Artificial Intelligence In Education

Promises and Implications for Teaching and Learning

Wayne Holmes, Maya Bialik, Charles Fadel

CENTER FOR CURRICULUM REDESIGN

EXCERPT

© Center for Curriculum Redesign – All Rights Reserved

The Center for Curriculum Redesign, Boston, MA, 02130
Copyright © 2019 by Center for Curriculum Redesign
All rights reserved.
Printed in the United States of America

Artificial Intelligence in Education
Promises and Implications for Teaching and Learning

ISBN-13: 978-1-794-29370-0
ISBN-10: 1-794-29370-0

A pdf version with active links is available at: http://bit.ly/AIED-BOOK

Keywords: Education technology, EdTech, artificial intelligence, machine learning, deep learning, AI, AIED, curriculum, standards, competencies, competency, CBL, deeper learning, knowledge, skills, character, metacognition, meta-learning, mindset, 21st century education, social-emotional skills, 21st century competencies, education redesign, 21st century curriculum, pedagogy, learning, jobs, employment, employability, eduployment, Education 2030, fourth Industrial Revolution, exponential technology, personalized learning, competency-based learning, individualized learning, adaptive learning.

Introduction: The Context

Artificial intelligence (AI) is arguably *the* driving technological force of the first half of this century, and will transform virtually every industry, if not human endeavors at large.[1] Businesses and governments worldwide are pouring enormous sums of money into a very wide array of implementations, and dozens of start-ups are being funded to the tune of billions of dollars.

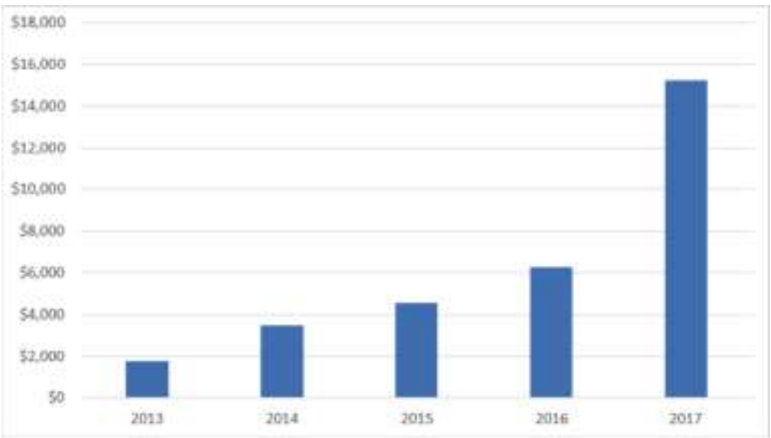

Funding of AI startup companies worldwide, from 2013 to 2017
(in millions of U.S. dollars). *Source*: Statista[2]

It would be naive to think that AI will not have an impact on education—*au contraire*, the possibilities there are profound yet, for the time being, overhyped as well. This book attempts to provide the right balance between reality and hype (per the Gartner diagram that follows), between true potential and wild extrapolations. Every new technology undergoes a period of intense growth of reputation and expectations, followed by a precipitous fall when it inevitably fails to live up to the expectations, after which there is a slower growth as the technology is developed and integrated into our lives. As visualized in the Gartner diagram, each technology can be said to reside somewhere on the curve

[1] Possibly matched only by biotechnology.

[2] https://www.statista.com/statistics/621468/worldwide-artificial-intelligence-startup-company-funding-by-year

at any given time (for example Deep Learning, which is part of AI, is currently peaking).

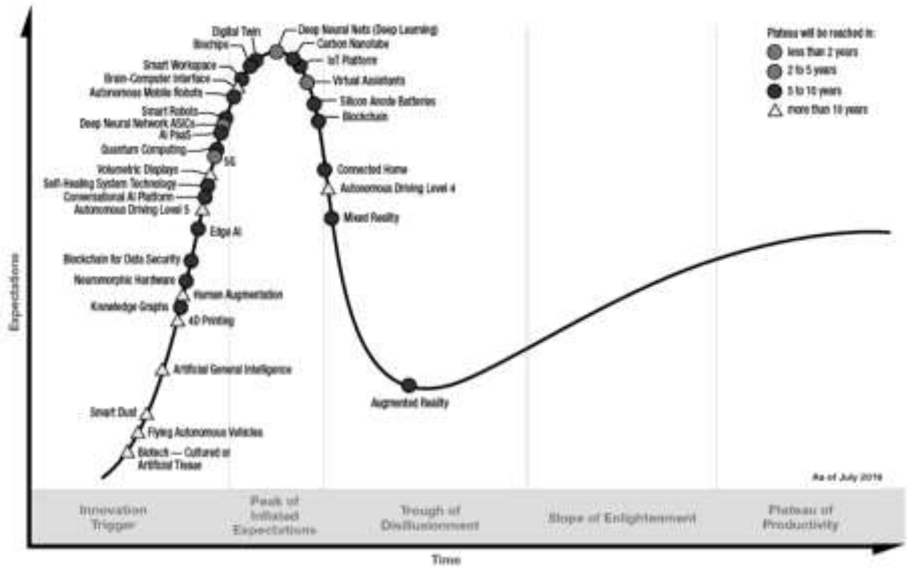

Source: Gartner Inc.[3]

It is of course a risky proposition then, in a field moving so fast, to attempt to predict the future. As such, this work will likely be updated periodically to keep up with the developments (just as you would expect from software/apps).

This book is organized around a somewhat glib quote: "There are only two problems in education: What we teach, and how we teach it."[4] Hence this book is divided into two parts, one focused on the *What*, and one on the *How* of AI in education.

[3] http://www.Gartner.com/SmarterWithGartner
[4] Dr Roger Schank, https://www.rogerschank.com/

The What

> *We're headed for a world where you're either going to be able to write algorithms … or be replaced by algorithms.*
>
> —Bridgewater hedge-fund billionaire Ray Dalio

The first part of this book explores the question: *What* should students learn in an age of AI? And all the corollary, provocatively phrased questions: "If you can search, or have an intelligent agent find, anything, why learn anything? What is truly worth learning?"

It is widely expected that AI will have an enormous impact on what we teach, as it will impact many occupations. Take for instance the Organization for Economic and Co-operative Development (OECD) Programme for the International Assessment of Adult Competencies (PIAAC)[5] survey, which measures adults' proficiency in key information-processing skills—literacy, numeracy and problem solving in technology-rich environments—and gathers information and data on how adults use their skills at home and at work. Already, AI is matching more than 50% of adult human-proficiency levels, and closing in on another 36%.

Proficiency Level	OECD Adults	Artificial Intelligence
2 and below	53%	Yes
3	36%	Close
4–5	11%	No

Source: Elliott Stuart, "Computers and the Future of Skill Demand."[6]

Such progress is bound to continue at an accelerating, pace. IBM's Open Leaderboard effort attempts to understand the progress being made by tracking many variables. According to IBM's Leaderboard, AI should be getting into the realm of deeper self-learning by the early 2020s

[5] https://www.oecd.org/skills/piaac/

[6] https://read.oecd-ilibrary.org/education/computers-and-the-future-of-skill-demand_9789264284395-en#page1

© Center for Curriculum Redesign – All Rights Reserved

and become capable of assisting, collaborating, coaching and mediating by the early 2030s.

