category theoretic formalism is needed to model such systems (cf. Louie 2009).

(see 2.1), showing that different types of AI have capabilities on different levels of this model.

Many of the influential econometric studies use the U.S. Occupational Information Network (O*NET) database as a starting point.[43] O*NET contains now about 1000 occupational definitions to help students, job seekers, and educators to understand skill requirements and work content in different occupations. An example of the task structure of one occupation, "Middle School Teachers, Except Special and Career/Technical Education," is shown in Figure 4.

Figure 4: O*NET task and skill structure for Middle School teacher occupation

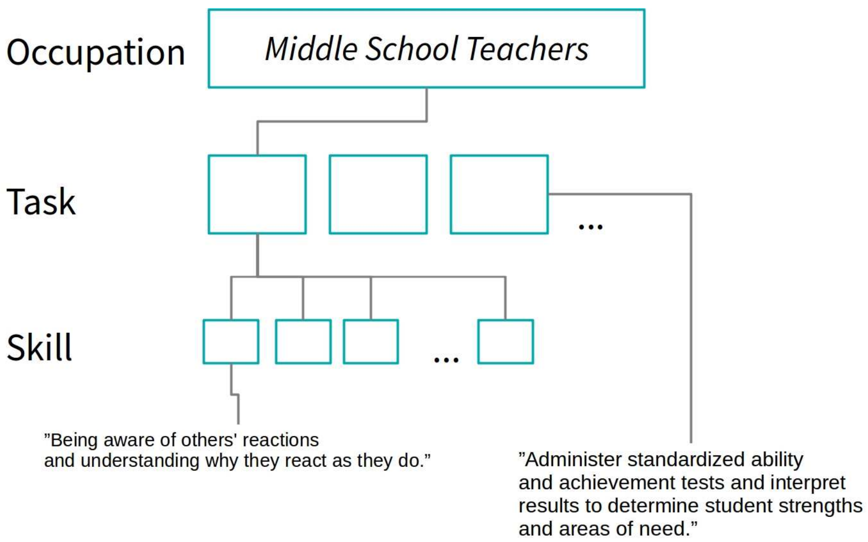

Source: Based on O*NET (www.onetonline.org)

The path-breaking study by Frey and Osborne asked experts in robotics and AI what are those technical bottlenecks that limit the automation of work tasks.[44] Using these automation bottlenecks as a starting point, they then asked the experts to classify a set of O*NET occupations based on whether automation of their tasks seemed possible. Those jobs that didn't contain hard-to-automate tasks were classified as having a high risk of being automated. One important outcome of the Frey and Osborne study is that it predicted that about half of U.S. occupations is at high risk of being automated in the near future using current technologies. Whether this estimate is accurate or not, it still highlights the point that educational systems will be under considerable pressure to address this wide-spread change. Traditional educational planning has tried to predict the future demand for different types of education based on estimated labour market developments. Frey and Osborne show that AI will have radical impact on the labour market, and create discontinuities in many trends that currently underpin educational planning and policies. We, therefore, need to reconsider both the content and the functions of education in this new environment.

[43] O*NET data can be accessed online at http://www.onetonline.org/.
[44] Frey and Osborne (2013).

2.4.2 Skill-biased and task-biased models of technology impact

Many earlier studies on the impact of computers and automation were based on skill-biased models of technological change. In skill-biased models, jobs that do not require educated, experienced, and skilled workers are susceptible to automation. In such models, computers are expected to be used mainly for tasks that require limited skill. It becomes then natural to assume that to avoid unemployment people need more and higher-level education. In contrast, recent studies on computerization have adopted a task-biased approach. It assumes that those tasks that can be exactly described can be programmed with a computer. In these studies, occupations that consist of routine tasks are susceptible to automation. This has typically led researchers to assume that occupations that require human-like intelligence are not susceptible to automation. The implication for educational policy could be that education should focus on non-routine cognitive tasks, often labelled as 21st century skills. Frey and Osborne used a task-biased model, but they argued for a different approach. In their view, the impact on AI and robotics should be studied based on current technological bottlenecks. AI is rapidly becoming able to perform tasks that have traditionally been understood to require human cognition. According to Frey and Osborne, it is therefore important to ask experts what computers cannot do. All those tasks where technical bottlenecks do not exist may be automated, and if an occupation consists of such tasks, it is susceptible to automation.

Beyond such an occupation-level analysis, it is interesting to drill down to specific occupations and consider how AI could change them. In Table 1 we do this for the O*NET Middle School Teachers. The table lists some of the teacher's tasks, as they are listed in O*NET, in their order of importance. The potential impact of AI on tasks is based on author's estimate, and should be taken as indicative.

Table 1: Potential impact, middle-school teacher tasks

	Task	AI impact
1	Adapt teaching methods and instructional materials to meet students' varying needs and interests	High
2	Establish and enforce rules for behaviour and procedures for maintaining order among students	?
3	Confer with parents or guardians, other teachers, counsellors, and administrators to resolve students' behavioural and academic problems	Low
4	Maintain, accurate, complete, and correct students records as required by laws, district policies, and administrative regulations	High
5	Prepare, administer, and grade tests and assignments to evaluate student's progress	High
6	Prepare material and classrooms for class activities	Medium
7	Instruct though lectures, discussions, and demonstrations in one or more subjects, such as English, mathematics, or social studies	Medium
8	Establish clear objectives for all lessons, units, and projects, and communicate these objectives to students	Medium
9	Assist students who need extra help, such as by tutoring, and preparing and implementing remedial programs	High
10	Assign lessons and correct homework	High
11	Enforce all administration policies and rules governing students	Medium
…		
15	Meet or correspond with parents or guardians to discuss children's progress and to determine priorities and resource needs	Medium

Source: I. Tuomi's estimate

Looking at this table, one might wonder why many of the listed tasks seem to be susceptible to automation. One explanation could be that technology has now advanced to a level where also some demanding human cognitive activities, such as performing tasks related to teaching, administrative and communication tasks, can be performed by computers. A more critical view might be that teachers are in the current educational systems burdened with rather mechanical tasks. The list of high-importance tasks also reflects deep beliefs about the functions of education and the social institutions around it. For example, comparative high-stakes testing and assessment of achievement may be highly important when educational systems are used for social selection. In educational systems that emphasize development and, for example, social competences, formative assessment might be higher on the list.

2.4.3 AI capabilities and task substitution in the three-level model

If we use the three-level model of activity (see 2.1), the econometric studies on future work and skill demand appear in a new light. First, as von Neumann argued half a century ago, if we can exactly and unambiguously describe a task, it is possible to program a computer to perform the task.[45] Von Neumann was talking about the capability of computers to simulate any system that can be simulated, although he also noted that we may need new forms of logic and new formalisms to do this. A simple conclusion from this might be that there are no fundamental technical bottlenecks that would make automation impossible. Indeed, well-known authors such as Kurzweil and Bostrom seem to adopt such a view.[46]

In the context of the three-level model of human activity and cognition, the level of *activity* is not directly accessible for individual human cognition. It provides a tacit cultural and social background which makes activities meaningful. As Polanyi and Hayek, among others, have emphasized, much of the knowledge that underpins social activity is contextual, distributed, embedded in social institutions and technologies, and enacted in practice.[47] It seems, therefore, that this social and cultural layer can, at best, be only partially articulated and made explicit. If von Neumann was right, and everything that can be explicitly described can be computed, it seems that the level of *acts* and *cognition* is the level where computing could have it main impact. This, indeed, is the level where most logic- and knowledge-based AI work has been done. In this view, the important bottleneck is not technical; instead, it is representational. Although we may convert some tacit knowledge to explicit knowledge, this requires a context that necessarily remains unarticulated.

An alternative way to approach the question of task substitution is to start from the statement by one of the leading AI experts, Andrew Ng. He summarizes the capabilities of neural AI and machine learning is a compact way:

> *"If a typical person can do a mental task with less than one second of thought, we can probably automate it using AI either now or in the near future."*[48]

This highlights the point that current neural AI and machine learning systems address the bottom level of the three-level hierarchy. Tasks that require habit formation and reflex reaction are well suited for supervised learning models.

