EDTECH MINDSET

YOUR MUST-HAVE EDUCATIONAL GUIDE TO THE FUTURE

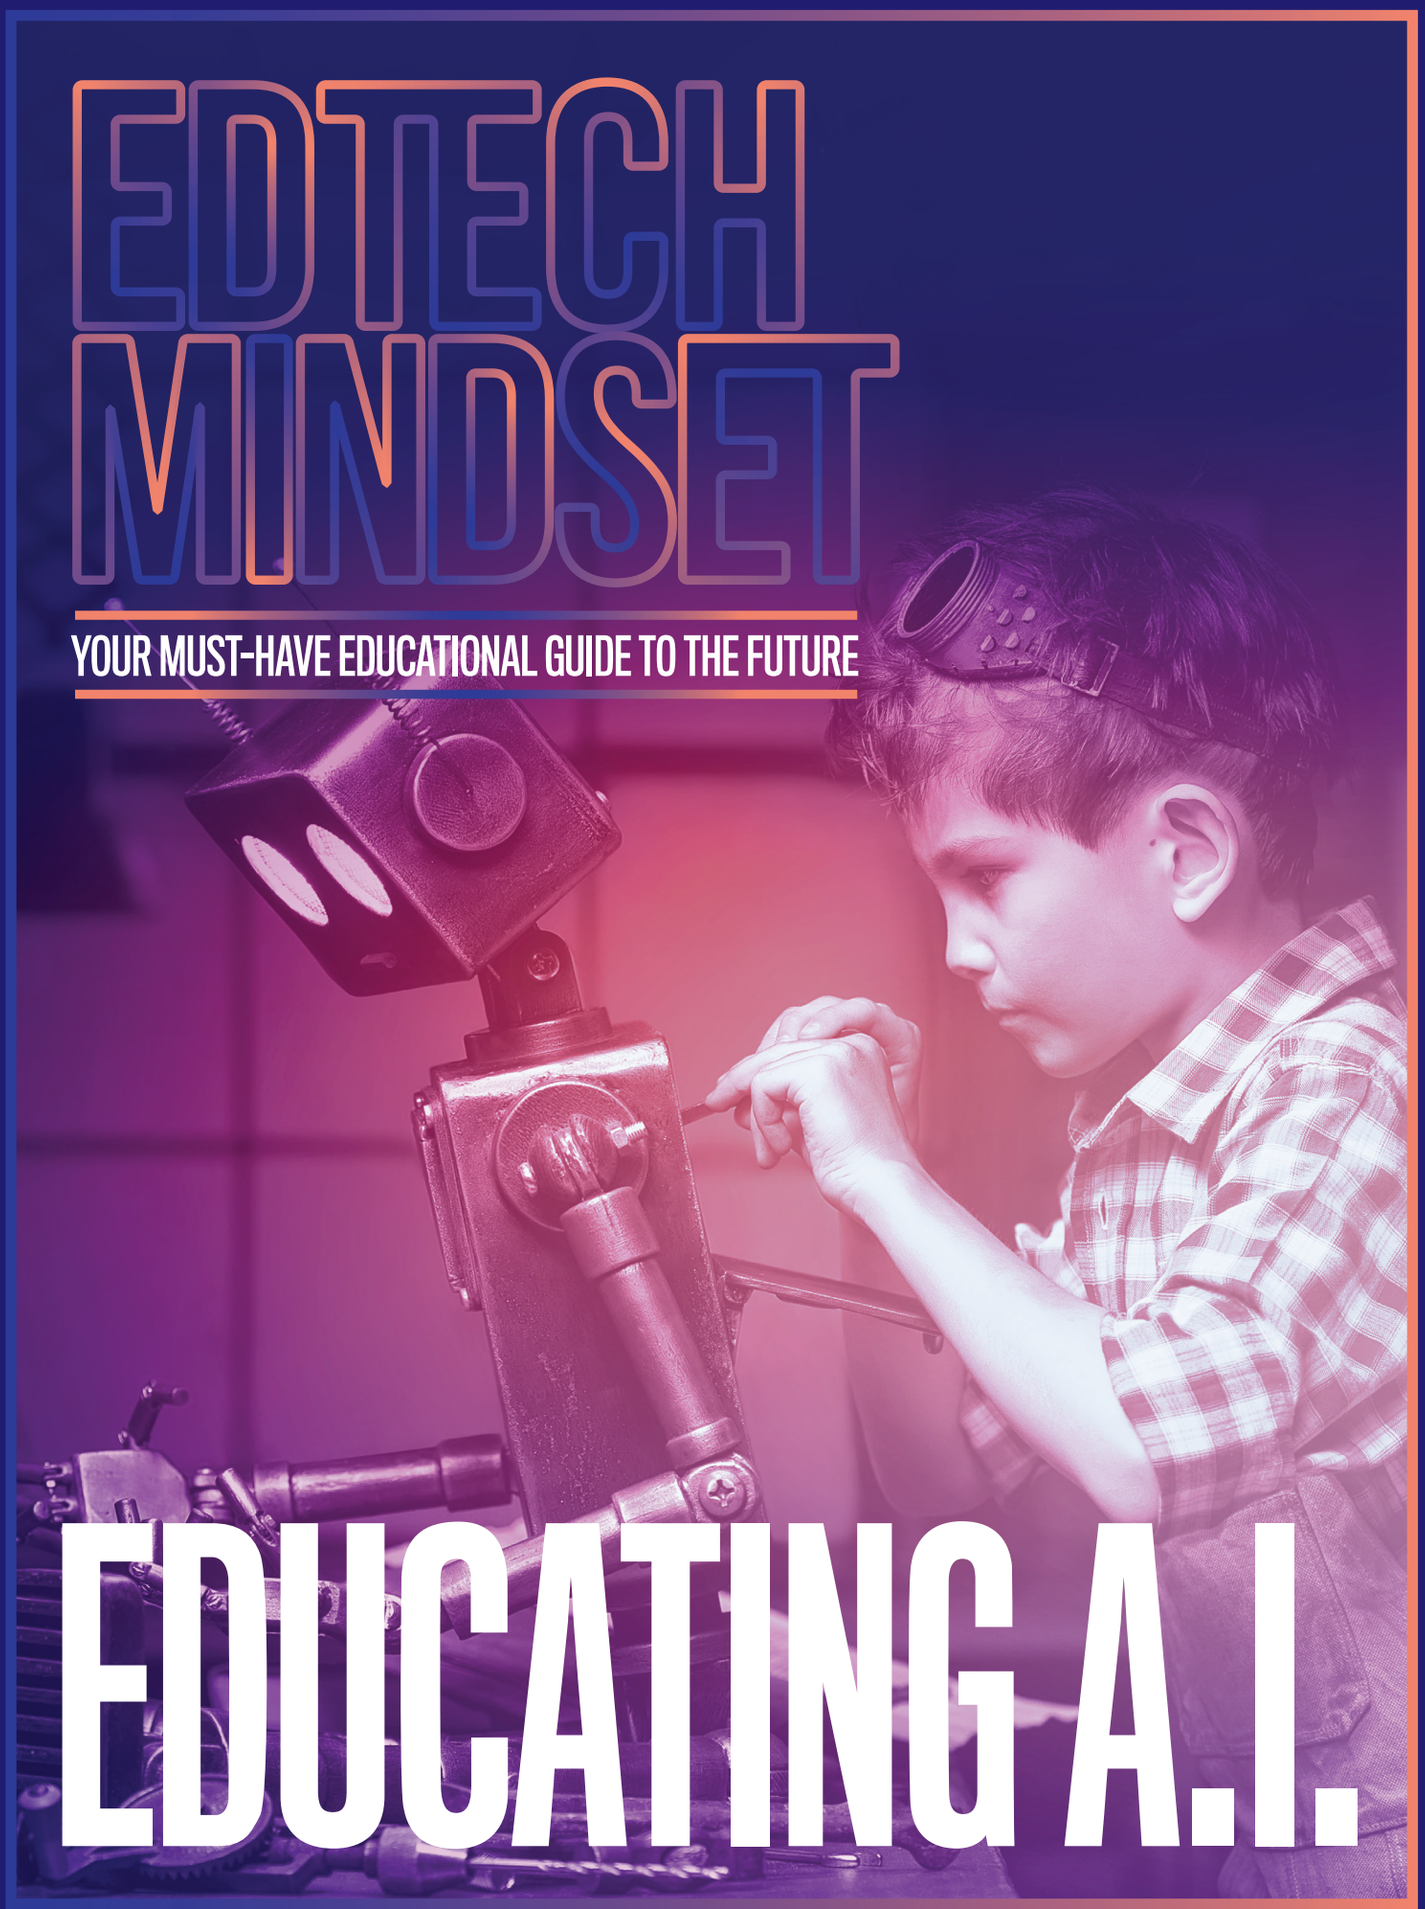

EDUCATING A.I.

JANUARY 2019 | SHAPING THE FUTURE 5 EDTECH SUMMIT

MindCet

CET • THE CENTER FOR EDUCATIONAL TECHNOLOGY

GESAwards

Global EdTech Startup Awards 2019

Join +2500 startups from around the world that are part of GESAwards community for a chance to be the most promising Edtech Startup of 2019

Apply at
www.globaledtechawards.org

CONTENT

Editor: **Cecilia Waismann** ceciliaw@cet.ac.il

Graphic Design: **Neriya Zur, Dov Abramson Studio**

EN Editing: **Nehama Uterman**

Translator: **Perry Zamek**

Content Contributors: **Yariv Adan, Maia Aron, Dr. Abe Asi, Cindy Chow, Dr. Orly Fuhrman, Orly Gan, Prof. Renee Hobbs, Ido Kenan, Guy Levi, Dr. Rami Puzis, Ben Lorica, Prof. Hagit Messer, Bob Rosenschein, Dr. David Weinberger, Dr. Cecilia Waismann, Avi Warshavsky**

Photo on cover: Shutterstock

THE EDITORIAL

DR. CECILIA WAISMANN, VP R&D MINDCET

"Last Friday, a portrait produced by artificial intelligence was hanging at Christie's New York opposite an Andy Warhol print and beside a bronze work by Roy Lichtenstein. On Thursday, it sold for well over double the price realized by both those pieces combined. 'Edmond de Belamy, from La Famille de Belamy' sold for $432,500 including fees, over 40 times Christie's initial estimate of $7,000-$10,000. The buyer was an anonymous phone bidder." (New York Times, Oct. 25, 2018)

Artificial Intelligence outcomes are definitely provoking amazement. There were many times during 2018 when headlines left us with our mouths wide open in awe: from self-driving cars to different art expressions, not to mention the constant, and even funny, media debate among our most prestigious thought leaders – is AI a monster or a savior?

On September 3, the largest annual EdTech Summit in Israel, Shaping the Future 5, decided to bring this debate to the education industry. A very impressive line-up of the latest developments was presented on stage to showcase potentially interesting pedagogical disruptions – from the technological giants such as Google, IBM, Microsoft, Amazon, or Intel, to Israeli startup companies at the forefront of international markets such as OrCam or Checkpoint. Leading educators discussed current challenges as well as hopes about the opportunities offered by AI developers' promises. This significant encounter, bringing together interested professionals from all over the world, tried to shed some light on an industry that is currently in crisis due to the tremendous difficulties it is experiencing in keeping up with the technological revolution.

This issue of EdTech Mindset "Educating AI" offers a few highlights from the Shaping the Future 5 EdTech Summit. These try to suggest to educators and startups to get on board in these magic times, to understand, to struggle, but overall, to take responsibility over the use of emerging technologies as an enabler that can narrow the gap between learners and the current unsatisfying educational system.