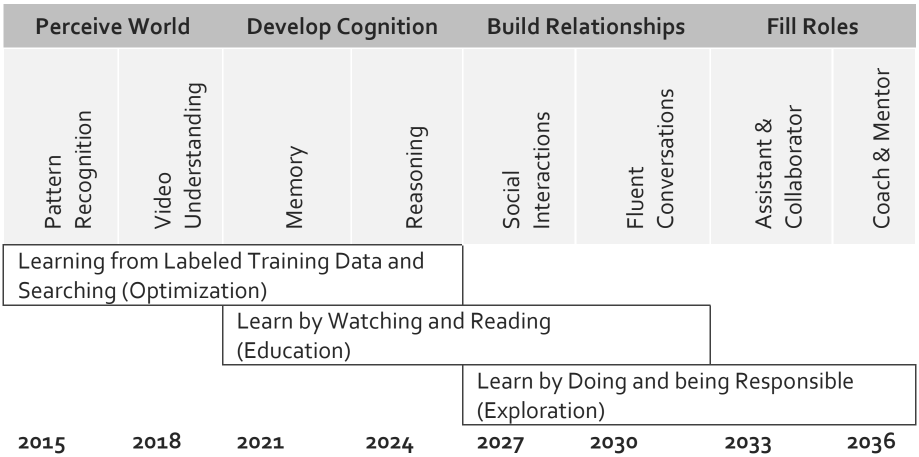

Source: Jim Spohrer, IBM[7]

Given all the above, the *What* section makes a case for the necessity to focus on a broad, deep, and versatile education as a hedge against uncertain futures, which in turn means a reinvigorated focus on the *deeper learning goals* of a modern education:

- Versatility, for robustness to face life and work.
- Relevance, for applicability, and student motivation.
- Transfer,[8] for broad future actionability.

All of which are to be developed via:

- Selective emphasis on important areas of traditional knowledge.
- The addition of modern knowledge.
- A focus on essential content and core concepts.
- Interdisciplinarity, using real-world applications.
- Embedded skills, character, and meta learning *into* the knowledge domains.

[7] IBM, 2017, Cognitive Opentech Group.
[8] This refers to the transfer of knowledge from a domain it was learned in to another domain.

The How

The second part of this book addresses the question: How can AI enhance and transform education? First, it is important to make the distinction between education technology (EdTech) at large and artificial intelligence in education (AIED) specifically. A quick summary of the affordances of EdTech is appropriate at this stage, as the taxonomy and ontology of the field is quite murky. Using the SAMR[9] model that follows, the How section showcases how AIED will span all layers, with its maximum impact growing as it moves up the stack.

Substitution, augmentation, modification, and redefinition model (SAMR).

Note that the examples shown in the preceding figure represent today's apps, not tomorrow's, and only serve to help explain the model. Often these apps are collapsed under one term, technology, and then there is much confusion about the potential of technology. This model helps us to delineate the different types of impact that technology can have, from mere substitution with no functional changes, all the way through the creation of new, previously inconceivable tasks as a result of technology.

[9] Dr. Ruben Puentedura, http://www.hippasus.com/

© Center for Curriculum Redesign – All Rights Reserved

The Role of Assessments

> *What gets measured gets managed.*—Lord Kelvin

Assessments have been the hidden villain behind a lot of education debates, and a powerful one at enshrining institutional inertia. Repurposing the famous Aristotelian syllogism:[10]

> Lack of, or poor, education is at the root of many human problems.
> Assessments define the education we get.
> *Therefore*, assessments are the root of many human problems.

Although not a focus of this book, it is clear that assessments have an oversized role to play in the change process, and as part of the AI-driven systems of (mostly formative) assessments.

Andreas Schleicher, director of the OECD's Directorate of Education and Skills, publicly stated "What is easy to measure is also easy to automate," thereby throwing the gauntlet to the assessment world to readjust its focus and thus *drive* change.

Lastly

Readers will have different priorities and interest in this topic. Policymakers and curriculum designers may initially favor the What section, while teachers and IT specialists may at first favor the How section.

The What and How sections are therefore written to be independent of each other; the appendices also reflect an emphasis on digestibility, particularly for the technical details.

Further, we are all pressed for time, so our writing philosophy is, to use Antoine de Saint-Exupéry's words, "Perfection is attained not when there is no longer anything to add, but when there is no longer anything to take away." This book is therefore not meant to be an in-depth

[10] "All men are mortal. Socrates is a man. Socrates is mortal." https://en.wikipedia.org/wiki/Syllogism

academic piece, but rather it is meant to be concise and to the point, and adhere to Yuval Harari's philosophy: "In a world deluged with irrelevant information, clarity is power."[11]

We wish you all very pleasant reading, and invite your feedback at: info@CurriculumRedesign.org

[11] Harari, Y. (2018). *21 Lessons for the 21st Century.* Spiegel & Grau.

© Center for Curriculum Redesign – All Rights Reserved

Part Two
The How: Promises and Implications of AI for Teaching and Learning

Seldom a day goes by without at least a mention in the news or entertainment media of AI. Perhaps an AI program has just beaten the world's leading player in a complex strategy game; perhaps a new Hollywood feature film depicts a dystopian future in which robots have overtaken the world; or perhaps a pair of leading tech entrepreneurs have a public disagreement.[12]

> I have exposure to the very cutting-edge AI, and I think people should be really concerned about it. … AI is a rare case where we need to be proactive about regulation instead of reactive. Because I think by the time we are reactive in AI regulation, it's too late.
>
> —Elon Musk

> I think people who are naysayers and try to drum up these doomsday scenarios … I just, I don't understand it. It's really negative and, in some ways, I actually think it is pretty irresponsible.
>
> —Mark Zuckerberg

In fact, as the exchange between Mark Zuckerberg (Facebook) and Elon Musk (SpaceX, Tesla) suggests, the future impact of AI remains very unclear (indeed, is artificial intelligence little more than the latest technical hype?).[13] Nevertheless, investments and developments continue to grow exponentially, such that AI has become an integral, pervasive

[12] E.g., https://www.nytimes.com/2018/06/09/technology/elon-musk-mark-zuckerberg-artificial-intelligence.html

[13] "Highly-publicized projects like Sophia [http://www.hansonrobotics.com/robot/sophia] [try to] convince us that true AI—human-like and perhaps even conscious— is right around the corner. But in reality, we're not even close. The true state of AI research has fallen far behind the technological fairy tales we've been led to believe. And if we don't treat AI with a healthier dose of realism and skepticism, the field may be stuck in this rut forever." Dan Robitzski quote (2018) at https://futurism.com/artificial-intelligence-hype

and inescapable, although often hidden, part of our daily lives: from Siri[14] to auto-journalism,[15] from forecasting stock movements[16] to predicting crime,[17] from facial recognition[18] to medical diagnoses[19] and beyond.

But of particular interest here, artificial intelligence has also quietly entered the classroom.[20] Whether students, teachers, parents and policy makers welcome it or not, so-called intelligent, adaptive, or personalized learning systems are increasingly being deployed in schools and universities[21] around the world, gathering and analyzing huge amounts of student big data, and significantly impacting the lives of students and educators.[22] However, while many assume that artificial intelligence in education (AIED) means students being taught by robot teachers, the reality is more prosaic yet still has the potential to be transformative. Nonetheless, the application of AI to education raises far-reaching questions.

> We should ask what happens when we remove care from education.... What happens to thinking and writing when... the whole educational process is offloaded to the machines—to "intelligent tutoring systems," "adaptive learning systems," or whatever the latest description may be? What sorts of signals are we sending students?
>
> —Audrey Watters[23]

---

[14] https://www.apple.com/uk/ios/siri

[15] E.g., https://www.washingtonpost.com/pr/wp/2018/06/12/the-washington-post-plans-extensive-coverage-of-2018-midterm-elections/?utm_term=.e66d88e4a716

[16] E.g., https://equbot.com

[17] E.g., http://www.predpol.com

[18] E.g., https://www.cbp.gov/newsroom/national-media-release/cbp-deploys-facial-recognition-biometric-technology-1-tsa-checkpoint

[19] E.g., https://www.babylonhealth.com

[20] Luckin, R., et al. (2016). *Intelligence Unleashed. An Argument for AI in Education*. Pearson. https://www.pearson.com/content/dam/one-dot-com/one-dot-com/global/Files/about-pearson/innovation/Intelligence-Unleashed-Publication.pdf

[21] "The time-to-adoption for adaptive learning technologies and artificial intelligence is estimated within two to three years, acknowledging the advances in these technologies and their promise to positively impact teaching and learning." Becker, S.A., et al. (2018). "Horizon Report: 2018." *Higher Education Edition 2*.