45 Von Neumann (1951, 310).
46 Kurzweil (1999), Bostrom (2014).
47 Cf. Polanyi (1967), Hayek (1952).
48 Ng (2016).

Yet, there is a caveat to Ng's definition: What counts as a "typical" person? Many "less-than-one-second" human tasks require years of learning. Some of these, for example, learning to walk, are rather behavioural, and can also be learned by AI-supported robots. Many of these tasks, however, also require long periods of cultural and social accommodation. It may, therefore, be possible, for example, to use AI to simulate a concert pianist playing Bach's Goldberg variations, and generate music that sounds similar. Meaningful interpretation of Goldberg variations, however, requires extensive knowledge about cultural history, reflection of the relation of Bach to other composers, knowledge about subsequent interpretations, as well as years of training. It may take less than a second to play a note, but it may take many years to be able to do that. Although it is clear that a concert pianist may not be a "typical" person, many very typical everyday tasks require similar enculturation and learning. Indeed, a central claim in Vygotsky's theory of cognitive development in the early 1930s was that those advanced cognitive capabilities that distinguish humans from other animals are exactly those capabilities that cannot be described as simple reflexes, but which require social and cultural learning. This suggests that Ng is really talking about instinctive behaviour, instead of intelligence. The fundamental automation bottleneck, therefore, is not about technical capability. It is in the qualitative difference between observed behaviour and its meaning. As soon as the meaning of activity is fixed, we may be able to mechanize the behaviour and learn to do this using a large number of examples of such behaviour. Many forms of human learning and advanced forms of human cognition, however, are based on creating meaning where it was not before. To address such areas of human intelligence, AI researchers will need models of intelligence that far exceed those that are currently used in artificial intelligence.

2.4.4 Trends and transitions

Econometric studies on the effects of automation, computerization, and AI are therefore interesting and important but they do not capture the future well. In general terms, there is no obvious reason why historical trends would remain valid in socio-economic transitions. Econometric models may be important for understanding the present in the light of the past, but they can predict the future only if nothing important changes. This is simply because these models are based on data, and we don't have empirical data about the future.[49] They are, however, important because they suggest that we can predict the future in a very specific way: If nothing important changes, wide use of already existing AI technologies will imply a future that will be very different from what it used to be. This somewhat paradoxical result shows that, if for nothing else, this is because paid labour used to be such a central factor in shaping the industrial age, its institutions, and our everyday life.

[49] More detailed discussion on this problem can be found in Tuomi (2012).Productivity is also often difficult to measure when quality change and innovation are important. This will be the case for AI, in particular, as it does not only replace existing functions but transforms existing ones and creates novel productive tasks. For example, the impact of computers has been measured using "quality adjusted prices" that take into account developments in technical characteristics of computer equipment, such as processor clock speed, memory bandwidth, and number of transistors on chips. Because of the almost exponential improvements in many of these technical features, computers have become important factors in productivity growth. It is, however, not clear how such productivity measures correlate with common-sense ideas of productivity. For example, it is difficult to say how much more productive a person is writing texts with a computer that has a thousand times faster processor than two decades ago.

2.4.5 Neural AI as data-biased technological change

A recent study by Nedelkoska and Quintini[50] at the OECD provides a good review of econometric research on the impact of automation, and extends the Frey and Osborne study using the results of the OECD Survey of Adult Skills (PIAAC). Nedelkoska and Quintini matched the technical bottlenecks from Frey and Osborne to PIAAC variables on job tasks, such as frequency of complex problem solving and advising or teaching others. The variables used by Nedelkoska and Quintini are shown in Table 2. For the overall sample of 32 countries, they found that the median job had a 48 per cent probability of being automated, with large variations across countries.

Table 2: Technical bottlenecks for automation

Engineering bottlenecks	**Variable in PIAAC**	**Description**
Perception manipulation	Fingers (dexterity)	How often - using skill or accuracy with your hands or fingers?
Creative intelligence	Problem solving, simple	How often - relatively simple problems that take no more than 5 minutes to find a good solution?
	Problem solving, complex	Problem solving - complex problems that take at least 30 minutes thinking time to find a good solution?
Social intelligence	Teaching	How often - instructing, training or teaching people, individually or in groups?
	Advise	How often - advising people?
	Plan for others	How often - planning the activities of others?
	Communication	How often - sharing work-related information with co-workers?
	Negotiate	How often - negotiating with people either inside or outside your firm or organization?
	Influence	How often - persuading or influencing people?
	Sell	How often - selling a product or selling a service?

Source: Adapted from Nedelkoska & Quintini, 2018

Economists have used both skill-biased and task-biased models to study the impact of automation, computers and AI. Neural AI and machine learning, however, do not fit these models well. The critical bottleneck is not whether a task is routine or non-routine, or whether it requires complex problem solving; instead, it is whether the task can be learned by a computer. This, in turn, depends on whether there are data that can be used for learning. The impact of AI on occupations can, therefore, best be understood in a "data-biased" model. If data are available and history repeats itself, current machine learning algorithms can at least in principle simulate the past. To the extent that learning, innovation and knowledge creation is about combining existing pieces of knowledge, machines may also be able to do that. From a technical point of view, such operations are purely syntactic. There are good reasons to expect that social, economic, and cognitive processes, as well as other systems that can be called living, cannot be simulated using such an approach.[51]

2.4.6 Education as a creator of capability platforms

As a result, AI will probably have its biggest impact when it is used to augment human cognition, and in supporting human learning and knowing. This suggests a general principle of keeping humans in the loop when AI is used for educational purposes and in

50 Nedelkoska and Quintini (2018).

51 Sophisticated mathematical formalisms are needed to appropriately study the possibility of building computational models of human cognition, and many AI experts remain agnostic whether this will ever be possible. See, e.g., Rosen (1998), Loiue (2007, 2009).

educational settings. Assuming that some occupations, perhaps such as truck drivers, data entry keyers or utilities meter readers, will become obsolete in the near future, an important question for education policy is how people in these occupations can move to new jobs. A recent study by Royal Bank of Canada (RBC) focused on this question, locating six skill clusters that can be used to group occupations in Canada.[52] Also this study used O*NET data, but focused on skills, instead of tasks as was done in the Frey and Osborne study. The RBC study argued that as many occupations overlap in their skill requirements, it is relatively easy to complement skills within these clusters in ways that enable people to move to new jobs when their old jobs become automated. These clusters are shown in Table 3. This approach, thus, complements the view that there are key transversal skills and competences that are necessary for future.

Table 3: Skill clusters and probability of disruption in their occupations

Skill cluster	**Description**	**Probability of disruption**
Technicians	High on technical skills	Moderate
Crafters	Medium in technical skills, low in management skills	Very high
Doers	Emphasis on basic skills	High
Solvers	Emphasis on management skills and critical thinking	Minimal
Facilitators	Emphasis on emotional skills	Moderate
Providers	High in analytical skills	Low

Source: Adapted from RBC, 2018

Similar questions may be asked for key competences as defined in the EU Key Competences for Lifelong Learning, as well as for the European Framework for Digital Competence of Educators.[53] Figure 5 lists some example capabilities that could have impact on the key competence on languages. In general, studies on future work and skill demand suggest that education cannot easily focus on specific work-related skills in the future. Instead, education needs to create competence platforms that enable effective life-long learning. Somewhat paradoxically, such a view on "platform education" suggests that we may be moving back towards the medieval trivium[54] and quadrivium[55], with their seven liberal arts. Business executives have already for many years argued that we need educational systems that teach people grammar, logic, rhetoric, arithmetic, and geometry. Although music and astronomy have not been high on the list, perhaps this is because they are now subsumed under terms such as creativity and science.

52 RBC (2018).

53 European Commission (EC 2018a), Redecker (2017).

54 The lower division of the seven liberal arts and comprises grammar, logic, and rhetoric, see: https://en.wikipedia.org/wiki/Trivium

55 Consisted of arithmetic, geometry, music, and astronomy, see: https://en.wikipedia.org/wiki/Quadrivium

Figure 5: Skills of the languages key competence and some associated AI capabilities

Key skills

- Ability to understand spoken messages
- Initiate, sustain and conclude conversations
- Read, understand and drafts texts

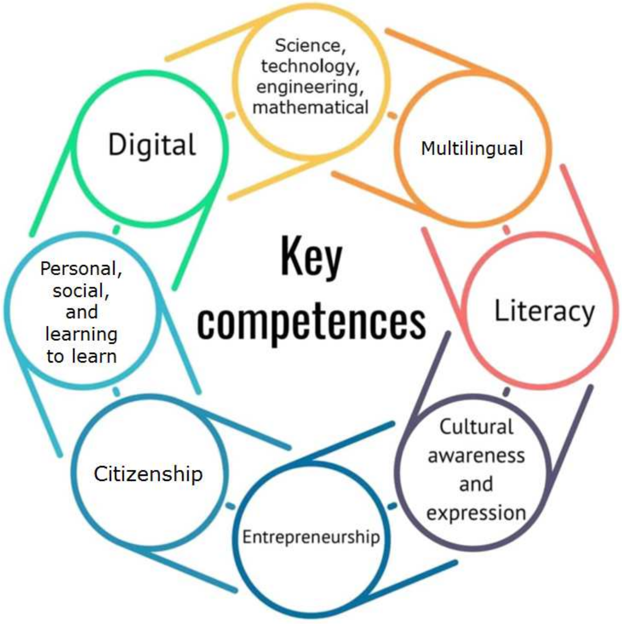

- Real time translation
- Semantic search
- Speech-to-text
- Grammar and spelling
- Rhetoric impact
- Online dictionaries