I hope you enjoy and feel inspired.

Photo by Shutterstock

AVI WARSHAVSKY, CEO OF MINDCET EDTECH INNOVATION CENTER

ARTIFICIAL INTELLIGENCE, SELF-IMAGE, AND THE MYTHOLOGY OF THE FUTURE OF EDUCATION

WHEN THE FISH DECIDES TO ROLL OVER

There is a well-known Talmudic tale that tells of a group of sailors who landed on a lonely island in the middle of the sea, lit a fire, and sat down to eat. A short time passed, and the island rolled over; it turned out that they had not landed on an island, but rather on the back of a giant fish, on which sand had accumulated and vegetation had grown. The public's interest with artificial intelligence is somewhat reminiscent of this picture. We talk at length about artificial intelligence and how it will influence reality – some of us with a messianic gleam in our eyes, and others with a look of terror. But artificial intelligence has long since arrived, and all of us are riding on its back (so to speak), without discerning its enormous influence on our lives. Google, Facebook and

AVI WARSHAVSKY
THE MYTHOLOGY OF THE FUTURE OF EDUCATION

IBM are, in effect, giant artificial intelligence factories, creating new knowledge on the basis of our vast use of the various services that they offer us. When we come to examine the influence of artificial intelligence on various areas of life, the major challenge is not only to imagine how the techniques of the future will look, but also to identify techniques that have all too fast become obvious to us. We very quickly become used to solutions that, in a more planned and systematic state of affairs, would have raised numerous question marks. If someone, twenty years ago, would have suggested that we give up the human skills of navigation, and instead rely blindly on software, such a suggestion would have generated debate. But in fact, from the moment that solutions such as Waze reached a critical mass of users, they swept all of us (apart from a handful of stubborn individuals) along on an enormous wave, without any reflection. The real debate does not begin, and cannot really begin, unless the fish decides to roll over.

When we talk of the encounter between information technologies and education in general, and the integration of artificial intelligence in education in particular, our tendency is to react, rather than taking a strategic view. We rush to ask questions about the way in which artificial intelligence can offer better tools for learning, and about the dangers that it brings with it, or about the implications of artificial intelligence for the labor market and the optimal way to prepare for it. All of these questions are important and productive, but artificial intelligence demands that we think on a more strategic level, where we outline our vision for our schools, and answer the big questions that direct it.