[22] Holmes, W., et al. (2018). *Technology-Enhanced Personalised Learning. Untangling the Evidence*. Robert Bosch Stiftung.

[23] http://hackeducation.com/2015/08/10/digpedlab

© Center for Curriculum Redesign – All Rights Reserved

In fact, AI technologies have been researched in educational contexts for around fifty years.[24] More recently, companies as influential as Amazon, Google and Facebook have invested millions of dollars[25] developing AIED products, joining well-established multimillion dollar-funded AIED companies such as Knewton[26] and Carnegie Learning,[27] while the $15 million Global Learning XPrize[28] called for software that empowers children to take control of their own learning (AIED by another name). Meanwhile, AI is being introduced into some mainstream schools as a curriculum in its own right,[29] is being developed to improve online tutoring,[30] and is being researched as a way of enhancing teacher training.[31] In short, the application of AI in educational contexts is growing exponentially,[32] such that by 2024 it is predicted to become a market worth almost $6 billion.[33]

While we may have some limited knowledge or experience of mainstream AI, either from the media or in our daily lives, for many the use of AI in education remains a mystery. A multitude of yet-to-be-answered questions spring to mind. How exactly can AI work in classrooms, and what can be achieved? With AI requiring so much data, how is student privacy maintained? What will be AI's long-term effects on teacher roles? Are the proponents of AIED promising more than can

---

[24] Woolf, B. (1988). "Intelligent tutoring systems: A survey.' In *Exploring Artificial Intelligence*: 1–43; Cumming, G., and McDougall, A. (2000). "Mainstreaming AIED into education?" *International Journal of Artificial Intelligence in Education* 11: 197–207; du Boulay, B. (2016). "Artificial intelligence as an effective classroom assistant." *IEEE Intelligent Systems* 31 (6): 76–81. https://doi.org/10.1109/MIS.2016.93

[25] https://www.linkedin.com/pulse/tech-giants-quietly-invest-adaptive-learning-system-rd-drew-carson

[26] http://www.knewton.com

[27] http://www.carnegielearning.com

[28] https://learning.xprize.org

[29] http://www.gettingsmart.com/2018/07/coming-this-fall-to-montour-school-district-americas-first-public-school-ai-program

[30] Following a $300m investment, the Chinese online tutoring company Yuanfudao has set up a research institute for artificial intelligence, which aims to train its homework app to be smarter. https://techcrunch.com/2018/12/26/yuanfudao-raises-300-million/

[31] O'Connell, S. (2018). "New Project Aims to Use Artificial Intelligence to Enhance Teacher Training." Center for Digital Education. http://www.govtech.com/education/higher-ed/New-Project-Aims-to-Use-Artificial-Intelligence-to-Enhance-Teacher-Training.html

[32] https://www.eschoolnews.com/2017/05/22/brace-ai-set-explode-next-4-years

[33] https://www.gminsights.com/industry-analysis/artificial-intelligence-ai-in-education-market

be delivered? What is the impact of AI on student agency and outcomes? And what are the social and ethical consequences?

However, we begin with a tentative response to an ostensibly simpler question: what does AIED actually look like?

AI in Education

As a brief review of AIED conference and journal papers will confirm, AIED includes everything from AI-driven, step-by-step personalized instructional and dialogue systems, through AI-supported exploratory learning, the analysis of student writing, intelligent agents in game-based environments, and student-support chatbots, to AI-facilitated student/tutor matching that puts students firmly in control of their own learning. It also includes students interacting one-to-one with computers, whole-school approaches, students using mobile phones outside the classroom, and much more besides. In addition, AIED can also shine a light on learning and educational practices.

The field of AIED is both derivative and innovative. On the one hand, it brings theories and methodologies from related fields such as AI, cognitive science, and education. On the other hand, it generates its own larger research issues and questions: What is the nature of knowledge, and how is it represented? How can an individual student be helped to learn? Which styles of teaching interaction are effective, and when should they be used? What misconceptions do learners have?[34]

While AIED tools necessarily instantiate specific learning theories (such as Gagné's "instructionalism"[35] or Vygotsky's "zone of proximal development"[36]), some AIED researchers question the assumptions behind those theories, applying AI and data analysis techniques to try to open the "black box of learning."[37] In other words, AIED effectively

---

[34] Woolf, B.P. (2010). *Building Intelligent Interactive Tutors: Student-Centered Strategies for Revolutionizing e-Learning.* Morgan Kaufmann, 11.

[35] Gagné, R.M. (1985). *Conditions of Learning and Theory of Instruction, 4th Revised Edition*. Wadsworth Publishing Co Inc.

[36] Vygotsky, L. S. (1978). *Mind in Society: Development of Higher Psychological Processes*. Harvard University Press.

[37] Luckin, R., et al. *Intelligence Unleashed. An Argument for AI in Education*.

© Center for Curriculum Redesign – All Rights Reserved

involves two main complementary strands: developing AI-based tools to support learning, and using these tools to help understand learning (how learning happens and other questions that have long been investigated by the learning sciences, and which might be applied in classrooms whether or not AI is being used). For example, by modeling how students go about solving an arithmetic problem and identifying misconceptions that might have been previously unknown to educators, researchers and teachers can begin to understand much more about the process of learning itself, which might then be applied to mainstream classroom practices.

In fact, new AIED applications and approaches, addressing old and newly identified problems, are being researched and released all the time—such that, what AIED looks like and can do is still emerging. Accordingly, here we adopt an alternative approach. Rather than trying to define AIED, within some relatively easy-to-identify broad areas we will discuss a wide range of AIED examples—existing AIED tools and AIED tools that might be available in the not-too-distant future.

What this book does not include, but what might nonetheless have a major impact on education, is the use of AI to support school and university administrative functions such as class timetabling, staff scheduling, facilities management, finances, cybersecurity, safety and security. Instead, we focus on the use of AI to support learning, what might be called the academic (or system-facing) functions of AIED.

However, before doing so, it will be helpful to have at least a working understanding of AI itself.[38] That is where we go now, before returning to look in more detail at how AI works in educational contexts. We conclude by considering the various challenges, pragmatic and ethical, from the perspectives of AIED researchers and developers, as well as of educators, students, funders and policy-makers.

---

[38] It has been argued that "Artificial Intelligence should be accessible to all of us, even without a math background." https://www.youtube.com/watch?v=LqjP7O9SxOM&list=PLtmWHNX-gukLQlMvtRJ19s7-8MrnRV6h6

The Background of AI

AI is one of those aspects of modern life about which most of us have some awareness, and yet recognize we have little knowledge.[39] In fact, for many AI is synonymous with humanoid robots,[40] which might be because news about AI is almost always illustrated with a picture of a robot or a digital brain. However, while robotics (embodied AI, that can move and physically interact with the world) is a core area of AI research, AI is being applied in many different ways and different contexts. Meanwhile, the dystopian images of futuristic robots remain firmly in the realm of science fiction (which is why for the most part we leave robotics well alone). In the next few pages, we provide a brief background to artificial intelligence; interested readers will find more information about the origins and development of AI and its various techniques in appendix 2.

However, first, we should acknowledge that the very name artificial intelligence is sometimes seen as unhelpful. Instead, some researchers prefer augmented intelligence, which retains the human brain as the source of intelligence, and positions the computer and its programs as a sophisticated tool with which humans might enhance or augment our intellectual capabilities. In this approach, computers are employed to do what humans find more difficult (such as finding patterns in huge amounts of data). The debate contrasting augmented and artificial will inevitably run and run, with artificial intelligence winning at least on popular usage even if augmented intelligence is more accurate or useful. Accordingly, hereafter we will take the ultimate pragmatic approach and refer almost exclusively to AI, leaving the reader to decide for themselves what the A in AI represents.