- Culture-aware machine translation
- Sentiment analysis
- Personalised messaging

Source: Author’s elaboration. Council recommendation on Key Competences for Lifelong Learning

2.4.7 Direct AI impact on advanced digital skills demand

The development of new AI and machine learning models requires very high levels of competences in several areas. This is one of the reasons why AI experts are now being paid extreme salaries. The number of neural AI experts is perhaps doubling annually, but the basic knowledge needed for state-of-the-art work in this area requires advanced levels of scientific, mathematical and technical skills that are demanding to acquire. Development of new AI methods requires good understanding of statistics, linear algebra, differential equations, as well as computer architectures and emerging chip technologies[56], programming approaches and tools. The required skill set is rather scarce, and recent estimates put the number of people with this set at some tens of thousands.[57] There are some 5,000 persons who have written academic articles or presented at AI conferences in recent years.

It may be expected that the high visibility of AI and the current demand will relatively rapidly direct talent to this area. As an example, since its launch in May 2018, about 90 000 students from over 80 countries have enrolled to the six-week Elements of AI –course organised as part of the AI Education programme of the Finnish Center of AI.[58] This introductory course has been popular among policymakers and in private and public sector organisations who struggle to make sense of developments in AI. High-level skills in AI, however, cannot be acquired quickly, and the scarcity of AI-related skills may have serious indirect implications for teaching and learning. In 2017, AI related business

[56] One key bottleneck for neural AI is its energy consumption. As a result, many chip designers are now trying to develop semiconductor chips that can be used for specific AI applications, see e.g. (Salvo 2018).

[57] Element AI has recently calculated the number of people with the required skill set at 22,000, see (Kahn 2018).

[58] https://www.elementsofai.com

mergers and acquisitions were about 21.8 billion USD worldwide, and start-ups without revenue fetched prices that amount to $5-10 million per AI expert.[59] As highly-qualified experts can now earn very high annual salaries, universities will have great difficulties in finding competent teachers for this specialty. Some practical implementation work can be done by relative novices using openly available development tools and learning materials, but the development of mission-critical applications requires quite advanced skills.[60]

One rather immediate result of this situation is that high-level AI talent and compute capability will probably be provided as a service. This would perhaps mean that there is not going to be massive needs for high-level AI competences. Due to the high wage differentials, many current students of statistics, mathematics, mathematical physics, computer and chip design, and perhaps neurophysiology may, however, reconsider their career paths and find new identities as experts in AI. Moreover, in the current informal learning environment, easy access to state-of-the-art technologies and research could also mean that high-level AI competences may emerge from unexpected places, for example, through open software and open hardware communities.

[59] Data from PitchBook, quoted in (Bass 2018).

[60] One key bottleneck for neural AI is its energy consumption. As a result, many chip designers are now trying to develop semiconductor chips that can be used for specific AI applications, see e.g. (Salvo 2018).

3 Impact on learning, teaching, and education

Since the beginning of the 1980s, and until recently, educational applications of AI have mainly focused on the knowledge-based approach.[61] The most prominent line of research has been concerned with intelligent tutoring systems, or ITS.[62] These systems use a knowledge-based architecture. A typical ITS architecture has a *domain model* that describes the area to be learned and a *student model* that describes the current state of student's knowledge and learning. An expert system or *pedagogical model* manages the introduction of learning materials to the student through an adaptive and interactive user interface.

These systems have traditionally used the knowledge-based approach, now commonly known as "gofai" (good-old-fashioned-AI). They have been successful mainly in relatively limited and unambiguous domains, such as mathematics and physics.[63] As student behaviour and learning can also be monitored in ITS environments in great detail, intelligent tutoring environments have also been an important source of data for research on learning.[64] The difficulty in developing ITS for broad learning domains has also switched the focus to the more narrow problem of using AI and machine learning to generate teacher interfaces for student and learning monitoring, and learning diagnostics. This is commonly known as learning analytics and educational data mining (EDM).[65]

3.1 Current developments

In special needs education, AI-based approaches have shown potential, for example, in the early detection of dyslexia.[66] A well-published example is the Swedish company "*Lexplore*" that has developed a system that quickly scans for students at risk and detects dyslexia by tracking reader eye movements. The system uses data-based pattern recognition, and the company is now expanding to the US and UK, offering school and school-district wide scanning.[67] AI-based systems have also been successfully developed for the diagnosis of autism spectrum disorder and attention deficit hyperactivity disorder (ADHD). In particular, child-robot interaction seems to enable new forms of diagnostics and special needs educational applications.[68]

As student testing plays an important role in many educational systems, many projects are trying to explore the use of AI for automatic test generation and assessment. Much of this work is aimed at automating summative assessment, with a promise of reducing teacher workloads. A possible unintended consequence of this work is that high-stakes testing will be increasingly displaced by frequent low-stakes formative assessment, as the effort and cost required for assessment decreases. Current AI systems are very good in combining evidence from complex and varied sources of data and using them for real-time pattern recognition. For example, student homework can relatively easily be checked and diagnosed by an AI system that has data on both individual student history and peer responses. Accumulated formative assessments could, therefore, to a large extent make high-stakes testing redundant. AI is also beginning to be used to diagnose student attention, emotion, and conversation dynamics in computer-supported learning

[61] For an early example, see Sleeman and Brown (1982).
[62] E.g., Woolf (2009).
[63] E.g. Ritter et al. (2007), Graesser et al. (2005).
[64] E.g., Porayska-Pomsta (2015).
[65] For a compact review of some relatively recent developments, see Luckin et al. (2016) and a JRC report on Learning Analytics by Ferguson et al. (2016)
[66] See, e.g., Drigas and Ioannidou (2012).
[67] Jakobsson (2017). For English version, see http://www.lexplore.com/
[68] E.g., Scassellati (2012), Boccanfuso et al. (2016).

environments, for example for course development and management, in an attempt to generate optimal groups for collaborative learning tasks, and to recognize patterns that predict student drop-out.[69] To do this effectively, large datasets are needed for training the systems. As was pointed out above, this is a major technical bottleneck. Student behavior also has to be actively monitored to provide feedback for learning. This creates technical needs to unobtrusively monitor students, for example, using video processing and remote eye tracking, with associated ethical and regulatory challenges. Ethically less problematic are systems that use less granular data to provide recommendations. For example, at UC Berkeley students can now get course recommendations using a system that relies on neural AI technologies originally developed for natural language processing and machine translation.[70]

3.1.1 "No AI without UI"

A core idea in intelligent tutoring systems is that a student interacts with adaptive interfaces that personalize learning experiences based on the student and her current level of learning. The core strength of data-based AI systems, on the other hand, is that they can process very complex data streams in real time. For next-generation ITS this means that these systems will need user interfaces (UI) that collect real-time input from learner behaviour and also historical data that can be used to model the learner. In informal terms, this can be called the principle of "no AI without UI." There will, therefore, be considerable commercial interest to push various kinds of sensor technologies and user interfaces to classrooms, as well as to gain access to data from other learner related data sources, such as social media and game platforms.

Although many ITS systems have been developed in the cognitivist tradition and based on an instructivist approach to pedagogy, also other pedagogical models have frequently been used. For example, the idea that technology can be used to support and scaffold learning and act as a competent guide and companion has been influential. Related research on social learning and knowledge building and construction has also shaped research in this area.[71] As constructivist and constructionist models have gained popularity, the emphasis has shifted from teaching to more student-centric approaches, including support for peer-to-peer social learning. It can be expected that, as conversational natural language systems such as the Google Duplex are now becoming commercially available, teachable conversational agents will be one area where educational AI start-ups try to create new business in the near future.

3.2 The impact of AI on learning

In formal education, AI can have both positive and negative impact on learning. As AI is now high on the policy agenda, it may appear that AI should be applied in as many educational settings as possible. When a new promising technology emerges, and when the limitations of technology and the challenges of applying it are often not perfectly understood, technology may seem to open radically new possibilities for solving old problems.