THE HIDDEN MYTHS

The word "myth" carries with it a lot of erroneous baggage. We may tend to see in a myth an undeveloped, perhaps even primitive, way of explaining reality. In popular terms, "myth" is often use synonymously with fantasy, whose key characteristic is its underlying lack of truth. A series of leading 20th century thinkers, ranging from Ernst Cassirer to Roland Barthes, taught us to relate more seriously to myths. According to Cassirer, the myth is another type of glasses through which we view reality, just like science, art, or religion. Through these glasses, so Cassirer claims, we reflect that which we cannot express by other means. Neil Postman devoted a significant portion of his book, **The End of Education**, to the important role of myths in the education system. Postman preferred not to use the word "myth" in this context, because of its problematic connotations; he spoke instead of narratives, or of "gods" that lead the education systems. Postman demonstrates how, in the Christian Middle Ages, it was the religious-ecclesiastic narrative that was the constitutive story of the education system, while the Enlightenment brought with it a scientific narrative, and from there on to the next narratives/gods – technology and the consumer culture. These major narratives, these myths, play an essential role in education systems, and answer the big question of "To what end?" – why are we in school, why is our school built the way it is, and what kind of world is it trying to prepare us for? Thus, for example, the function of the school, according to the Thomas Jefferson narrative, is to ensure that the citizenry should know when and how to defend their freedom, while the Protestant ethic, for example, wants the school to teach us that we need to stick to hard work and develop our ability to delay gratification. We are also familiar with "smaller" narratives, such as that in which we learn arithmetic so as not to be cheated at the grocery store, or Talmud so as to sharpen our minds. In the complex, multicultural reality of the present day, there is no single founding myth for our education systems, other than what is commonly referred to as "popular education," nourished by an eclectic fabric of beliefs, which may not even be consistent, were we to apply a sterile, academic

> WE DIDN'T JUST GET UP ONE MORNING AND DISCOVER, TO OUR SURPRISE, A TECHNOLOGY THAT WAS "SMARTER" THAN US. WE HAD LOOKED FORWARD TO THAT MOMENT, WE DREAMED OF IT, AND WE ADVANCED TOWARD IT WITH OUR EYES WIDE OPEN.

analysis to them. In spite of their lack of coherence, and in spite of the fact that these myths are often not formulated expressly, they play a key role in directing the world of education. Artificial intelligence may have an enormous influence on a broad spectrum of educational narratives and myths, and in this article we will focus on one important, central myth – the myth of explanation. In order to understand how artificial intelligence influences the myth of explanation, we should recall the dramatic moments when an artificial intelligence machine bested the most gifted of human players.

MACHINE BEATS MAN: THE MYTH THAT WAS WRITTEN BACKWARDS

A chess player sits before a chessboard, beads of perspiration dotting his brow and running down his neck; he takes out a handkerchief to wipe them away. The audience around is spellbound, anxiously following every move. The surprising aspect of this picture is actually the other side of the table – the chair opposite the player is empty; he is playing against a machine which, were it not for the tendency toward the dramatic, might have been represented by an almost invisible box. Our player is still struggling along, but the audience already knows that his loss is a foregone conclusion. This picture is familiar to us from the loss by the Korean Go champion, Lee Sedol, to Google's AlphaGo in 2016, and the loss by the world's leading chess master, Gary Kasparov, to IBM's Deep Blue in 1997. The picture is, in fact, a lot older. The Turk was a mechanical chess player that played in Europe in the 18th and 19th centuries, defeating celebrated chess players and notable statesmen such as Benjamin Franklin and Napoleon Bonaparte. The Turk, however,

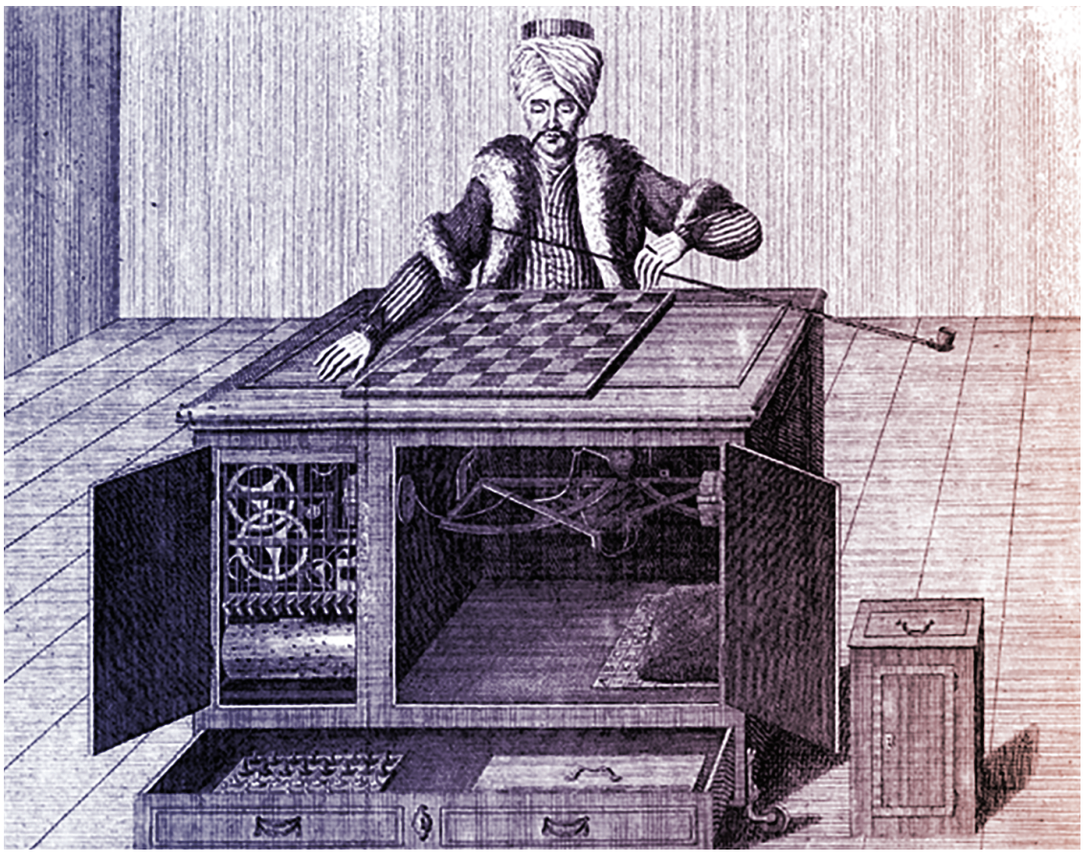

A copper engraving of the Turk, a fake chess-playing machine constructed in the late 18th century.

was a fraud – secreted within the machine was a diminutive man who made the actual moves, while the audience thought it was the machine. However, the story of the Turk is not merely an amusing old-time tale of deception. The story shows that the image of a man-made machine, whose performance exceeds that of a human being, predates the situation in which such a capability exists, and demonstrates just how interested we are in such a story. After all, we didn't just get up one morning and discover, to our surprise, a technology that was "smarter" than us. We had looked forward to that moment, we dreamed of it, and we advanced toward it with our eyes wide open. One of the fundamental papers in computer science was written by Claude Shannon in 1950, and it dealt with the possibility of a chess game between a man and a machine. Shannon wrote his paper in the years when computers were taking their first steps, in the same year that Alan Turing formulated what would later be known as the Turing test, and at a time when the computing power of the enormous computers that existed in those days was smaller than that of the most negligible of apps on the phone in our pocket. However, this did not stop Shannon from being sufficiently visionary to be fascinated by the idea of a competition between man and machine. Such a competition was mythical, and we are used to myths that hark back to the past. However, myth of artificial intelligence is one that was written backwards – it was a myth of the future. And like every myth, it involved drama, and was the story of a struggle. It is not for nothing that the picture that immediately comes to mind when we hear the words "artificial intelligence" is that of a chess match between man and machine, which is the image of competition; it is not for nothing that we also seek a deterministic aspect in this myth, which usually ends with man's loss, as in the ancient depictions of Greek tragedy. But this is a good point to stop and ask: Why did we so much want the machines to defeat us? From where does this deep sentiment, that creates this myth, come? What part of our humanity does this myth perpetuate? In order to touch on this question, we should look at later manifestations of this myth – the moments in which this victory actually took place.