[39] E.g., "What is artificial intelligence?" https://www.brookings.edu/research/what-is-artificial-intelligence

[40] The question "When you think of AI, what is the first thing that comes into your head?" has been posed in numerous lectures and surveys about AI in education. The evidence is anecdotal, but overwhelmingly participants answer: robots.

© Center for Curriculum Redesign – All Rights Reserved

A 1956 workshop held at Dartmouth College, a US Ivy League research university, is widely considered to be AI's foundational event.[41] It was where what is thought to be the first AI program, the Logic Theorist, was presented and discussed. Over the following decades, AI developed in fits and starts with periods of rapid progress (for example, rule-based expert systems) interspaced with periods known as AI winters where confidence and funding all but evaporated.

In recent decades, thanks to three key developments (the advent of faster computer processors, the availability of large amounts of big data, and advances in computational approaches) AI has entered a period of renaissance—AI has now become an integral, pervasive, and inescapable, although often hidden, part of our daily lives. In fact, paradoxically, the more that it is integrated, the less we tend to think of it as AI.

> A lot of cutting edge AI has filtered into general applications, often without being called AI because once something becomes useful enough and common enough it's not labeled AI anymore.[42]

Instead, AI is often known as an advanced computer program (such as email spam filtering),[43]a personal assistant (such as Cortana),[44] a recommendation system (such as in Netflix),[45] or perhaps a language-learning app (such as Duolingo).[46] Having said that, recent voice-activated smart speakers, such as Google Home[47] and Amazon Echo,[48] have made AI more visible in our living rooms. In fact, many recent developments in AI are both groundbreaking and in many ways transformative. Relatively recent AI computational approaches, such as machine learning

---

[41] Crevier, D. (1993). *AI. The Tumultuous History of the Search for Artificial Intelligence*. Basic Books.

[42] http://edition.cnn.com/2006/TECH/science/07/24/ai.bostrom/index.html (Professor Nick Bostrom, director of the Future of Humanity Institute, University of Oxford).

[43] E.g., https://www.mailwasher.net uses Bayesian techniques to learn which emails are spam and which are not.

[44] https://www.microsoft.com/en-us/cortana

[45] https://help.netflix.com/en/node/9898

[46] https://www.duolingo.com

[47] https://store.google.com/gb/product/google_home

[48] https://www.amazon.com/b/?ie=UTF8&node=9818047011

(supervised, unsupervised, and reinforcement learning), neural networks (including deep learning), and evolutionary algorithms have all been used in a diverse range of applications (interested readers will find more information about these techniques in appendix 2).

For example, recent advances in face recognition (ensuring that faces in smartphone photographs are always in sharp focus, and identifying travelers at e-passport gates) are thanks to the application of neural networks and machine learning. Google researchers presented a brain-inspired AI neural network with 10 million randomly selected video thumbnails from YouTube.[49] By using deep-learning techniques, and despite not being told how to recognize anything in particular, this machine learning system soon learned how to detect human faces in photographs. Two years later, Facebook introduced a nine-layer deep AI neural network, involving more than 120 million parameters, to *identify* (not just detect) faces in timeline photographs.[50] It was trained on a dataset of four million images of faces that had previously been labeled by humans (the Facebook users, who had been happily labeling their friends in uploaded photographs over several years).

Another area that has seen much AI development in recent years is autonomous vehicles, with neural networks being used to enable cars, trucks, and taxis to drive without human intervention. A complex rig of cameras and sensors collate massive amounts of real-time data (the road's edges and markings, road signs and traffic lights, other vehicles including bicycles, other potential obstacles, and pedestrians), while a neural network-driven intelligent agent, drawing on massive of computing power, controls the car's steering, acceleration, and braking. A probably less well-known use of AI is in journalism. News organizations around the world are developing AI technologies to support their news gathering and news reporting. For example, AI agents continually monitor global

[49] https://www.nytimes.com/2012/06/26/technology/in-a-big-network-of-computers-evidence-of-machine-learning.html?_r=1

[50] Facebook introduced a nine-layer deep AI neural network, involving more than 120 million parameters, to identify (not just detect) faces in timeline photographs. It was trained on a dataset of four million images.

© Center for Curriculum Redesign – All Rights Reserved

news outlets and use semantic analysis to automatically extract key information that is made available to the journalists to write their stories.[51] There are even some AI technologies that go one step further and automatically write the stories themselves.[52]

Another application of AI is in law, where e-Discovery tools are being used by lawyers to help process the huge amounts of documentation that need to be reviewed as potential evidence in civil or criminal legal cases.[53] One technique involves a machine learning analysis of a sample of documents that have been reviewed and labeled by an expert. The outcomes enable the AI to then identify which of the remaining documents need to be prioritized for in-depth review. A final brief example is the use of AI in medical diagnoses. For example, AI techniques are used by radiologists to help them identify anomalies in medical images more quickly while making fewer mistakes.[54] One system looks for irregularities in X-ray images. If, for example, it finds nodules on an image of a pair of lungs, it sends it to a pulmonary radiologist for further checks.

AI Techniques and Terminology

While it is relatively straightforward to understand what the applications of AI outlined in the previous section are doing, understanding how they are doing it can require some highly technical knowledge—exacerbated by the fact that any one AI application might draw on several different AI techniques. This is one reason why many people involved in AI have advanced degrees in mathematics or physics (although AI is increasingly being offered as a service: for example, Amazon's Machine Learning on AWS,[55] Google's TensorFlow,[56] IBM's Watson,[57] and Microsoft's

[51] E.g., http://bbcnewslabs.co.uk/projects/juicer

[52] E.g., https://narrativescience.com/Products/Our-Products/Quill

[53] E.g., https://talkingtech.cliffordchance.com/en/emerging-technologies/artificial-intelligence/ai-and-the-future-for-legal-services.html

[54] Hosny, A., et al. (2018). "Artificial intelligence in radiology." *Nature Reviews Cancer* 18 (8): 500–510. https://doi.org/10.1038/s41568-018-0016-5

[55] https://aws.amazon.com/machine-learning

[56] https://www.tensorflow.org

Azure).[58] Nonetheless, because some AI techniques have already been repeatedly mentioned, and because they play important roles in AIED and so will be mentioned again, some key and closely interlinked AI techniques and terminologies will next be introduced.[59] At times (despite our best efforts) our discussion will be somewhat technical; so please feel free to jump to the next section, to move directly onto our discussion of the application of AI in education (which, after all, is the reason we are all here).

Algorithms

Algorithms are at the core of AI, such that the history of AI might be thought of as the history of the development of increasingly sophisticated and increasingly efficient (or elegant) algorithms. Probably the most famous algorithm of recent times is PageRank,[60] which was developed in 1996 by the founders of Google while they were students at Stanford University. It ranks the relative importance of a website, by counting the number of external links to the website's pages, to determine where the website appeared in a Google search. In fact, all computer programs are algorithms. They comprise hundreds if not thousands of lines of code, representing sets of mathematical instructions that the computer follows in order to solve problems (compute a numerical calculation, grammar-check an essay, process an image, or explain patterns that we see in nature).[61] All that makes AI algorithms distinct from other computer programs is that they involve some specific approaches and, as we have noted, they are applied to areas we might think of as essentially human—such as visual perception, speech recognition, decision-making and learning.

---

[57] https://www.ibm.com/watson

[58] https://azure.microsoft.com

[59] Readers wishing to learn more about AI techniques might be interested in Russell, S. and Norvig, P. (2016). *Artificial Intelligence: A Modern Approach, 3rd Edition.* Pearson; and Domingos, P. (2017). *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World.* Penguin.

[60] PageRank of Site = Σ [PageRank of inbound link/Number of links on that page].

[61] Turing, A. (1952). "The chemical basis of morphogenesis." *Philosophical Transactions of the Royal Society* 237 (641): 37–72.