This is what happens at the early phases of the life-cycle of general-purpose technologies, and it leads to technology push. Visionary entrepreneurs and policymakers

[69] See, e.g., Nkambou et al. (2018), Rosé et al. (2018).
[70] E.g., Pardos et al. (2018).
[71] See, e.g., Scardamalia and Bereiter (2006), Paavola and Hakkarainen (2005), Thomas and Brown (2011).

realize the potential of new technology and see all the possibilities of how it could make a difference. In the domain of learning, this enthusiasm will be mitigated when people realize that AI will not only make existing education more efficient but that it will also change the context where learning occurs and where it becomes socially relevant. Many current learning practices address the needs of an industrial society that is currently being transformed. It is easy to automate things that merely institutionalize old habits. In a changing world, this often creates frustration as the solutions can become obsolete already before they are implemented.

In the stage of technology push, technology experts possess scarce knowledge. Because it is scarce, it often dominates and overrides other types of knowledge. In the domain of education and training, this can become a problem as technologists easily transfer their own experiences and beliefs about learning to their designs. For example, in the field of machine learning, learning is often understood as simple association between system inputs and outputs. For learning scientists, such a concept of machine learning may be an oxymoron. Using technology, it may be possible to revolutionize learning but it is also possible to automate ideas and replicate practices that have little to do with learning.

For example, the promise of MOOCs has been widely noted but we still know very little about their impact on "delivering desired learning outcomes." As it is possible for one teacher to teach very many students in online environments,[72] but difficult to know what the students learn, one of the great promises of AI is to do large-scale learning analytics in such environments. For example, it is often suggested that AI could be used to objectively assess student learning by scoring test results without teacher bias. Given enough human-labelled examples of data, neural AI and machine learning can easily learn to categorize students based on their test results. Yet, it is not clear that test results are accurate indicators of learning. To support learning, it may be more important to measure individual development than average performance in standardized tests.[73] Neural AI, however, strongly prefers large datasets and standardized testing. Current neural AI systems are a natural fit with learning models that view learning as transfer of knowledge to student's mind. If learning is understood as the development of skills and competences, AI my need to be incorporated in learning processes in different ways.

For example, IBM's Watson Classroom promises cognitive solutions that help educators gain insights into the learning styles, preferences, and aptitudes of each student, "bringing personalized learning to a whole new level."[74] It is, however, not obvious that such objectives would be beneficial or relevant for learning. As Vygotsky pointed out long time ago, the development of many cognitive capabilities that define advanced forms of thinking are based on their social relevance and have little immediate relevance for an individual learner. For example, mediated communication through written text is unnatural for a child who is perfectly able to use speech from an early age.[75] Without a complex system of social interests and practices, advanced conceptual systems such as those used in mathematics would make little sense for an individual learner. AI may thus provide exciting new opportunities for adapting learning content based on student's individual characteristics and learning style, even when large bodies of empirical research show that the concept of learning style is perhaps best characterized as an urban myth.[76] In short, computer programs scale up very well, and AI can easily scale up bad pedagogical ideas.

[72] See e.g., Tuomi (2013).
[73] See, e.g. Mislevy (2018), Gane et al. (2018).
[74] https://www.ibm.com/watson/education
[75] Vygotsky (1986).
[76] E.g., Riener and Willingham (2010).

3.2.1 Impact on cognitive development

On a more fundamental level, we can ask what is the impact of AI on the development of human cognition and human brain[77]. More broadly, this is a question about co-evolution of technology and human mind. Friedrich Engels' influential unfinished essay "The Part Played by Labour in the Transition from Ape to Man" emphasized the specialization of knowledge, division of productive labour, and the role of technology, arguing that the development of human brain and society were intrinsically connected.[78] Labour, states Engels in the beginning of his essay, "is the prime basic condition for all human existence, and this to such an extent that, in a sense, we have to say that labour created man himself."

The idea that new ways to organize production lead to new forms of "consciousness" became one of the driving forces in the revolutionary movements towards the end of the 19th century. The original idea, however, was essentially a Darwinian explanation about how human brain has evolved. This idea of linkages between cognitive development and social division of knowledge and practical labour is also today influential in the post-Vygotskian learning theory, and Vygotsky himself was highly interested in the role of material artefacts and tools in thinking.[79]

Recent research on neuroplasticity takes this idea one step further, showing that tools and technology do not only shape the way we think but they can also shape the brain itself. One could, therefore, ask how the use of AI technologies in learning changes the structure of human brains.[80] In particular, recent research shows that there are critical phases in the development of the brain. Cognitive technologies may, therefore, have quite fundamental consequences if used during such critical periods. At present, we don't know whether this is the case.[81]

In general, AI can be used in three essentially different ways that may have different implications for the development of human cognitive capabilities both in children and adults. First, AI can support existing capabilities. When competences are understood as combinations of domain specific expertise and behavioural repertoires,[82] AI can reduce the need for human knowledge, experience, and skill, and emphasize the importance of behavioural repertoires. As a result, humans do not necessarily need to learn domain specific knowledge that earlier was required for competent behaviour. In particular, as

---

77 See for instance: Gómez, E., Castillo, C., Charisi, V., Dahl, V., Deco, G., Delipetrev, et al. (2018). Assessing the impact of machine intelligence on human behaviour: an interdisciplinary endeavour. arXiv preprint arXiv:1806.03192.

78 Engels (1966, chap. 6). A similar historical approach is more recently adopted by Morrison and Miller (2017), who argue that human learning is a species-specific capability that is in many ways built in to human biology, culture and social structures.

79 E.g. Bruner (1986), Engeström (1987).

80 There are now large bodies of empirical research on structural change in the human brain. Often quoted studies in this area are by Maguire et al. (2000; Woollett and Maguire 2011). They measured the structural changes in the hippocampus of London taxi-drivers, showing changes in this area associated with spatial navigation.

81 For example, it has been shown that musical training in infancy leads to an expanded auditory cortical representation, but only if practicing begins before the age of 9 (Pantev et al. 1998). Whereas the classical studies focused on the period where normal development occurs, abnormal input can have a permanent deleterious effect also after the period of normal development is over. Lewis and Maurer (2005) called these the "sensitive periods for damage," and showed that visual deprivation up to 10 years of age leads to a permanent deficit in visual acuity.

82 This is suggested, for example, by Hoekstra and van Slujis (2003). In the context of the three-level model presented here, such a model of competences appears too narrow, and would need to be augmented by both cultural and technical elements that make expressions of competence possible and relevant.

domain-specific knowledge becomes less important for competence, transversal and domain-independent generic competences may become relatively more important.

Second, AI can speed-up cognitive development and create cognitive capabilities that would not be possible without technology. The mechanization or human work has made possible things that would be impossible without technology; similarly, the mechanization of cognitive work makes possible new activities that have not been possible before. This, of course, is something that already has happened. It would be entirely impossible to design a modern microprocessor or a neural chip without computer-aided design tools that use extensive bodies of design knowledge.

Third, AI may reduce the importance of some human cognitive capabilities, or make them obsolete. For example, as AI can convert speech to text and vice versa, dyslexia may become socially less important than it has been in the past. However, although in cases such as dyslexia and dyscalculia AI may have clear benefits for individuals, the overall impact is not easy to predict. For example, computers may support people in adding and multiplying numbers; if they became reliant on computational machines, it may, however, become more difficult to develop more advanced mathematical skills that require mental arithmetic and number skills. From a pedagogic point of view, it may sometimes be more beneficial to use AI to help people to develop competences that allow them to overcome difficulties in reading and counting, instead of using AI to make redundant skills that underpin important cognitive capabilities.

3.3 The impact of AI on teaching

If we think how AI can most effectively be used in the current educational context, we easily automate things that used to be important in the past. It is therefore important to understand the impact of AI in the context of future learning and education, instead of in current systems of education and forms of learning. The analysis of the impact of AI on teaching will, therefore, be inherently linked to foresight-oriented work on the future of learning.

Yet, there are some educational tasks where AI can have a clear impact. One such task is assessment in its various forms. In the conventional intelligent tutoring systems a central component is a student model that maintains information about the current state of the learner and which, based on the student model, tries to infer possible bottlenecks in student's way of understanding a domain that she or he is learning.

3.3.1 AI-generated student models and new pedagogical opportunities

In principle, neural AI is well suited for diagnostic tasks. Traditional knowledge-based intelligent tutoring systems have struggled with the challenge of creating student models partly because there is no obvious way to create representations of student models in complex domains and in realistic context of learning. Neural AI, however, may generate student models if sufficient amounts of data are available. As discussed above, words in natural languages can often be represented using a 300-dimensional space where millions of words are located based on billions of examples (see 2.3). Machine learning can generate such complex representations in ways that work in practice, despite all their conceptual and technical inadequacies. Given enough data, machine learning can probably create student models that are good enough to be of practical value.