FROM BRUTE-FORCE ARTIFICIAL INTELLIGENCE TO INTUITION

We will not go into a complicated discussion of the exact definition of intelligence in general, and that of intelligence in the context of artificial intelligence. For the purpose of discussion, we will make use of a narrow, somewhat imprecise but reasonably practical definition, under which intelligence is the ability to perform tasks. In this sense a cat has a higher intelligence than a spider, a chimpanzee has a higher intelligence than a cat, and man has a higher intelligence than a chimpanzee. The good news about artificial intelligence-based programs is that they have a higher intelligence than man. It is important to note that, based on the definition that we have adopted, we are talking about the measurement of task performance – a program that is capable of defeating a man at chess is better than a human being in performing this task, but this tells us nothing about other mental qualities that it may or may not have. For the purposes of this discussion, it may be able to defeat the world champion at chess, yet still be as a sensitive as a block of wood, or profound as a bowl of whipped cream.

MYTH OF ARTIFICIAL INTELLIGENCE IS ONE THAT WAS WRITTEN BACKWARDS - IT WAS A MYTH OF THE FUTURE

In his article, Shannon attempted to characterize, on the theoretical level, the path that would be followed, in the future, in chess games between man and machine. Shannon distinguished between two types of possible victory by the computer over the human being:

Type A, also known as brute-force artificial intelligence, is based on an algorithm that traverses all the possible states of a chess game, and tries all of them, until victory is achieved. A chess game has about 300 billion possibilities in only the first four moves, most of them not particu-

larly successful. In other words, a computer that wins using a type A strategy will have to consider an enormous number of states within a very short time.

Type B, which is more sophisticated, is able to focus on smaller number of successful moves, and achieve victory through them. This ability might be referred to, by a somewhat inaccurate analogy, as intuition.

MACHINE BEATS MAN: A PLAY IN TWO ACTS

Type B play requires much greater sophistication, yet Shannon assumed that victory over a human chess player would take place specifically through artificial intelligence of type B, rather than type A, mainly because type A requires enormous computing power, the likes of which did not seem, during the 1950s, to be achievable in the foreseeable future.

THE AGE OF THE SMART MACHINES PLACES US, FOR THE FIRST TIME, OPPOSITE EFFECTIVE MACHINES WHO LOGIC WE DO NOT UNDERSTAND.

Shannon erred in his prediction – in 1997, almost fifty years following the writing of the article, the world's champion chess player, Gary Kasparov, was beaten by IBM's Deep Blue. Deep Blue's victory was a type A victory – the computer was sufficiently fast and powerful to preview all possible moves, and to choose the most appropriate one. It certainly wasn't the most sophisticated program; or, as Kasparov himself put it, it was intelligent in the same way that an alarm clock set to ring at a particular time is intelligent.

Almost twenty years later, however, a type B victory was also achieved. Google's AlphaGo defeated Lee Sedol, world champion Go player.

Go is a traditional Chinese game with an enormous number of possible moves, more than the number of atoms in the whole of the universe, and this was a much greater challenge than chess.

For AlphaGo to be able to play this complex game, it learned from about 160,000 games, and practiced more than three million board positions, many more than a human being could dream of grasping. But AlphaGo was not the final stage in the story. The next version, Alpha Zero, taught itself to play Go without its learning being based on anyone teaching it. Alpha Zero took three days to learn the game, following which it won 100 out of 100 games.

WHO UNDERSTANDS ARTIFICIAL INTELLIGENCE?

The victories by AlphaGo and Alpha Zero reveal an impressively broad range of aspects of machine learning and artificial intelligence, but the most astonishing phenomenon in these victories lies, as shown by internet thinker David Weinberger, in our inability to explain how they won. We know how machine learning operates, but we are unable to rationalize or recreate the specific learning process. AlphaGo offered moves that no human player had ever made. For every task, up to the age of artificial intelligence, we have been able to define certain regularities, which allow us to create a technology that addressed the task. In many instances the technology was more effective than us; often it was fearsome in its power, but we always understood the logic and the consistency underlying its actions. Till now, technology served to amplify the human body, but the rules and the models that were the basis of its activity, and its logic, were totally human. The age of the smart machines places us, for the first time, opposite effective machines who logic we do not understand.

THE MYTH OF EXPLANATION

The victories by AlphaGo and Alpha Zero might serve to challenge one of the more fundamental myths of the world of education and learning – the myth of explanation.

Kurt Vonnegut expressed this myth in lyrical terms in his book, Cat's Cradle:

"Tiger got to hunt,
Bird got to fly;

THE MYTHOLOGY OF THE FUTURE OF EDUCATION

Illustration of Nicolaus Copernicus

Man got to sit and wonder, 'Why, why, why?'
"Tiger got to sleep,
Bird got to land;
Man got to tell himself he understand."

Vonnegut's words should be read very carefully, from beginning to end:

Just as it is the nature of the tiger to hunt, and the bird to fly, so it is man's nature to sit and ask "Why?" – to seek explanations for the reality that surrounds him. The act of seeking explanations disturbs our rest, while finding an explanation is a source of calm. Just as the tiger sleeps and the bird lands, so too man finds rest when he tells himself that he has understood. This doesn't mean that he has actually found the ultimate explanation, only that he has reached a subjective state in which feels he has understood. This idea is inherent in the Hebrew language – we seek an explanation that מניח את הדעת (settles our mind), a place in which our restless consciousness can rest. That place is the explanation.

Explanations generally show how a specific occurrence is subject to general rules. Explanations have mechanisms for justification, reasoning, proof, and theories that support them. We would like to see education systems, among others, as a place for explanations – a space which teaches us to seek explanations, presents us with convincing rationales that are comprehensible to us, and primarily delineates boundaries to the question of what is a satisfactory explanation. As with the giant fish in the Talmudic tale, the myth of the explanation is so deeply entrenched within our culture, until we barely notice it. The entry of artificial intelligence into our world is one of those moments in which the fish rolls over, and we see in a new light that which we had erroneously seen, till that moment, as stable land.

ARTIFICIAL INTELLIGENCE AND SELF-IMAGE

Artificial intelligence's challenge to the institution of explanation is, first and foremost, a challenge to our self-image. As Weinberger shows, since the time of Plato, and especially since the Age of Enlightenment, our ability to perform tasks and achieve goals went hand in hand with our understanding. We always had the ability, po-

ARTIFICIAL INTELLIGENCE CHALLENGES THE BOUNDARIES OF OUR UNDERSTANDING, IN AN INESCAPABLE WAY. ONE MAY COMPARE THIS SHAKEUP TO TWO EARLIER SCIENTIFIC REVOLUTIONS: THE COPERNICAN REVOLUTION, AND THE DARWINIAN REVOLUTION.

tentially at least, to understand what works, what is effective. Artificial intelligence challenges the boundaries of our understanding, in an inescapable way. One may compare this shakeup to two earlier scientific revolutions: the Copernican revolution, and the Darwinian revolution. Copernicus taught us that the planet on which we live is not the center of the universe, and that, like other planets, it circles the sun. This is not just an important astronomic discovery, but also – and foremost – a revolutionary humanistic discovery, since it removes man from the center of the universe to one of its back rows. It is a discovery that had a dramatic influence on our collective self-image. In a similar way, the processes that Darwin uncovered, and the broad context that his theory offered for the development of life, are not just theories in biology, but also theories that reposition the human species among the rest of the creatures, all of us – according to Darwin – having developed through the identical principle, from simpler to more complex life forms. Darwin's Man is not the pinnacle of creation, but rather another link in an impressive, yet blind, symphony of natural selection. In both instances, the strong vocal opposition to these theories was often not based on purely scientific grounds, but on the shake-up and the attack on our self-image as human beings.