© Center for Curriculum Redesign – All Rights Reserved

Machine Learning

Much early, rule-based AI involves writing in advance the steps that the computer will take to complete a task, rules that will be followed exactly. Machine learning, on the other hand, is about getting computers to act without being given every step in advance. Instead of the algorithms being programmed exactly what to do, broadly speaking they have the ability to learn what to do. This is not to suggest that machine learning does not require large amounts of programming, because it does. But rather that, instead of direct commands leading to direct outputs, machine learning involves large amounts of input data to predict novel outcomes.

Machine learning algorithms analyze the data to identify patterns and to build a model which is then used to predict future values (for example, by identifying patterns in historical stocks data, AI predicts future stock movements; by identifying patterns in photographs of named people, it predicts who is shown in other photographs; and by identifying patterns in medical symptoms, it predicts a specific diagnosis). In other words, machine learning may be considered a three-step process (analyze data, build a model, undertake an action) that is continuously iterated (the outcomes of the action generate new data, which in turn amends the model, which in turn causes a new action). It is in this sense that the machine is learning.

Many recent applications (including natural language processing, self-driving cars, and the Google DeepMind AlphaGo program that beat the world's number one player of Go)[62] have all been made possible thanks to machine learning. In fact, machine learning is so widespread today that, for some commentators, AI and machine learning have become synonymous—whereas machine learning is more properly a sub-field of AI. What is true, however, is that the renaissance and exponential growth of AI over the last decade, has come about because of significant

[62] https://www.theguardian.com/technology/2016/mar/15/googles-alphago-seals-4-1-victory-over-grandmaster-lee-sedol

advances in machine learning (based on, as we have noted, faster computer processors, the availability of large amounts of big data, and new computational approaches).[63]

There are three main categories of machine learning: supervised, unsupervised, and reinforcement learning.

Supervised Learning

Most practical machine learning involves supervised learning. The AI is first provided large amounts of data for which the output is already known—in other words, data that has already been labeled. For example, the AI might be given many thousands of photographs of streets in which the numerous visible objects (bicycles, road signs, pedestrians, etc.) have already been identified and labeled by humans. The supervised learning algorithm aims to identify the function that links the data to the labels, from which it builds a model that can be applied to new similar data. This is broadly speaking the approach, mentioned earlier, used by Facebook to identify people in photographs, which used millions of photographs submitted and labeled by Facebook users to identify and label automatically the same people in new photographs.

Unsupervised Learning[64]

In unsupervised learning, the AI is provided with even larger amounts of data, but this time data that has not been categorized or classified, that is to say data that is not labeled. By analyzing this unlabeled data, unsupervised learning algorithms aim to uncover hidden patterns in the underlying structure of the data, clusters of data that can be used to classify new data (this is broadly the approach, mentioned earlier, used by Google to detect faces in photographs). Example applications of unsupervised learning include dividing online shoppers into groups so

[63] Interestingly, the origins of machine learning can be traced back to at least 1959, with the publication of "Some Studies in Machine Learning Using the Game of Checkers" by an IBM researcher.

[64] A comprehensive list of the algorithms available on one of the leading "AI as a service" platforms, Microsoft Azure, is available at http://download.microsoft.com/download/A/6/1/A613E11E-8F9C-424A-B99D-65344785C288/microsoft-machine-learning-algorithm-cheat-sheet-v6.pdf

© Center for Curriculum Redesign – All Rights Reserved

they can be served tightly targeted advertisements;[65] identifying different letters and numbers from examples of handwriting; and distinguishing between legitimate and fraudulent financial transactions.

Reinforcement Learning

In some senses, reinforcement learning is the most powerful of the machine learning categories. In both supervised and unsupervised learning, the model derived from the data is fixed, and if the data changes the analysis has to be undertaken again (in other words, the algorithm is run once more). However, reinforcement learning involves continuously improving the model based on feedback—in other words, this is machine learning in the sense that the learning is ongoing. The AI is provided with some initial data from which it derives its model, which is evaluated, assessed as correct or incorrect, and rewarded or punished accordingly (to use a computer game metaphor, its score is increased or reduced). The AI uses this positive or negative reinforcement to update its model and then it tries again, thus developing iteratively (learning and evolving) over time. For example, if an autonomous car avoids a collision, the model that enabled it to do so is rewarded (reinforced), enhancing its ability to avoid collisions in the future.

Artificial Neural Networks

An artificial neural network is an AI algorithm that is based on the structure and functions of biological neural networks (i.e. animal brains), that might be applied in advanced supervised, unsupervised, or reinforcement learning. Our brains are made up of billions of individual neurons, each of which is connected to as many as a thousand other neurons, giving trillions of connections. Memory is thought to emerge from complex combinations of these connections across the brain, while learning is thought to involve the strengthening of those connections.

---

[65] In a now-infamous story, the US retailer Target automatically identified a teenager as being pregnant, before she had told anyone, just by her store purchases, and some unsupervised learning. https://www.forbes.com/sites/kashmirhill/2012/02/16/how-target-figured-out-a-teen-girl-was-pregnant-before-her-father-did/#31650c296668

Although artificial neural networks have been trained to do some incredible things (such as identifying faces in moving crowds of people), they remain primitive in comparison to higher-order animal brains. Unlike for example the human brain's billions of neurons, they usually involve only a few thousand neurons (in some exceptional cases, a few million).

As illustrated in the following figure, artificial neural networks each comprise three types of layers: an input layer that takes stimuli from the environment, in the form of millions of data points, perhaps pixels from images; at least one, but often many more, hidden intermediary layers that together undertake the computation; and an output layer that delivers the result. During the machine learning process, weightings given to the connections are adjusted in a process of reinforcement learning, which allows the artificial neural network subsequently to compute outputs for new stimuli.

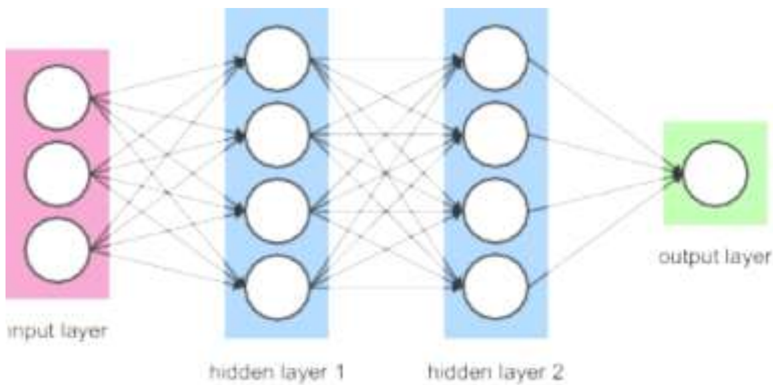

A representation of a typical, simple artificial neural network, with two hidden layers.

The hidden layers are the key to the power of artificial neural networks, but they also bring an important problem. It isn't possible (or at the very least it isn't easy) to interrogate an artificial neural network to find out how it came up with its solution—for example, how did it identify a particular person in a photograph? In other words, artificial neural networks can lead to decision making for which the rationale is

© Center for Curriculum Redesign – All Rights Reserved

hidden and unknowable, or un-inspectable, and possibly unjust,[66] a critical issue that is the subject of much research.[67]

Finally, the impressive results of neural networks and other machine learning technologies should not beguile us:

> A neural network of today no more "learns" or "reasons" about the world than a linear regression of the past. They merely induce patterns through statistics. Those patterns may be opaquer, more mediated and more automatic than historical approaches and capable of representing more complex statistical phenomena, but they are still merely mathematical incarnations, not intelligent entities, no matter how spectacular their results.[68]

---

[66] O'Neil, C. (2017). *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy.* Penguin.