Neural AI can also learn patterns of interaction and associate these with pedagogically relevant clusters so that a teacher can have a better understanding of the ways in which students think and where they could be effectively guided. AI systems can also provide such diagnostic data also to the students so that they can reflect on their metacognitive approaches and possible areas in need of development. Neural AI will therefore have important potential in learning diagnostics, analytics and educational data mining.

The rapid advances in natural language processing and AI-based human-machine interfaces will generate new pedagogical possibilities, too. For example, as conversational robots and learning companions are becoming more and more available, learning by teaching robots shows some potential[83]. Affective computing and emotion AI will be important components of such systems. Additionally, real-time machine translation opens up new possibilities in language learning, and AI systems can be used, for example to interpret texts written by students thus helping them to write texts that communicate better what the student intended to communicate.

3.3.2 The need for future-oriented vision regarding AI

It is possible to imagine many exciting possibilities for AI in teaching. Without clear pedagogic principles, it is, however, probable that AI vendors will provide products and services that address key decision-makers' perceived immediate problems, instead of more fundamental social and economic challenges. For an AI start-up in the educational sector, it is difficult to offer products and services that require change in current educational practices.

Therefore, without clear visions and policies that put emerging technical possibilities in the broader context of the transformation of education and the future of learning, educational AI will probably mainly be provided as solutions to existing problems. Instead of renewing the system and orienting it towards the needs of a post-industrial economy and knowledge society, AI may therefore mechanize and reinvent outdated teaching practices and make them increasingly difficult to change. It may, therefore, be necessary to develop appropriate visions and policies by simultaneously creating future-oriented models for education and teaching. Creating concrete experimentations in an authentic context with teachers and experts in education is important. As AI is now very high on the policy agenda, it is too easy to generate high-level visions of the future that claim that AI is the next technical revolution. AI is now frequently called "the new electricity." It is therefore important that teachers, who often struggle with concrete demands of everyday teaching practice and new initiatives, will not be electrocuted by this new technology.

3.4 Re-thinking the role of education in society

On a more systemic level, AI will have a profound impact on education systems. This is not because of any specific characteristics of AI; Instead, AI is one expression of an ongoing broader transformation that results from digitalization, global real-time networking of communication and production, and automation of productive processes.

[83] E.g. see projects such as http://de-enigma.eu/ and https://www.dream2020.eu/

This has variously been called the information society, the knowledge economy and the algorithmic revolution.[84] One of the reasons why AI has emerged as major policy topic in recent years is that it is becoming clear that AI will have a radical impact on the world of work. As the current educational institutions have to a large extent emerged as answers to problems of the industrial age, many of these answers are now becoming outdated.

It is possible that those economists are right who argue that automation and AI will not increase unemployment in the future. In the 20th century context, this would be good, as unemployment was a major economic challenge in industrialized societies. Such arguments are supported by economic theories that start from the assumption that economies tend toward equilibrium. They are also supported by common sense that says that of course people have to work. Adopting such views, one may say that of course there will be work in the future although we do not yet know how it will look like and what the jobs will be. It is also possible that work in the future will no longer be what it used to be. In the history of educational thinking, there has been a constant battle between views that see education from an instrumental point of view—as a way of preparing future workers for future jobs—and a more developmental view that sees education as a way of realizing human potential. Whether there will be jobs in the future or not, AI seems to push education towards these more developmental models of education. Assuming that AI will transform the labour market, a potentially useful way of imagining the future of education and educational systems is to start from the latter possibility. If we imagine education in a world where work is not a central factor in life or where jobs, as we knew them, do not exist, what would be the role of education? How could we organize it? What would be its aims and what needs would it address?

[84] The concept of algorithmic revolution is perhaps the least known of these. It has been discussed by Zysman (2006).

4 Policy challenges

The current excitement about AI easily leads to technology push, where AI is viewed as a solution to a wide variety of problems in education and learning. It is probably fair to say that the potential and challenges of AI in education are still not adequately understood. AI can be understood as a general-purpose technology, and it can be applied in many different ways. Although the characteristics of technology itself may push development towards specific directions, it is always possible to use technology in many ways and for many different purposes, also in education. For policy development, it is therefore probably more important to understand why and for what we use technology than how it is used. The future promises of technology, in this view, have to be justified by making explicit the motivation of using the technology, as well as the key assumptions that underpin the stated motivation. This lifts technology to a level of policy, and we have to ask what are the objectives and goals of using it. Only if we have such a birds-eye view on technical development, we can say where we want to go and how technology can help us on the way. When the assumptions and motivations are made explicit, they can also be critically assessed.

A continuous dialogue on the appropriate and responsible uses of AI in education is therefore needed. As technology and its uses change, important contributions to this dialogue may emerge from "outsiders" who do not represent current stakeholder interests. Enabling and funding independent research on, for example, the politics, ethics, social implications, and economy of AI may be a practical way to create useful inputs to this dialogue.

In the domain of educational policy, it is important for educators and policymakers to understand AI in the broader context of the future of learning. To a large extent, the debate about AI is now about the ongoing informationalization, digitalization, and computer-mediated globalization. The current estimates of the impact of AI and other digital technologies on the labour market highlight the point that the demand for skills and competences is changing fast, and the educational system needs to adapt, in particular when education aims to create skills for work. AI enables the automation of many productive tasks that in the past have been done by humans. **As AI will be used to automate productive processes, we may need to reinvent current educational institutions.** It is, for example, possible that formal education will play a diminishing role in creating job-related competences. This could mean that the future role of education will increasingly be in supporting human development.

For example, the current AI systems make almost continuous assessment of student progress possible. Instead of high-stakes testing that functions as a social filter, AI supported assessment can be used to help learners to develop their skills and competences and keep students on effective learning paths. With such ongoing assessment, **high-stakes testing may become redundant, and broader evidence may be used for assessing skills and competences**. This may be important in particular for assessing transversal key competences that are now relatively difficult to assess. As AI and other information technologies facilitate informal learning, it also becomes important to ask what the division of labour between formal and informal learning will be in the future.

In general, the balance may thus shift from the instrumental role of education towards its more developmental role. Perhaps more importantly, it is possible that the industrial age link between work and education is changing. Current institutions of education to a large extent address the needs of an industrial world. As knowledge and data are now created, used, and learned in ways that have not been possible before, it is important that AI is not understood only as a solution to problems in the current educational systems.

In general, the profound changes in the society and economy that AI and related technologies are now making possible will create a world where many social institutions will change, and people have to adapt. When a similar broad change occurred almost two centuries ago, the social and human costs were high. Although we now with hindsight often neglect the negative consequences of technical development and emphasize its positive consequences, it is important to realize that general-purpose technologies can have fundamental transformative impact on social life and human development. The rather poetic declaration in 1848 that "all that is solid melts into air," was not just a vision but it was based on careful empirical observation of the everyday consequences of industrialization.[85] **A general policy challenge, thus, is to increase among educators and policymakers awareness of AI technologies and their potential impact.** One way of doing this is to participate in processes that generate images of future, develop concepts that can be used to describe them, and design scenarios and experiments where such imagined futures can be tested. A rather simple proposal for policy development, thus, is to launch explicitly future-oriented processes that generate understanding of the possibilities of the present.

AI provides new means for research on learning, but it is also important to rethink the capabilities of AI systems using existing knowledge about learning.[86] In particular, almost all currently developed AI systems rely on associative and behaviouristic models of learning. The long history of neural AI contains many attempts to go beyond these simple models of learning. **Learning sciences could have much to offer to research on AI, and such mutual interaction would enable better understanding about how to use AI for learning and in educational settings, as well as in other domains of application.**

Data that is needed for machine learning is often highly personal. If it is used for assessing student performance, data security can become a key bottleneck in using AI, learning analytics, and educational data mining. As neural AI systems do not understand the data they process, it is also easy to forge data that fools the decision process.[87] AI security is an important topic, but it is also challenging as neural AI systems typically use complex internal representations of data that are difficult or impossible to interpret. Because of this there is now considerable interest in creating "explainable AI." The current systems, however, lack all the essential reflective and metacognitive capabilities that would be needed to explain what they do or don't do.[88] To rephrase Descartes, it is, therefore, as futile to ask a clock on the wall why it just struck seven or eight as it is to ask a deep learning AI system why it gave a specific grade to a student. Clocks are not built to explain their ticking, and AI systems, as we know them, have no explanatory capabilities. At best they can support humans in explaining what happened and why. **As there may be fundamental theoretical and practical limits in designing AI systems that can explain their behaviour and decisions, it is important to keep humans in the decision-making loop.**

---

[85] The quote is from the Manifesto of the Communist Party by Marx & Engels, 1848.