ARTIFICIAL INTELLIGENCE AS A COPERNICAN REVOLUTION

One might shake one's head with patronizing scorn at the primitive nature of these objections, but this would overlook the pain that they reflect – a pain caused by the fundamental change in our place in the world. The difficulty in accepting these theories may be compared to the difficulty experienced by a baby, as it grows, in recognizing the existence of other people, separate entities in the world around him. It is not surprising that we still find an astonishing percentage of graduates of Western education systems who believe that the sun revolves around the earth, or bitter objections to the teaching of evolution in schools in the western worlds, almost two hundred years after Darwin. It is almost certain that many of those who, consciously or unconsciously, object to these theories, understand that they are not just professional, scientific theories. A scientific theory such as evolution is accompanied by a whole fabric of ideological teachings which, for example, may generate ecological sensitivity on the one hand, or cruelly deterministic approaches on the other. Artificial intelligence is no different in these senses. Copernicus taught us to be humble in the face of the universe, Darwin taught us to be humble in the face of other living creatures, and artificial intelligence teaches us a lesson in humility in the face of the devices that we ourselves construct. These understandings of the human species require that we reformulate the fundamental myths of our schools. It sounds like a somewhat distant, philosophical mission, of little urgency, but the fact that we are neglecting the great narratives of education in favor of pinpoint responses, may incur for us a heavy price. The growing alienation lies less in our ability to amuse and arouse curiosity in the students, and more in our difficulty to provide convincing answers to questions of "For what purpose," and to adapt our pedagogic strategy and the major narratives to the world in which they operate. Students who will live in a world in which most of their tasks are performed by smart machines, can no longer rely on educational narratives that are full of holes, whose key aspects have not been tested since the Age of Enlightenment. This opportunity, to become partners in composing our world's new educational myths, is given into the hands of all those involved in the educational enterprise – parents, teachers, and policy makers alike.

Q&A

WITH DR. **BEN LORICA**

CHIEF DATA SCIENTIST AT O'REILLY MEDIA, PROGRAM DIRECTOR OF STRATA DATA AND ARTIFICIAL INTELLIGENCE CONFERENCES.

Q With all the notions about AI that are going around – of a catastrophic threat that will create displacement and overthrow humans and, on the other hand, that it is an enabler and an enhancement – which would you choose in talking about AI and humans today?

B.L. "I think at this point, one can think of AI as an assistant to humans. For one, many of the things we read about AI tend to be about systems that are limited to one task; for instance, pattern recognition, or perception – detecting whether or not this moving object is a pedestrian, or a bicycle, or a car, if you're in a self-driving car. And I think over time, as the systems become imbued with knowledge and reasoning capability, they can do more. But already, actually, even those types of technologies can help to partially automate many routine tasks and workflows. So I would say, at this point, the best metaphor is an assistant, and that assistant will become smarter over time. One example is chatbots. Last year, or maybe two years ago, there was much hype about chatbots, and it turned out they were easy to build, but also they were very limited in what they could do. A lot of them were rule based, based on state machines. So if the user starts interacting with a chatbot, and goes off in a direction that the chatbot isn't expecting, the chatbot gets stuck. As the underlying building block technologies for chatbots

get better – that would be in natural language understanding – one can expect some of these solutions to get better. But I think, for the foreseeable future, a lot of them will be much more focused, domain specific; they will not be offering general intelligence, but will be an assistant for a very specific role inside a company or an organization. And then things get better over time as the building block technologies for that solution get better."

Q Let's talk about personal assistants. How do you see the impact of these, especially in education? What would be the effect you foresee of using personal assistants? Is it a way of outsourcing human relations to more automated ones?

B.L. "You can imagine them getting better over time, and expanding in scope over time. So, just as the self-driving car industry has settled on five levels of autonomy, with level five being the true self-driving car, I think these chatbots will have different levels too. Maybe at the most routine level it just notifies the student, 'Hey, your homework is due. Hey, you should read this for tomorrow.' And then maybe at the second level, it will start answering routine questions – if you're familiar with the acronym FAQ, the frequently asked questions in tech support. And then it gets better and better over time every day, and can start handling contextual questions – given that there are multiple possibilities to answer or frame a problem, which one would you choose? And then over time, it gets even more personalized to you as the student. I think of it as an evolution of the technology. I like the notion that as the technology gets better, it can do more, but it will always kind of evolve in terms of capability. We can't get around the fact that these technologies rely on basic building blocks. So in the example of the chatbot, that would be natural language understanding."

AS THE TECHNOLOGY GETS BETTER, IT CAN DO MORE, BUT IT WILL ALWAYS KIND OF EVOLVE IN TERMS OF CAPABILITY. WE CAN'T GET AROUND THE FACT THAT THESE TECHNOLOGIES RELY ON BASIC BUILDING BLOCKS.

Q AlphaGo, two years ago, really surprised us with the capacity of an intelligent system to take its own decisions, to think coherently, or to use common sense. How do you see the implications of such developments? Which areas of our lives do you think they will affect the most?

B.L. "AlphaGo is an impressive achievement, but it's limited to a game with very well-defined rules. There's a lot of computation that they had to use in order to get to that level. So, the question is, what other tasks do we have where we can afford to throw that much computation to automate something, and secondly, where the rules of the game are so well defined? I think maybe there are certain tasks inside a company – the phrase people use is enterprise workflow automation – there might be a series of tasks that are somewhat repeatable, confined, and well defined, and with enough simulations and examples, you can automate them. The question at the end of the day is: One, do you have enough data? Two, do you have the scale to justify automation? Because if you don't have the scale, if you only have to do something a few times a week, then there's no point automating it. But if you have the scale, you have to start looking at the problem, and you see if it fits into the framework of the technologies we have today. So AlphaGo, for one, relied on a mix of underlying technologies that may or may not apply to the problem that you have."

Q So in a way, even very intelligent systems do have limitations, or still depend on the capacity that we have in feeding the system. So would you say there are core human learning aspects that a machine would not be able to develop?

B.L. "I think right now, the systems we have rely on a lot more data than humans, so you and I can look at one or two examples of something and internalize that pattern. We also rely on prior knowledge and domain knowledge. So we know that we know when we enter a situation; we know certain laws of physics, we know that

something can't just disappear, right? So I think that right now, we are in a situation where our systems are good if we have a lot of data and a lot of compute. One interesting example is language and natural language. Deep learning is a great approach, and it has proven to be very successful in computer vision and speech recognition. It has had some success in natural language, although it hasn't led to natural language understanding. If you talk to the people who work in computational linguistics, they're all using deep learning, but they also feel that deep learning is producing models that are not the most efficient. Because linguists come with a lot of prior knowledge from linguistics, they want models that are much more efficient, require less data, take advantage of linguistic rules and patterns, and things like that. I am hoping that we'll come up with hybrid solutions where deep learning is one part of the answer, but there are other techniques that take advantage of prior knowledge and similar things."