[67] Morcos, A.S., et al. (2018). "On the importance of single directions for generalization." ArXiv:1803.06959. http://arxiv.org/abs/1803.06959

[68] https://www.forbes.com/sites/kalevleetaru/2018/12/15/does-ai-truly-learn-and-why-we-need-to-stop-overhyping-deep-learning/#edd206168c02

How AI Works in Education

Having established a working understanding of AI, we will now look in more detail at how AI works in educational contexts, beginning with a brief history. However, what will not be discussed but what might nonetheless have a major impact on education is the use of AI to support school administration (system-facing AI that addresses things like class timetabling, staff scheduling, facilities management, finances, cybersecurity, safety and security).[69] Our focus is the use of AI to support learning (student- and teacher-facing AI).

A Brief History of AI in Education

Precursors of the application of AI in education can be found in the work of the psychologists Sidney Pressey, who was a professor at Ohio State University in the 1920s, and B. F. Skinner, known as the father of behaviorism, who was a professor at Harvard University from 1948 until his retirement in 1974. For Pressey, the challenge was to leverage the potential of multiple-choice tests to consolidate student learning as well as to evaluate it. Drawing on Edward Thorndike's law of effect,[70] he argued that, for tests to support learning, immediate feedback was essential—which is not usually possible when tests are marked by hand. However, a mechanical approach could ensure that no learning opportunities were missed.

> Devices which at once inform a student about the correctness of his answer to a question, and then lead him to the right answer, clearly do more than test him; they also teach him.[71]

---

[69] Readers who are interested in the use of AI technologies to support administrative functions might like to read about Ofsted, the UK's school inspection service. Ofsted's use of "artificial-intelligence algorithm to predict which schools are 'less than good'." https://www.tes.com/news/ofsted-use-artificial-intelligence-algorithm-predict-which-schools-are-less-good

[70] Thorndike. E.L. (1927) "The Law of Effect." *The American Journal of Psychology* 39 (1/4): 212–22. https://doi.org/10.2307/1415413

[71] Pressey, S.L. (1950). "Development and appraisal of devices providing immediate automatic scoring of objective tests and concomitant self-instruction." *Journal of Psychology* 30: 417–447.

© Center for Curriculum Redesign – All Rights Reserved

Pressey made various versions of his machine (and made several unsuccessful attempts to commercialize his idea), the most sophisticated being based on a mechanical typewriter. Inside this device was a rotating drum around which was wrapped a card printed with a list of questions and hole-punched (much like the perforated rolls used in self-playing pianos) to represent the correct answers. Meanwhile, the casing featured a small window, which showed the number of the current question, and five typewriter keys, one for each possible answer. As the student worked through a printed sheet of questions and answers, they would press one of the keys on the device to select their answer for each question. The machine was configured so that the student immediately knew whether they had made the right choice, and it prevented them from moving onto the next question until they had.

Interestingly, Pressey was also one of the first to make the case that, in addition to supporting learning, a teaching machine could make a teacher's life easier and more fulfilling—by relieving them of one of their least interesting tasks (marking tests) and giving them more time to engage with their students.

> Lift from [the teacher's] shoulders as much as possible of this burden and make her [sic] free for those inspirational and thought-stimulating activities which are, presumably, the real function of the teacher.[72]

Pressey's approach was later extended by Skinner, who argued that the techniques he pioneered for training rats and pigeons (in operant conditioning chambers now known as Skinner Boxes) might be adapted for teaching people. Skinner's teaching machine, which he devised in 1958, was a wooden box with a windowed lid. Questions written on paper disks appeared in one window, and the student wrote a response on a roll of paper accessible through a second window (for later marking by a teacher). Advancing the mechanism automatically covered the

[72] Pressey, S.L. (1926). "A simple device for teaching, testing, and research in learning." *School and Society* 23: 374.

student's answer, so that it could not be changed, and simultaneously revealed the correct answer. In this way, Skinner's teaching machine provided automatic, immediate reinforcement. Students were required to compose their own answers, rather than choose from a limited selection (as with Pressey's multiple-choice questions), because Skinner found that learning is more effectively reinforced by recalling a correct response than by simply recognizing it. This approach also gave the student the opportunity to compare their answer with the given model answer, which if properly designed by the teacher and actively undertaken by the student could also contribute to learning.

Skinner argued that his teaching machine in effect acted like a personal tutor, foreshadowing AIED's intelligent tutoring systems.

> The machine itself, of course, does not teach ... but the effect upon each student is surprisingly like that of a private tutor.... (i) There is a constant interchange between program and student.... (ii) Like a good tutor, the machine insists that a given point be thoroughly understood ... before the student moves on.... (iii) Like a good tutor, the machine presents just that material for which the student is ready.... (iv) Like a skillful tutor, the machine helps the student to come up with the right answer.... (v) Lastly, of course, the machine, like the private tutor, reinforces the student for every correct response, using this immediate feedback ... to shape his behavior most efficiently.[73]

Skinner's teaching machine might be thought to have also foreshadowed something else later taken up by AI in education researchers, dividing automated teaching into separate components (in Skinner's case, distinguishing between the subject content, which was pre-programmed into the machine, and the student's achievements, whether or not they answered a question correctly). However, although in a sense Skinner's teaching machine was responsive to individual students, it could not be considered adaptive. That is to say, it did not adapt either the questions, or the order in which they were presented,

---

[73] Skinner, B.F. (1958). "Teaching machines." *Science* 128 (3330): 969–77.

© Center for Curriculum Redesign – All Rights Reserved

according to the achievements or needs of the individual students. Instead, question delivery was pre-scripted. While a student could proceed at their own pace, they went through the same list of questions as every other student and in the same order.

Adaptive Learning

Also working in the 1950s, Norman Crowder, who was interested in communication rather than psychology, devised a paper-based alternative to the early teaching machines, known as intrinsic or branching programmed instruction.[74] In Crowder's system (which he developed for training U.S. Air Force engineers to find malfunctions in electronic equipment), the user is presented with a short page of information followed by a multiple-choice question, with each possible answer directing the student to a new page. If the correct answer was chosen, the new page would present new information, building upon that which was correctly understood; if an incorrect answer was chosen, the new page would contain feedback designed to help the student understand the cause of their error, based on what the student had chosen. The system might also branch through one or two additional pages of corrective materials before returning the student back to the main pages. In short, Crowder's system adapted the pathway through the teaching materials according to the individual student's developing knowledge, such that each student might see quite different sets of pages.

However, a British polymath, Gordon Pask, probably developed the first truly adaptive teaching machine in the early 1950s. Known as SAKI (the self-adaptive keyboard instructor), it was designed for trainee keyboard operators learning how to use a device that punched holes in cards for data processing.[75] What distinguished SAKI from the other early teaching machines was that the task presented to a learner was

---

[74] Crowder, N.C. (1960). "Automatic tutoring by means of intrinsic programming." In *Teaching Machines and Programmed Learning: A Source Book*. Vol. 116. Lumsdaine, A.A., and Glaser, R. (eds.) American Psychological Association, 286–298.

[75] Pask, G. (1982). "SAKI: Twenty-five years of adaptive training into the microprocessor era." *International Journal of Man-Machine Studies* 17 (1): 69–74. https://doi.org/10.1016/S0020-7373(82)80009-6

adapted to the learner's individual performance, which was represented in a continuously changing probabilistic student model.

When you interact with the system, learning which keys represent which numbers:

> the machine is measuring your responses, and building its own probabilistic model of your learning process. That "7," for instance, you now go to straight away. But the "3," for some obscure reason, always seems to elude you. The machine has detected this, and has built the facts into its model. And now, the outcome is being fed back to you. Numbers with which you have difficulty come up with increasing frequency in the otherwise random presentation of digits. They come up more slowly, too, as if to say: "Now take your time." The numbers you find easy, on the contrary, come up much faster: the speed with which each number is thrown at you is a function of the state of your learning.[76]

Computer-Aided Instruction

SAKI went through many iterations, taking advantage of developments in computers and the new microprocessors, and was one of the first adaptive systems to be commercialized. However, over the following years, other than in the various iterations of SAKI, adaptive learning made few advances, and the focus shifted to what became known as computer-aided instruction (CAI) systems. The 1960s and 1970s saw many CAI systems being built, an early influential example being PLATO (programmed logic for automatic teaching operations), which was developed at the University of Illinois. PLATO involved students accessing standard teaching materials, some of which were interactive, on a central mainframe computer via remote terminals, with as many as a thousand students working at the same time.