[86] There have been very few attempts to analyse AI from the point of view of learning theories. The learning capabilities of convolutional neural networks have been compared with Vygotsky's model of conceptual development in Tuomi (2018).

[87] Pattern matching systems can be very fragile in their decision-making capabilities. It is possible, for example, to fool image recognition programs by modifying image pixels (e.g., Yuan et al. 2017; Kurakin, Goodfellow, and Bengio 2016).

[88] Luckin (2018).

As several recent reports have emphasized, ethical considerations become highly relevant when AI is applied in the society or in educational settings.[89] From a policy perspective, **the ethics of AI is a generic challenge, but it has specific relevance for educational policies.**

From the regulatory point of view, ethical considerations provide the fundamental basis from which new regulations and laws are created and justified. From a developmental point of view, ethics and value judgements underpin fundamental concepts such as agency, responsibility, identity, freedoms, and human capabilities. In supervised AI learning models, the possible choice outcomes need to be provided to the system before it starts to learn. This means that the world becomes described in closed terms, based on predefined interests and categories. Furthermore, the categories are based on data that are collected in the past. Neural AI categorizes people in clusters where data from other people, considered similar by the system, is used to predict individual characteristics and behaviour.

From political and ethical points of view, this is highly problematic. Human agency means that we can make choices about future acts, and thus become responsible for them. When AI systems predict our acts using historical data averaged over a large number of other persons, AI systems cannot understand people who make true choices or who break out from historical patterns of behaviour. **AI can therefore also limit the domain where humans can express their agency**.

As has been emphasized above, the recent successes in AI have to a large extent been based on the availability of vast amounts of data. AI-based products and services can be created in the educational sector only if appropriate data is available. At present, some of the existing datasets can be considered as natural monopolies, and they are often controlled by few large corporations. **An important policy challenge is how such large datasets that are needed for the development and use of AI-based systems could be made more widely available.** One potential solution is to build on the current General Data Protection Regulation which requires that data subjects can have a copy of their personal data from data controllers in a commonly used electronic form. Technically this would make it possible for users to access their personal data, anonymize it locally, and submit it in an appropriate format to platforms that are used for AI learning and educational purposes. Such functionality might be relatively easily embedded, for example in commonly used web browsers, if platforms for data aggregation would be available. One possibility could be to pilot such aggregation platforms on a suitable scale and, if successful, provided at the EU level.

[89] See, e.g., Demiaux and Si Abddallah (2018). The U.K. House of Lords special committee on AI suggests that the ethical use of AI could become the differentiating factor for AI research in the U.K. (House of Lords 2018). Also commercial actors have highlighted the importance of ethical considerations (Microsoft 2018). The European group of ethics in science and technology has well emphasized the importance of agency for understanding ethical and political implications of AI (EGE 2018). Also the European Commission's High-Level Expert Group on Artificial Intelligence (AI HLEG) is currently developing AI ethics guidelines.

References

Acemoglu, Daron, and Pascual Restrepo. 2016. "The Race Between Machine and Man: Implications of Technology for Growth, Factor Shares and Employment." Working Paper 22252. National Bureau of Economic Research. https://doi.org/10.3386/w22252.

Agrawal, Ajay, Joshua Gans, and Avi Goldfarb, eds. 2018. "Introduction to : 'Economics of Artificial Intelligence.'" In *Economics of Artificial Intelligence*. Toronto: nber.org. http://www.nber.org/chapters/c14005.pdf.

Anderson, James, A., and Edward Rosenfeld, eds. 1988. *Neurocomputing: Foundations for Research*. Cambridge, MA: The MIT Press.

Autor, David H. 2015. "Why Are There Still So Many Jobs? The History and Future of Workplace Automation." *Journal of Economic Perspectives* 29 (3): 3–30. https://doi.org/10.1257/jep.29.3.3.

Autor, David H., Frank Levy, and Richard J. Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration." *The Quarterly Journal of Economics* 118 (4): 1279–1333. https://doi.org/10.1162/003355303322552801.

Bass, A.S. 2018. "Non-Tech Businesses Are Beginning to Use Artificial Intelligence." *Financial Times*, March 31, 2018.

Boccanfuso, Laura, Erin Barney, Claire Foster, Yeojin Amy Ahn, Katarzyna Chawarska, Brian Scassellati, and Frederick Shic. 2016. "Emotional Robot to Examine Differences in Play Patterns and Affective Response of Children with and Without ASD." In *The Eleventh ACM/IEEE International Conference on Human Robot Interaction*, 19–26. HRI '16. Piscataway, NJ, USA: IEEE Press. http://dl.acm.org/citation.cfm?id=2906831.2906837.

Boden, Margaret A. 2016. *AI: Its Nature and Future*. Oxford, New York: Oxford University Press.

Bostrom, Nick. 2014. *Superintelligence: Paths, Dangers, Strategies*. 1 edition. Oxford: Oxford University Press.

Brown, J.S., A. Collins, and P. Duguid. 1989. "Situated Cognition and the Culture of Learning." *Educational Researcher* 18 (1): 32–42.

Bruner, J. 1986. *Actual Minds, Possible Worlds*. Cambridge, MA: Harvard University Press.

Brynjolfsson, Erik, and Andrew McAfee. 2012. *Race Against the Machine: How the Digital Revolution Is Accelerating Innovation, Driving Productivity, and Irreversibly Transforming Employment and the Economy*. Brynjolfsson and McAfee.

Cole, M. 1986. *Culture in Mind*. Cambridge, MA: Harvard University Press.

Cole, M., and J.V. Wertsch. 1996. "Beyond the Individual-Social Antinomy in Discussions of Piaget and Vygotsky." *Human Development* 39 (5): 250–56.

Demiaux, Victor, and Yacine Si Abdallah. "How Can Humans Keep the Upper Hand? The Ethical Matters Raised by Algorithms and Artificial Intelligence." Report on the public debate led by the French Data Protection Authority (CNIL) as part of the ethical discussion assignment set by the Digital Republic Bill. Paris: CNIL, December 2017.

Dewey, J. 1991. *How We Think*. Buffalo, NY: Prometheus Books.

Dreyfus, H.L. 1979. *What Computers Can't Do: A Critique of Artificial Intelligence*. New York: Harper & Row.

Drigas, Athanasios, and Rodi-Eleni Ioannidou. 2012. "Artificial Intelligence in Special Education: A Decade Review." *International Journal of Engineering Education* 28 (6): 1366–72.

EC. 2018a. "Proposal for a Council Recommendation on Key Competences for Lifelong Learning." COM(2018) 24 final. Brussels: European Commission.