Q Experience, you would say.

B.L. "Yeah, domain expertise. Understand some prior structure."

Q In an article you talk about the promise and pitfalls of AI and deep learning – what would you say is the main pitfall today?

B.L. "I think right now, one is that it requires a lot of data. Two, it's a bit of a black box. Some of the more famous talks towards the end of last year were around people being frustrated with understanding how deep learning works. Let's say you're a deep learning expert in a company and someone joins the team, and they don't have enough understanding to have a way to pass on a lot of the knowledge, other than you need to get your hands dirty and try things at this point. I think that's getting better. One of the good signs over the last year is that the people who work in theoretical computer science have gone into machine learning and are trying to understand how machine learning works, and when it fails and when it excels".

Q When people talk about the use of AI in education, the main focus is adaptive and personalized learning. And as you've mentioned about other areas, the frustration here is even more critical, especially in adaptive learning. With the current trend of technology personalizing everything, is there a realistic hope of AI being used in education? What are the implications of that? There is a lot of value in being in a class with a group of people, there's a lot of learning. Do you think we will adapt and learn how to deal with negative and positive effects, or do you see a danger there?

B.L. "I don't know if I would say that it's either-or. I think it will definitely be an important part of the picture. I think part of it is just a recognition that people learn differently, so there should be some amount of personalization involved. People learn in different ways at different speeds, and respond to different styles of teaching, so I think that from that perspective, personalization should help education. As to your other question, there's more to education than just rote learning; there's interaction with your peers, developing social skills, developing emotional intelligence. So I think that there's room for both, depending on how an institution actually deploys these technologies. If they deploy it in a way where you don't have to come to school anymore, you all just stay at home, then I think maybe that might be going too far, particularly for people in a certain age bracket. I think though, arguably, that if you're already an adult learner, you're already working and you want to enhance your skills, and you want to take courses or get credentialed, and you want to do it in the comfort of your home, technology can help. But I don't know if we're at that point where you can just turn it over to an AI; you still need instruction. As we mentioned before, these personal assistants are not there yet. An important

part of education, people forget, is networking. So even for adult learners, you can imagine learning the material using this great AI system at home that can help you learn at your own pace, using the style you like. But one of the things you want to be able to get out of education is to meet people who share your interests and can become part of your professional network over time."

THE QUESTION IS, WHAT OTHER TASKS DO WE HAVE WHERE WE CAN AFFORD TO THROW THAT MUCH COMPUTATION TO AUTOMATE SOMETHING

Q One last question, related to a new buzz, which is the ethical implications of AI. There is a huge trend to use AI across industries and society. For example, using AI to help judges take decisions. Would you prefer, if you could choose, an emotional human decision with all of the dangers that ensue, or to pass this over to an objective machine? And should we be concerned about this trend?

B.L. "I think I would respond by not responding. Again, it's not an either-or situation. On the one hand, I think the black box AI system is unemotional and can take things so-called objectively, but time and again, there are examples of AI systems that have exhibited bias or haven't been exactly fair. And actually, this is an area of personal interest to me; I've been doing a lot of reading and studying in this area. The machine learning community has over time established certain metrics or statistical tools for ensuring that an AI system is fair. But each of those metrics has problems. So there are exceptions to each of those metrics. A classic example is what they call NP classification, which doesn't use variables that are protected, such as gender, age, and race. But then people point out that maybe there are some situations where you'll have to use those variables, because the distribution of women might be different than the distribution of men. So if you apply the same rule, which is 'if above a certain level, we decided this way,' but the two populations have different distributions, you might be actually inadvertently penalizing the women. If you look at each of these statistical metrics, there are exceptions. The main takeaway is that we're still at the point where the machine learning community is developing these tools. And the main thing that I tell people is that if you're serious about ethics and fairness, then there's no substitute for having to get in there. You can't rely on a statistical metric and statistical procedures to make your ethical dilemma go away. Because for one, there are also a lot of papers coming out now which say that even though you have a statistical procedure that you deem ethical, there might be impacts that are delayed over time that make it less ethical. In other words, humans are in the loop, humans are still involved. You will have to set up processes in place where you take the best of the statistical advice and procedure for creating ethical AI, but make sure you have teams of data scientists who can audit and make sure that the AI is behaving accordingly. And actually, one of the things that I've come around to is this notion of risk management, in general, for machine learning and AI. Now that we're deploying many of these systems in mission critical, real world applications, there are many considerations beyond statistical machine learning and business metrics: fairness, ethics, privacy, security. All of these come with risks. Just as we want software and financial services that are risk free, we want AI that's also risk free, so we need to start thinking in terms of risk management for AI, which might mean, in this particular case of privacy and ethics, having in the team that builds your AI system a team of data scientists, on one hand, and then an independent team of data scientists who serve as validators, so after you build the model this team that wasn't involved in the model building process will independently validate your AI system to make sure that it is fair and unbiased."


$\min_G \max_D \mathbb{E}_x \left[\log(D(x)) \right]$

DR. **CECILIA WAISMANN,** VP R&D MINDCET

A NEW WORLD ORDER THAT IS NOT NEW FOR ALL

The illustration is "Edmond de Belamy", a generative adversarial network (artificial intelligence unsupervised machine learning) portrait painting constructed in 2018 by Paris-based arts-collective Obvious.

2018

"Elon Musk warns A.I. could create an 'immortal dictator from which we can never escape'" (CNBC, Nov. 2018)**; "Reddit Co-founder Mocks Elon Musk's Warnings About AI: But then he issued a warning of his own"** (Futurism, Dec. 2018); **"How we can prepare for catastrophically dangerous AI and why we can't wait"** (Gizmodo, Dec. 2018); **"Bill Gates: A.I. can be our friend"** (CNBC, Feb. 2018).

The year 2018 has witnessed a continuous debate over alarming calls for a public "scare" about the latest developments of Artificial Intelligence (AI). A significant number of headlines, especially in non-scientific media, were clearly semantically structured to call for our dismay and apprehension. The big question is whether this is a true civic concern to raise public awareness, an irresponsible abuse of the power of media regarding concerns taken out of context, or a real concern of the growing incertitude about our future. In any case, AI has become a center of attention, provoking skeptical reactions towards the constant and never-ending technological marvelous developments.

However, I believe that the true reaction to such headlines relies on our adoption of smart systems, on our seamless and inadvertent integration of emerging technologies in our daily habits, leaving those headlines to our conceptual and moral debates. Moreover, a stronger reaction we currently observe is the younger generations'

expression of a confident know-how, a natural behavior towards the ongoing technological revolution, accompanied by a lack of awe, and instead the suggestion of a new world order that is not new for them! We are witnessing the birth of a much more inclusive natural environment. Let us try to look a bit closer into it.