This system was also notable for being the first to introduce in an educational technology many tools and approaches still common today, such as user forums, email, instant messaging, remote screen-sharing, and

[76] Beer, S. (1960). *Cybernetics and Management.* The English Universities Press, 124.

© Center for Curriculum Redesign – All Rights Reserved

multiplayer games. Around the same time, Stanford University and IBM developed a computer-aided instruction system that was made available via remote terminals to a few local elementary schools. This system involved a linear presentation of teaching materials, for mathematics and language arts, together with drill and practice activities. A third prominent example was TICCIT (time-shared interactive computer-controlled information television), developed by Brigham Young University, which was used to teach freshman-level mathematics, chemistry, physics, English, and various language courses. Each subject area was broken down into topics and learning objectives, which in turn were represented as screens of information. TICCIT then provided a predetermined sequence, although learners could also use the keyboard to navigate through the screens in any order that they found helpful.

Although in other ways successful, during the 1960s and 1970s only very few of these CAI systems were widely adopted, mainly due to the cost and accessibility of the university mainframes that were needed to host the software. The arrival of personal computers in the 1980s changed everything, with the number of CAI programs quickly mushrooming. Very soon, CAI programs addressing every aspect of learning were being widely used in schools, universities and family homes. Nonetheless, of particular relevance for our present purposes, almost all of these systems were severely hampered by the same flaw—a lack of adaptivity. The sequence of topics, the information provided, and the system's response to student actions, were predefined and the same for each student, ignoring the individual student's successes, misconceptions, and interests, and thus inhibiting their potential learning.

AI and CAI

At around the same time that mainstream CAI was developing, other researchers, such as John Self[77] and William Clancey,[78] were wondering

[77] Self, J.A. (1974). "Student models in computer-aided instruction." *International Journal of Man-Machine Studies* 6 (2): 261–76. https://doi.org/10.1016/S0020-7373(74)80005-2

[78] Clancey, W.J. (1983). "GUIDON." *Journal of Computer-Based Instruction* 10 (1–2): 8–15.

how CAI could become adaptive to the needs of individual students, and whether the application of techniques from AI might be useful. The first implementation of AI techniques in computer-aided instruction is usually credited to Jaime Carbonell, who, in his 1970 PhD thesis, introduced a system called SCHOLAR. This novel system was designed to enable a student to review their knowledge of the geography of South America, by means of a dialogue using a subset of English. SCHOLAR was qualitatively different from standard CAI because, unique for the time, it was capable of generating individual responses to student statements by drawing from a semantic network (a way of representing knowledge, in this case knowledge of geography, in which concepts are linked according to their semantic relationships).

> In [a] conventional CAI, the database consists of many "frames" of specific pieces of text, questions, and anticipated answers entered in advance by the teacher. By contrast, [SCHOLAR] is based on the utilization of an information network of facts, concepts, and procedures; it can generate text, questions, and corresponding answers. Because [it] can also utilize its information network to answer questions formulated by the student, a mixed-initiative dialogue between student and computer is possible with questions and answers from both sides.[79]

A SCHOLAR dialogue[80] might look something like this:

> SCHOLAR: Use one of the following: Sucre, French, Argentina, Spanish, to answer the question: What is the language in Chile?
> SCHOLAR: You are taking too much time...
> Student: *Spanich*
> SCHOLAR: You should have written "Spanish." Very good.

Although by today's standards the dialogue (and especially its pedagogy) appears primitive, SCHOLAR is usually considered to be the

---

79 Carbonell, J. R. (1970). "AI in CAI: An Artificial-Intelligence Approach to Computer-Assisted Instruction." *IEEE Transactions on Man-Machine Systems* 11 (4): 190–202. https://doi.org/10.1109/TMMS.1970.299942
80 Carbonell, "AI in CAI," 192.

© Center for Curriculum Redesign – All Rights Reserved

first example of what came to be known as Intelligent Tutoring Systems, to which we turn next.

Applications of AI in Education

Intelligent Tutoring Systems

> Under the best learning conditions we can devise (tutoring), the average student is two sigma above the average control student taught under conventional group methods of instruction. The tutoring process demonstrates that most of the students do have the potential to reach this high level of learning. I believe an important task of research and instruction is to seek ways of accomplishing this under more practical and realistic conditions than the one-to-one tutoring, which is too costly for most societies to bear on a large scale. This is the 2 sigma problem.[81]

So-called intelligent tutoring systems (ITS) are among the most common applications of AI in education (in any case, as we have seen, they have probably been around the longest). Generally speaking, ITS provide step-by-step tutorials, individualized for each student, through topics in well-defined structured subjects such as mathematics or physics.[82] Drawing on expert knowledge about the subject and about pedagogy, and in response to individual student's misconceptions and successes, the system determines an optimal step-by-step pathway through the learning materials and activities. As the student proceeds, the system automatically adjusts the level of difficulty and provides hints or guidance, all of which aim to ensure that the student is able to learn the given topic effectively.

ITS come in many shapes, although typically they involve several AI models, an approach that we will unpack here. As we saw in our earlier discussion of AI technologies, AI models are highly simplified computational representations (in semantic networks, as used by

[81] Bloom, Benjamin S. (1984). 'The 2 Sigma problem: The search for methods of group instruction as effective as one-to-one tutoring." *Educational Researcher* 13 (6): 4. Note, however, that according to VanLehn, "human tutors are 0.79 sigma more effective than no tutoring and not the 2.0 sigma found in the Bloom (1984) studies" VanLehn, K. (2011.) "The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems." *Educational Psychologist* 46 (4): 209. https://doi.org/10.1080/00461520.2011.611369
[82] Alkhatlan, A. and Kalita, J. (2018). "Intelligent tutoring systems: A comprehensive historical survey with recent developments." ArXiv:1812.09628. http://arxiv.org/abs/1812.09628

© Center for Curriculum Redesign – All Rights Reserved

SCHOLAR, in ontologies,[83] or in knowledge graphs)[84] of specific knowledge about the real world (just like a model car is a simplified representation of a real car). The models used by ITS represent knowledge specific to teaching and learning. Typically, knowledge about the topic to be learned is represented in what is known as a domain model, knowledge about effective approaches to teaching is represented in a pedagogical model, and knowledge about the student is represented in a learner model.[85] The ITS algorithm draws on these three models in order to adapt a sequence of learning activities for each individual student. A fourth model found in some ITS is the open learner model, to which we will return later.

The Domain Model

A domain model represents knowledge about the subject that the ITS aims to help the students learn (much like the subject knowledge in a standard, non-educational, expert system). This might, for example, be knowledge about mathematical procedures, genetic inheritance, or the causes of World War I. In fact, over the years, mathematics for primary and secondary school students has dominated ITS. Mathematics, along with physics and computer science, are AIED's low-hanging fruits because they are, at least at a basic level, well-structured and clearly defined.