https://ec.europa.eu/education/sites/education/files/recommendation-key-competences-lifelong-learning.pdf.
———. 2018b. "Artificial Intelligence for Europe." COM(2018) 237 Final. Brussels: European Commission. https://ec.europa.eu/digital-single-market/en/news/communication-artificial-intelligence-europe.
EGE. 2018. "Statement on Artificial Intelligence, Robotics and 'Autonomous' Systems. European Group on Ethics in Science and New Technologies." Brussels: European Commission. https://ec.europa.eu/research/ege/pdf/ege_ai_statement_2018.pdf.
Engels, Friedrich. 1966. *Dialectics of Nature*. Moscow,.
Engeström, Y. 1987. *Learning by Expanding: An Activity Theoretical Approach to Developmental Work Research*. Helsinki: Orienta Konsultit.
Engeström, Y., J. Virkkunen, M. Helle, J. Pihlaja, and R. Poikela. 1996. "The Change Laboratory as a Tool for Transforming Work." *Lifelong Learning in Europe* 1 (2): 10–17.
EPSC. 2018. "The Age of Artificial Intelligence: Towards a European Strategy for Human-Centric Machines." 29. EPSC Strategic Notes. European Political Strategy Centre. https://ec.europa.eu/epsc/sites/epsc/files/epsc_strategicnote_ai.pdf.
Ferguson, R, Brasher, A., Clow, D., Cooper, A.,Hillaire, G., Mittelmeier, J., Rientes, B., Ullmann, T. , Vuorikari, R. 2016. "Research Evidence on the Use of Learning Analytics: Implications for Education Policy." JRC Science for Policy Report. JRC. http://europa.eu/!cB93Gb
Freire, P. 1972. *Pedagogy of the Oppressed*. Harmondsworth: Penguin.
Frey, Carl Benedikt, and Michael A. Osborne. 2013. "The Future of Employment." Working Paper. Oxford: Oxford Martin Programme on Technology and Employment. https://www.oxfordmartin.ox.ac.uk/downloads/academic/future-of-employment.pdf.
———. 2017. "The Future of Employment: How Susceptible Are Jobs to Computerisation?" *Technological Forecasting and Social Change* 114 (January): 254–80. https://doi.org/10.1016/j.techfore.2016.08.019.
Gane, Brian D., Sania Z. Zaidi, and James W. Pellegrino. 2018. "Measuring What Matters: Using Technology to Assess Multidimensional Learning." *European Journal of Education* 53 (2): 176–87. https://doi.org/10.1111/ejed.12269.
Gardner, H. 1987. *The Mind's New Science: A History of Cognitive Revolution*. New York: Basic Books.
Gómez, E., Castillo, C., Charisi, V., Dahl, V., Deco, G., Delipetrev, et al. 2018. *Assessing the impact of machine intelligence on human behaviour: an interdisciplinary endeavour*. arXiv preprint arXiv:1806.03192
Graesser, A. C., P. Chipman, B. C. Haynes, and A. Olney. 2005. "AutoTutor: An Intelligent Tutoring System with Mixed-Initiative Dialogue." *IEEE Transactions on Education* 48 (4): 612–18. https://doi.org/10.1109/TE.2005.856149.
Harré, Rom, David Clarke, and Nicola De Carlo. 1985. *Motives and Mechanisms: An Introduction to the Psychology of Action*. London: Methuen & Co. Ltd.
Hayek, F.A. 1945. "The Use of Knowledge in Society." *American Economic Review* 35 (4): 519–30.
———. 1952. *The Sensory Order: An Inquiry into the Foundations of Theoretical Psychology*. Chicago, IL: Chicago University Press.
Heinämaa, S., and I. Tuomi. 1989. *Ajatuksia Synnyttävät Koneet: Tekoälyn Unia Ja Painajaisia (Thought Provoking Machines: Dreams and Nightmares of Artificial Intelligence; in Finnish)*. Porvoo: Werner Söderström Osakeyhtiö.
Hoekstra, H.A., and Van Slujis, E. 2003. *Managing Competences: Implementing Human Resource Management*. Assen: Koninklijke Van Gorcum.
House of Lords. 2018. "AI in the UK: Ready, Willing and Able?" HL Paper 100. London: House of Lords, Select Committee on Artificial Intelligence. https://publications.parliament.uk/pa/ld201719/ldselect/ldai/100/100.pdf.
Hutchins, E. 1995. *Cognition in the Wild*. Cambridge, MA: MIT Press.

Jakobsson, Josefin. 2017. “Nya miljoner ska ta deras dyslexi-startup till USA.” *Di Digital*, March 29, 2017. https://digital.di.se/artikel/nya-miljoner-ska-ta-deras-dyslexi-startup-till-usa.

Kahn, Jeremy. 2018. “Just How Shallow Is the Artificial Intelligence Talent Pool?” *Bloomberg.Com*, February 7, 2018. https://www.bloomberg.com/news/articles/2018-02-07/just-how-shallow-is-the-artificial-intelligence-talent-pool.

Kurakin, Alexey, Ian Goodfellow, and Samy Bengio. 2016. “Adversarial Examples in the Physical World.” *ArXiv:1607.02533 [Cs, Stat]*, July. http://arxiv.org/abs/1607.02533.

Kurzweil, R. 1999. *The Age of Spiritual Machines: When Computers Exceed Human Intelligence*. New York: Viking.

Leont’ev, A.N. 1978. *Activity, Consciousness, and Personality*. Englewood Cliffs, NJ: Prentice-Hall.

Lewin, Kurt. 1946. “Action Research and Minority Problems.” *Journal of Social Issues* 2 (4): 34–46. https://doi.org/10.1111/j.1540-4560.1946.tb02295.x.

Lewis, Terri L, and Daphne Maurer. 2005. “Multiple Sensitive Periods in Human Visual Development: Evidence from Visually Deprived Children.” *Developmental Psychobiology* 46 (3): 163–83. https://doi.org/10.1002/dev.20055.

Louie, A. H. 2007. “A Living System Must Have Noncomputable Models.” *Artificial Life* 13 (3): 293–97.

———. 2009. *More than Life Itself: A Synthetic Continuation in Relational Biology*. Frankfurt: Ontos Verlag.

Luckin, Rosemary. 2018. *Machine Learning and Human Intelligence: The Future of Education for the 21st Century*. London: UCL Institute of Education Press.

Luckin, Rosemary, Griffiths, M., and Forcier, L.B. 2016. “Intelligence Unleashed. An Argument for AI in Education.” London: Pearson.

Luria, A.R., and L. Vygotsky. 1992. *Ape, Primitive Man, and Child: Essays in the History of Behavior*. Hemel Hempstead: Harvester Wheatsheaf.

Mace, W.M. 1977. “James J. Gibson’s Strategy for Perceiving: Ask No What’s inside Your Head, but What Your Head’s inside Of.” In *Perceiving, Acting and Knowing: Toward Ecological Psychology*, edited by R. Shaw. New York: John Wiley & Sons.

Maguire, Eleanor A., David G. Gadian, Ingrid S. Johnsrude, Catriona D. Good, John Ashburner, Richard S. J. Frackowiak, and Christopher D. Frith. 2000. “Navigation-Related Structural Change in the Hippocampi of Taxi Drivers.” *Proceedings of the National Academy of Sciences* 97 (8): 4398–4403. https://doi.org/10.1073/pnas.070039597.

McCorduck, P. 1979. *Machines Who Think: A Personal Inquiry into the History and Prospects of Artificial Intelligence*. San Francisco, CA: W.H. Freeman and Company.

McCulloch, W.S., and W.H. Pitts. 1943. “A Logical Calculus of the Ideas Immanent in Nervous Activity.” *Bulletin of Mathematical Biophysics* 5: 115–33.

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.

Microsoft. 2018. “The Future Computed: Artificial Intelligence and Its Role in Society.” Redmond, WA: Microsoft Corporation.

Mislevy, Robert J. 2018. *Sociocognitive Foundations of Educational Measurement*. New York: Routledge.

Morrison, Donald M., and Kenneth B. Miller. 2017. “Teaching and Learning in the Pleistocene: A Biocultural Account of Human Pedagogy and Its Implications for AIED.” *International Journal of Artificial Intelligence in Education*, September, 1–31. https://doi.org/10.1007/s40593-017-0153-0.

Nedelkoska, L., and G. Quintini. 2018. “Automation, Skills Use and Training.” 202. OECD Social, Employment and Migration Working Papers. Paris: OECD.

Neumann, John von. 1951. "The General and Logical Theory of Automata." In *Cerebral Mechanisms in Behavior; the Hixon Symposium*, 1–41. Oxford: Wiley.
Ng, Andrew. 2016. "Andrew Ng: What AI Can and Can't Do." Harvard Business Review. November 9, 2016. https://hbr.org/2016/11/what-artificial-intelligence-can-and-cant-do-right-now.
Nilsson, Nils J. 2009. *The Quest for Artificial Intelligence: A History of Ideas and Achievement*. Cambridge: Cambridge University Press. https://doi.org/10.1017/CBO9780511819346.
Nkambou, Roger, Roger Azevedo, and Julita Vassileva, eds. 2018. *Intelligent Tutoring Systems: 14th International Conference, ITS 2018, Montreal, QC, Canada, June 11–15, 2018, Proceedings*. Programming and Software Engineering. Springer International Publishing.
Norman, D.A. 1993. "Cognition in the Head and in the World: An Introduction to the Special Issue on Situated Action." *Cognitive Science* 17: 1–6.
Paavola, Sami, and Kai Hakkarainen. 2005. "The Knowledge Creation Metaphor – An Emergent Epistemological Approach to Learning." *Science & Education* 14 (6): 535–57. https://doi.org/10.1007/s11191-004-5157-0.
Pantev, C, R Oostenveld, A Engelien, B Ross, L E Roberts, and M Hoke. 1998. "Increased Auditory Cortical Representation in Musicians." *Nature* 392 (6678): 811–14. https://doi.org/10.1038/33918.
Papert, Seymour. 1980. *Mindstorms: Children, Computers, and Powerful Ideas*. New York, NY, USA: Basic Books, Inc.
Papert, Seymour, and Idit Harel. 1991. *Constructionism*. Ablex.
Pardos, Zachary A., Zihao Fan, and Weijie Jiang. 2018. "Connectionist Recommendation in the Wild." *ArXiv:1803.09535 [Cs]*, March. http://arxiv.org/abs/1803.09535.
Pennington, Jeffrey, Richard Socher, and Christopher D. Manning. 2014. "GloVe: Global Vectors for Word Representation." In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1532–43. Doha: Association for Computational Linguistics.
Polanyi, M. 1967. *The Tacit Dimension*. New York: Anchor.
Porayska-Pomsta, Kaska. 2015. "AI in Education as a Methodology for Enabling Educational Evidence-Based Practice." In *Seventeenth International Conference on Artificial Intelligence in Education (AIED 2015)*, 52–61. Madrid.
Rajpurkar, Pranav, Jeremy Irvin, Kaylie Zhu, Brandon Yang, Hershel Mehta, Tony Duan, Daisy Ding, et al. 2017. "CheXNet: Radiologist-Level Pneumonia Detection on Chest X-Rays with Deep Learning." *ArXiv:1711.05225 [Cs, Stat]*, November. http://arxiv.org/abs/1711.05225.
Rashevsky, Nicolas. 1954. "Topology and Life: In Search of General Mathematical Principles in Biology and Sociology." *Bulletin of Mathematical Biophysics* 16: 317–48.
———. 1960. *Mathematical Biophysics : Physico-Mathematical Foundations of Biology*. 3rd rev. ed. Dover.
———. 1972. *Organismic Sets: Some Reflections on the Nature of Life and Society*. Holland, Michigan: Mathematical Biology, Inc.
RBC. 2018. "Humans Wanted: How Canadian Youth Can Thrive in the Age of Disruption." Royal Bank of Canada. http://www.rbc.com/humanswanted.
Redecker, Christine. 2017. "European Framework for the Digital Competence of Educators." EUR 28775 EN. JRC Science for Policy Report. Luxembourg: Publications Office of the European Union.
Riener, Cedar, and Daniel Willingham. 2010. "The Myth of Learning Styles." *Change: The Magazine of Higher Learning* 42 (5): 32–35. https://doi.org/10.1080/00091383.2010.503139.
Ritter, Steven, John R. Anderson, Kenneth R. Koedinger, and Albert Corbett. 2007. "Cognitive Tutor: Applied Research in Mathematics Education." *Psychonomic Bulletin & Review* 14 (2): 249–55. https://doi.org/10.3758/BF03194060.
Rosé, Carolyn Penstein, Roberto Martínez-Maldonado, Ulrich Hoppe, Rose Luckin, Manolis Mavrikis, Kaska Porayska-Pomsta, Bruce McLaren, and Benedict du Boulay, eds.