HUMANS' AND MACHINES' SYMBIOTIC PROCESS DESPITE APOCALYPTIC SCARES

Emerging technologies are quickly becoming more intelligent and helping bring machines and humans closer together. A significant milestone happened in 2016, during the Google project "AlphaGo," wherein a computer program played against Lee Sedol, the 18-time world champion in "Go," which is considered the most complex and intellectually challenging board game in existence. AlphaGo initiated moves that were unthinkable for a human, leaving everyone in total bewilderment. AlphaGo's moves provoked a state of expectancy that made programmers react to the game-play with the emotional and surprised reactions of mere observers. This event exemplified a stage of smart machines' development where they are not only able to process non-human amounts of data at non-human speed, but to behave with a non-human and maybe unpredictable intelligence.

Mixed reactions emerged. On one hand, more significant technological developments took over many industries; on the other hand, there was a burst of scary predictions even from original advocates of AI. The image of this apocalyptic Terminator, a super-intelligent monster that will take over jobs, lives, and even the world, started to take root in the public perception. Alongside that came the idea of a disconnection from our relevance and power over it, almost as if we had no responsibility for these developments, but were instead only passive user-victims. We forget sometimes that many of our modern conveniences were also initially met with such fear and dissent, and long later, acknowledged responsibility. Any technological development must be expected to have a double-edged-sword impact. We can take as an example the human use of plastic (a major development of the 20th century) which took half a century to truly revolutionize industries, and ended up helping inhibit bacterial spread and lower mortality rates as well as creating new human habits. At the same time, humans' irresponsible and exacerbated use of plastic contributed, and still does, to the destruction of earth's environment, quickly becoming a danger for our planet's ecosystem and survival of species. Does that make plastic a monster? Plastic itself isn't inherently evil, but our use of it has become dangerous, so who is at fault? To what extent is AI development any different from other technologies? Maybe the answer lies in the "intelligence" element – our descriptive differentiation from the other beings of our ecosystem. AI has so many basic commonalities with humans that it may require a new definition of "being."

EMERGING TECHNOLOGIES ARE QUICKLY BECOMING MORE INTELLIGENT AND HELPING BRING MACHINES AND HUMANS CLOSER TOGETHER

A MUCH MORE INCLUSIVE NATURAL ENVIRONMENT

The generational gaps we currently face are only of a few years, causing significant differences in the perception of the world around us. The youngest generations have incorporated as a natural part of the environment even the most sophisticated technologies, and have even become immune to novelties. Some explain it by the uncontrolled overload of information they learned to digest; others explain it by their navigation outreach accessibility; others by the constant and overriding virtual connections which displace (or place anywhere) us all. For whatever reason, smart machines are a natural part of the current landscape, which the unstoppable development race is re-defining as a dynamic, ever-changing, unpredictable landscape.

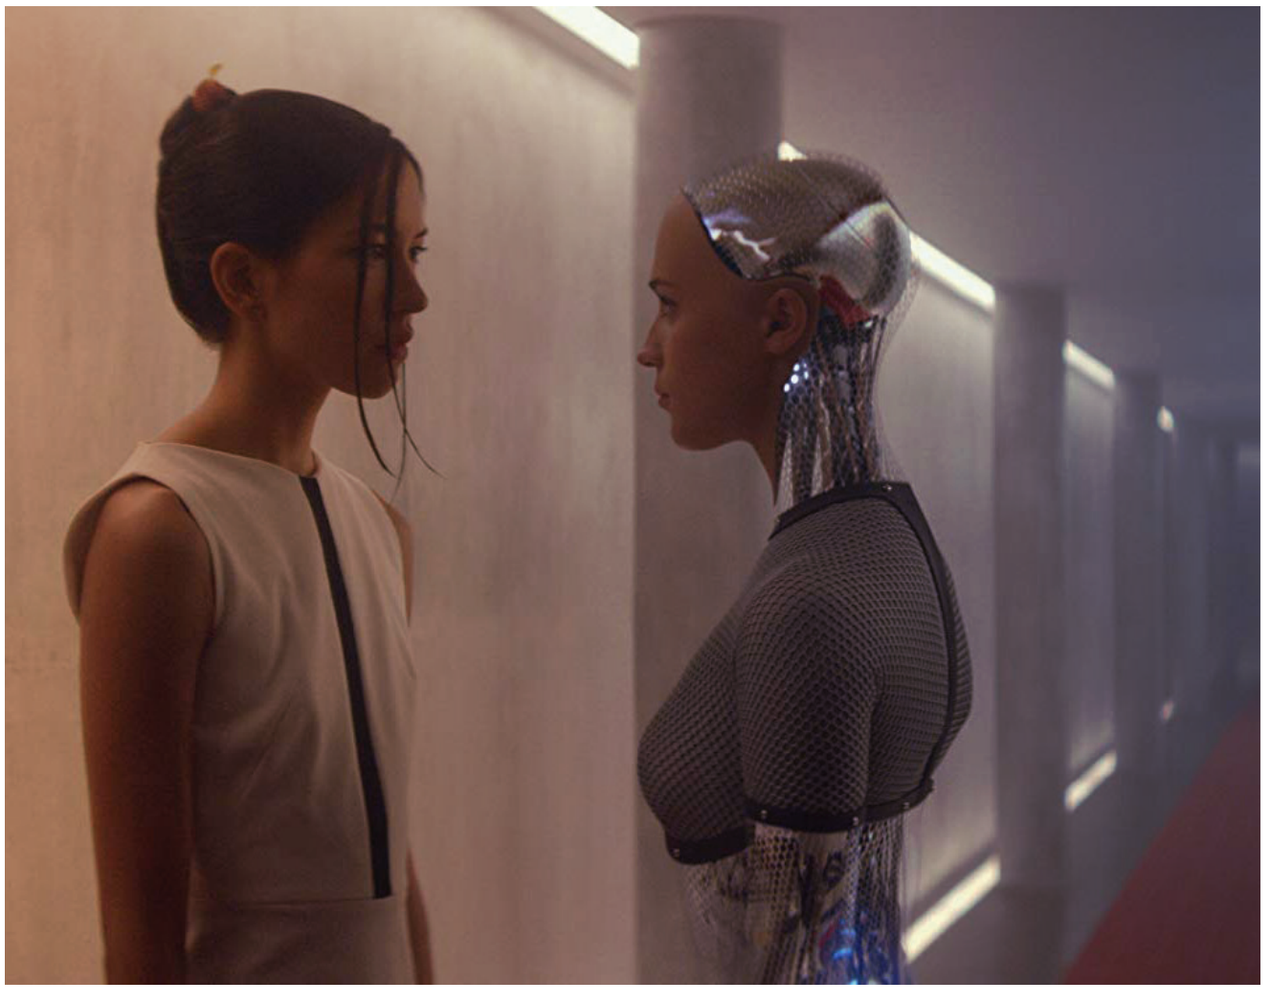

Alicia Vikander and Sonoya Mizuno in Ex Machina (2014) © 2015 - Universal Pictures International

ARE WE READY TO "EDUCATE" THE CURRENT GENERATION?

In such a scenario, it is no surprise to see educational systems in crisis, standing on a shaky ground without a clear vision of how and what role they play. Educators perceive a growing generation gap almost as something debilitating!

THE GENERATIONAL GAPS WE CURRENTLY FACE ARE ONLY OF A FEW YEARS, CAUSING SIGNIFICANT DIFFERENCES IN THE PERCEPTION OF THE WORLD AROUND US

Of all industries, educational systems probably face the biggest generational gap due to their incapacity to understand and adapt to the new landscape, while at the same time they still hold a major and defining role in the younger generations' lives. The challenges are many and the solutions not clearly defined. Artificial intelligence is an enabler that can help us change and improve our educational systems. As Ray Kurzweil argues, the melding of humans and machines as a result of the singularity and the growth of AI will be a significant enabler: "As machines become more intelligent, humanity will also grow to become smarter."