The Pedagogy Model

The ITS pedagogy model represents knowledge about effective approaches to teaching and learning that have been elicited from teaching experts and from research in the learning sciences (although it should be

[83] Ontologies are a way of representing a domain's concepts, data, components, entities and properties, and the relationships between them. Sowa, J.F. (1995). "Top-level ontological categories." *International Journal of Human-Computer Studies* 43 (5): 669–85. https://doi.org/10.1006/ijhc.1995.1068

[84] Knowledge graphs are an alternative approach to ontologies, https://ontotext.com/knowledgehub/fundamentals/what-is-a-knowledge-graph

[85] Luckin, R., et al. (2018). *Intelligence Unleashed. An Argument for AI in Education*, 18; Boulay, B. du., Poulovassilis, A., Holmes, W., and Mavrikis, M. (2018). "What does the research say about how artificial intelligence and big data can close the achievement gap?" 4. In Luckin, R. (ed.) (2018). *Enhancing Learning and Teaching with Technology*, 316–27. Institute of Education Press.

acknowledged that some ITS developers falsely assume that they have sufficient expertise in pedagogy).[86] Pedagogical knowledge that has been represented in many ITS include knowledge of instructional approaches,[87] the zone of proximal development,[88] interleaved practice,[89] cognitive load,[90] and formative feedback.[91] For example, a pedagogical model that implements Vygotsky's zone of proximal development will ensure that activities provided by the system to the student are neither too easy nor too challenging, one that implements individualized formative feedback will ensure that feedback is provided to the student whenever it might support the student's learning.

The Learner Model

As we have seen, some CAI effectively (although usually by another name) implemented versions of both domain and pedagogical models: knowledge of what was to be learned and knowledge of how to teach what was to be learned (for example, using linear or branching programmed instruction). However, what distinguishes AI-driven ITSs is that, as foreshadowed by Pask's SAKI, they also include a learner model: "a representation of the hypothesized knowledge state of the student."[92] In fact, many ITS incorporate a wide range of knowledge about the student—such as their interactions, material that has challenged the

---

[86] For example, many ITS set out to address student "learning styles" (Kumar, Amit, Ninni Singh, and Neelu Jyothi Ahuja. (2017). "Learning-styles based adaptive intelligent tutoring systems: Document analysis of articles published between 2001 and 2016." *International Journal of Cognitive Research in Science, Engineering and Education* 5 (2): 83–98. https://doi.org/10.5937/IJCRSEE1702083K This construct that has been widely discredited, e.g., Kirschner, P.A. (2017). "Stop propagating the learning styles myth." *Computers & Education* 106: 166–171. https://doi.org/10.1016/j.compedu.2016.12.006

[87] Bereiter, C. and Scardamalia, M. (1989). "Intentional learning as a goal of instruction." *Knowing, Learning, and Instruction: Essays in Honor of Robert Glaser*, 361–392.

[88] Vygotsky, *Mind in Society*, 86ff.

[89] Rohrer, D., and Taylor, K. (2007). "The shuffling of mathematics problems improves learning." *Instructional Science* 35 (6): 481–98. https://doi.org/10.1007/s11251-007-9015-8

[90] Mayer, R.E. and Moreno, R. (2003). "'Nine ways to reduce cognitive load in multimedia learning." *Educational Psychologist* 38 (1): 43–52.

[91] Shute, V.J. (2008). "Focus on formative feedback." *Review of Educational Research* 78 (1): 153–89. https://doi.org/10.3102/0034654307313795

[92] Self, J.A. (1974). "Student models in computer-aided instruction." *International Journal of Man–Machine Studies* 6 (2), 261–276. http://dx.doi.org/10.1016/S0020-7373(74)80005-2

© Center for Curriculum Redesign – All Rights Reserved

student, their misconceptions, and their emotional states while using the system—all of which can be used to inform what is being taught and how, together with what support needs to be provided and when. In fact, most ITSs go much further. The knowledge stored about the individual student is augmented with knowledge of all the students who have used the system so far, from which the system machine learns in order to predict which pedagogical approach and which domain knowledge is appropriate for any particular student at any specific stage of their learning. It is the learner model that enables ITS to be adaptive, and the machine learning that makes this adaptivity especially powerful.

A Typical ITS Architecture

The following figure shows how the domain, pedagogy, and learner models might be connected in a typical ITS.

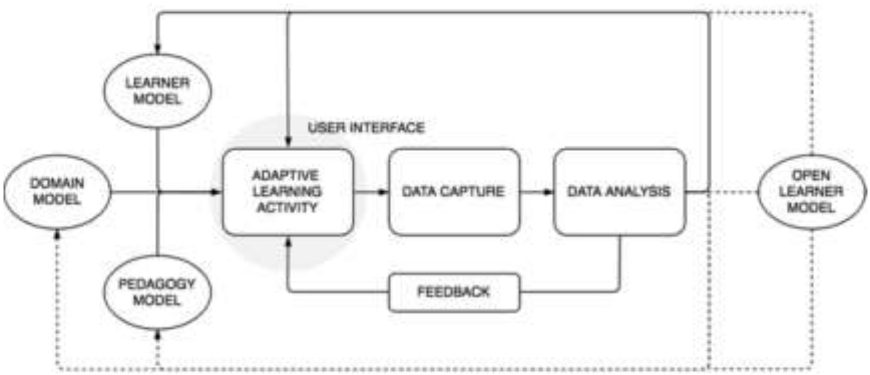

A typical ITS architecture, including the pedagogy, domain, learner, and open-learner models.

About CCR

Redesigning Education Standards

The Center for Curriculum Redesign (CCR) is an international convening body and research center seeking to expand humanity's potential and improve collective prosperity by redesigning K–12 education standards for the twenty-first century. In order to create a comprehensive set of frameworks, CCR brings together constituencies with diverse points of view—international organizations, jurisdictions, academic institutions, corporations, and nonprofit organizations including foundations—to consider and respond to the question: "*What* should students learn for the twenty-first century?"

The Center's Guiding Principles

A sustainable humanity—one in which collective potential is expanded, and collective prosperity improved—is orchestrated out of multiple social, economic, and environmental factors. Key among them: a relevant education, based on meaningful curriculum, is critical to creating sustainability, balance, and wellbeing.

While significant attention is being paid to teaching methods and pedagogy, the CCR argues that the *what* of K–12 education is at least as important as the *how*, and brings a singular focus to the what.

That twenty-first century what must take into account the accelerated pace of change we are experiencing, and shifts in societal and personal needs. Curriculum must be useful for the lives children will live and adapted accordingly.

Our ability to contribute a meaningful WHAT requires openness to different perspectives. Therefore, CCR avoids dogma and emphasizes innovation and synthesis—multiple inputs applied and organized for optimum clarity and impact.

We can—and will—shape the future we want.

© Center for Curriculum Redesign – All Rights Reserved

Focus on the What

Exponential changes in technology make specific predictions about the future all the more unreliable, but one thing is certain: we must prepare children to deal with greater complexity than ever before. The last major curriculum reform occurred in the late 1800s, also in a time of rapidly changing needs. Well into the twenty-first century, we can ill afford to depend on a nineteenth century curriculum. Indeed, we cannot expect our children to thrive unless we deeply examine, redesign and deliver a curriculum consistent with twenty-first century needs—one that is balanced and flexible. To thrive will mean to be adaptable, versatile and wise.

In designing a curriculum framework around adaptability, versatility and wisdom we accomplish two main goals:

- Enhance the chances of an individual's personal and professional success and fulfillment.
- Provide a common base of understanding and ability to participate in society, for a sustainable humanity.

The Center's Work

The Center for Curriculum Redesign is not a program or intervention. The staff and CCR's partners approach their work holistically, actively engaging with policymakers, standard setters, curriculum and assessment developers, school administrators, heads of schools, department heads, key teachers, EdTech experts and other thought leaders and influencers to develop a thorough understanding of the needs and challenges of all education stakeholders. This is essential to creating the vision of meaningful, relevant twenty-first century education, and to enabling practical implementation.

The organization's research, findings and recommendations are actively disseminated through a wide variety of formats: CCR-sponsored conferences and seminars, active web presence and social media, consulting engagements and keynoting.

The following video serves to summarize our views, and can be shared freely: http://bit.ly/CCRintrovideo[93]

CENTER FOR CURRICULUM REDESIGN

[93] For the video on Vimeo, go to http://bit.ly/CCRintrovideovimeo

© Center for Curriculum Redesign – All Rights Reserved