2018. *Artificial Intelligence in Education: 19th International Conference, AIED 2018, London, UK, June 27–30, 2018, Proceedings, Part I*. Lecture Notes in Artificial Intelligence. Springer International Publishing.

Rosen, Robert. 1958. "A Relational Theory of Biological Systems." *Bulletin of Mathematical Biophysics* 20: 245–60.

———. 1985. *Anticipatory Systems: Philosophical, Mathematical and Methodological Foundations*. Oxford: Pergamon Press.

———. 1991. *Life Itself: A Comprehensible Inquiry into the Nature, Origin and Fabrication of Life*. New York: Columbia University Press.

———. 1998. "Causal Structures in Brains and Machines." *International Journal on General Systems* 12: 107–26.

Rosenblatt, Frank. 1958. "The Perceptron: A Probabilistic Model for Information Storage and Organization in the Brain." *Psychological Review* 65: 386–408.

Rumelhart, D.E., and J.L. McClelland. 1986. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Vol. 1: Foundations; Vol 2: Psychological and Biological Models*. Cambridge, MA: The MIT Press.

Salomon, G. 1993. *Distributed Cognitions: Psychological and Educational Considerations*. Cambridge: Cambridge University Press.

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.

Scardamalia, M, and C Bereiter. 2006. "Knowledge Building: Theory, Pedagogy, and Technology." In *Cambridge Handbook of the Learning Sciences*, 97–118. New York: Cambridge University Press.

Scassellati, Brian, Henny Admoni, and Maja Matarić. 2012. "Robots for Use in Autism Research." *Annual Review of Biomedical Engineering* 14 (1): 275–94. https://doi.org/10.1146/annurev-bioeng-071811-150036.

Silver, David, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, et al. 2017. "Mastering Chess and Shogi by Self-Play with a General Reinforcement Learning Algorithm." *ArXiv:1712.01815 [Cs]*, December. http://arxiv.org/abs/1712.01815.

Sleeman, D., and J.S. Brown. 1982. *Intelligent Tutoring Systems*. New York: Academic Press.

Somers, James. 2017. "Is AI Riding a One-Trick Pony?" MIT Technology Review. September 29, 2017. https://www.technologyreview.com/s/608911/is-ai-riding-a-one-trick-pony/.

Steering Group of the Artificial Intelligence Progamme. 2017. "Finland's age of artificial intelligence: Turning Finland into a leading country in the application of artificial intelligence." MEAE reports 41/2017. Publications of Ministry of Economic Affairs and Employment. Helsinki. http://urn.fi/URN:ISBN:978- 952- 327- 24 8- 4.

Suchman, L. 1987. *Plans and Situated Actions: The Problem of Human-Machine Communication.* New York: Cambridge University Press.

Thomas, Douglas, and John Seely Brown. 2011. *A New Culture of Learning: Cultivating the Imagination for a World of Constant Change*. CreateSpace Independent Publishing Platform.

Tuomi, Ilkka. 1988. "Neural Networks as Dynamical Systems: Some Theoretical Reasons for Non-Algorithmic Information Processing." Proceedings of the Finnish Artificial Intelligence Symposium, STeP-88, Vol 2, 593–601.

———. 2002a. *Networks of Innovation*. Oxford: Oxford University Press.

———. 2002b. "The Lives and Death of Moore's Law." *First Monday* 7 (11). http://www.firstmonday.org/issues/issue7_11/tuomi/.

———. 2009. "The Future of Semiconductor Intellectual Property Architectural Blocks in Europe." EUR 23962 EN. JRC Scientific and Technical Reports. Luxembourg: European Commission. http://ftp.jrc.es/EURdoc/JRC52422.pdf.

———. 2012. "Foresight in an Unpredictable World." *Technology Analysis & Strategic Management* 24 (8): 735–51. https://doi.org/10.1080/09537325.2012.715476.

———. 2013. "Open Educational Resources and the Transformation of Education." *European Journal of Education* 48 (1): 58–78. https://doi.org/10.1111/ejed.12019.
———. 2018. "Vygotsky Meets Backgpropagation: Artificial Neural Models and the Development of Higher Forms of Thought." In *Artificial Intelligence in Education. AIED 2018*. Lecture Notes in Artificial Intelligence, Vol. 10947. Cham: Springer. https://doi.org/10.1007/978-3-319-93843-1_42.
U.S. GAO. 2018. "Artificial Intelligence: Emerging Opportunities, Challenges, and Implications." GAO-18-142SP. United States Government Accountability Office.
Veer, R. van der, and J. Valsiner. 1994. *Understanding Vygotsky: A Quest for Synthesis*. Cambridge, MA: Blackwell Publishers.
Vouloutsi, V., Blancas, M., Zucca, R., Omedas, P., Reidsma, D., Davison, D., ... & Cameron, D. 2016. *Towards a synthetic tutor assistant: the EASEL project and its architecture*. In Conference on Biomimetic and Biohybrid Systems (pp. 353-364). Springer, Cham.
Vygotsky, Lev. 1986. *Thought and Language*. Cambridge, MA: The MIT Press.
Winograd, T., and F. Flores. 1986. *Understanding Computers and Cognition: A New Foundation for Design*. Norwood, NJ: Ablex Publishing Corporation.
Woolf, Beverly Park. 2009. *Building Intelligent Interactive Tutors: Student-Centered Strategies for Revolutionizing e-Learning.* San Francisco, CA: Morgan Kaufmann.
Woollett, and E. A. Maguire. 2011. "Acquiring 'the Knowledge' of London's Layout Drives Structural Brain Changes." *Current Biology* 21 (24): 2109–14.
Yuan, Xiaoyong, Pan He, Qile Zhu, Rajendra Rana Bhat, and Xiaolin Li. 2017. "Adversarial Examples: Attacks and Defenses for Deep Learning." *ArXiv:1712.07107 [Cs, Stat]*, December. http://arxiv.org/abs/1712.07107.
Zysman, John. 2006. "The Algorithmic Revolution---the Fourth Service Transformation." *Communications of the ACM* 49 (7): 48.

KJ-NA-29442-EN-N

The European Commission's science and knowledge service

Joint Research Centre

JRC Mission

As the science and knowledge service of the European Commission, the Joint Research Centre's mission is to support EU policies with independent evidence throughout the whole policy cycle.

EU Science Hub
ec.europa.eu/jrc

@EU_ScienceHub

EU Science Hub - Joint Research Centre

Joint Research Centre

EU Science Hub

Publications Office

doi:10.2760/12297

ISBN 978-92-79-97257-7