We, educators, should be aware of the risks and opportunities of AI and help develop a generation to responsibly use technology. Smart systems as enablers of new educational solutions could narrow generational gaps, and definitely help educational systems meet the needs of a generation that is currently very desensitized and frustrated with the existing ones. This generation is abandoning educational systems in favor of the much more attractive and rich Internet, which, it could be argued, is currently becoming a parallel educational system. Is that the future educators wish to see?

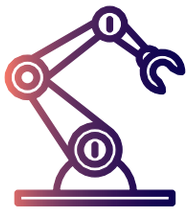

IDO KENAN, JOURNALIST

ARE BIONIC-PEOPLE BECOMING A REALITY?

WHEN AI ENHANCES HUMAN CAPABILITIES AND PROVIDES LOST LIFE OPPORTUNITIES

MobileEye is a driving assistant that uses AI to watch the road for what drivers might have missed, identify dangers and alert the drivers. Could that mechanism be improved to assist people in their lives? OrCam thinks so. In 2015, the company launched **MyEye**, a wearable camera with a speaker which, according to OrCam's site, is "designed to assist people who are blind, visually impaired, or have a reading disabilities [sic]" by recognizing faces, currency and products and conveying that knowledge to its users, as well as reading aloud texts from varied surfaces. Unsurprisingly, OrCam and MobileEye share their co-founders, Prof. Amnon Shashua and Ziv Aviram.

OCR (optical character recognition) existed before, but it was cumbersome and took time to process. Dr. Yonatan Wexler, OrCam's Executive

AI IS A FANTASTIC TOOL. IT CAN REDUCE ROAD ACCIDENTS, MAKE DISEASE DISCOVERY FASTER AND ALLOW BLIND PEOPLE TO READ

VP R&D, says AI has made it possible for MyEye, which is about the size of a finger, to read texts aloud in real time. "We have made the act of reading a technicality," he says, disagreeing with critics who claim that machines reading to us are ruining reading for us: "We were taught that reading is the ability to understand shapes on a piece of paper. But that's like saying work is sitting in front of a computer. Technically it is, but it's not the work. Reading is the ability to understand the material. Until recently, understanding the material depended on the ability to understand shapes on a paper. If you were dyslexic or hard of seeing, you couldn't do it. I'm dyslexic, and a prefer audiobooks to print books any day of the week. I'm not stupid, but by the time I put in the time to read what's written, I don't have any energy left to understand what it says."

EyeCam isn't just complemental, but supplemental. It enhances its users' senses and capabilities with AI, which in a sense, mind the pun, noninvasively turns them into demi-bionic people. "With language and text analysis and nimble understanding capabilities, we can have the hard of seeing get things even seeing people find hard to get," says Wexler, illustrating with a feasible future feature: "When a user looks at a menu in a restaurant, he could tell the EyeCam 'I want something dairy,' 'I feel like a hamburger,' or 'I'm vegetarian,' and it'll only read you the relevant dishes – faster than a regular person, who has to go through the entire menu."

While AI is a mainstay of EyeCam, Wexler, only half-jokingly, casts doubt on its very definition: "This domain of AI is not well defined – we know what A is, but what is the I? There's no definitive definition of intelligence. It's not a distinct thing. Let's take the Western Wall – it's pretty intelligent, it lasted 2,000 years. Maybe it's more intelligent than humans who live a 100 years. We're not even close to fulfilling our human capabilities. AI is another step in that direction. The tractor helped, too. Those are tools that help humanity reach farther, and we need to understand how to leverage it. If a kid can't understand shapes on paper, it's a shame. You just lost this kid. He'll tell himself he's stupid. But if you assist him, you haven't replaced his brain – it's just a better way to hand him the information. And if he's smart, he'll take that information and be even more successful."

SHOULD WE AS HUMANITY FULLY EMBRACE AI, OR HEDGE ITS POTENTIAL DANGERS?

"Once things are easy to do, they need to be dealt with cleverly. AI is a fantastic tool. It can reduce road accidents, make disease discovery faster, cheaper and more accurate, and allow blind people to read books better than some sighted people. Such achievements undoubtably propel humanity forward and should be embraced. The faster we understand the added value of technology, the better it is."

Wexler is optimistic when it comes to humans as well: "Machines can be built to keep working. Some of their work can be very fast. The combination allows machines to process more data that is possible by humans. The processing itself can be quite simple. We all know that computers can do basic math faster and with less mistakes than humans. General inference, and the ability to understand situations, are still very far off. We humans are amazing and we should never forget that. We are smarter than anything we create."

Photo by Franki Chamaki on Unsplash

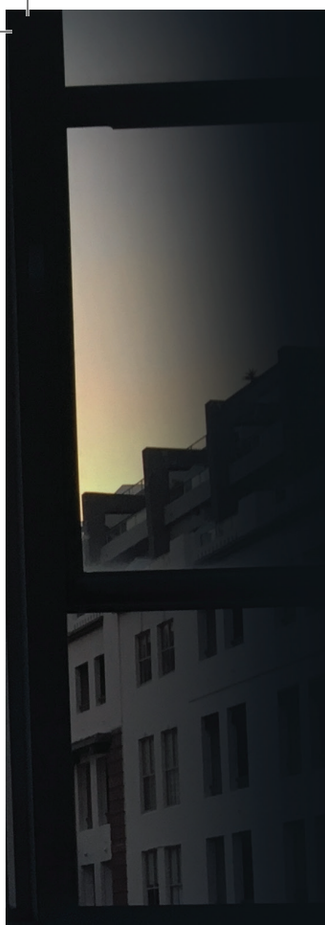

DR. DAVID WEINBERGER, SENIOR RESEARCHER AT HARVARD'S BERKMAN KLEIN CENTER FOR INTERNET & SOCIETY

MACHINE LEARNING WILL CHANGE THE POWER OF DYNAMICS IN EDUCATION!

Dr. David Weinberger is senior researcher at Harvard's Berkman Klein Center for Internet & Society, co-director of the Harvard Library Innovation Lab, a philosophy professor, journalist, strategic marketing consultant to high tech companies, Internet entrepreneur, advisor to several presidential campaigns, and a Franklin Fellow at the US State Department. Most of all, Dr. Weinberger is a thought leader who helps us make sense of all that the new world order technology is leading us to.

In a conversation with EdTech Mindset, Dr. Weinberger expressed his fascination for developments in AI and its consequences for humans' understanding of the world, at the same time pointing out the crucial role of humans in deciding how and where to take these developments. "It's crucial that existing democratic processes, not commercial interests, determine how artificial intelligence systems are optimized," Weinberger explains in an article published at Wired (Jan. 2018). Moreover, Weinberger shows how machine learning is significantly changing the power dynamics of knowledge, from an absolute to a much more shared "we are in this together" distribution.

"I am really, really excited about the prospects of machine learning and its effect on education in two sorts of ways, one of which is the way in which a technology can help scholars find information, make sense of information, help teachers and students personalize, discover weaknesses and strengths. I'm sure that it will introduce a whole set of bias into the system. But also, I think it will take a whole bunch of bias out as well, and we need to pay attention to that, of course. But the thing that has me most

DR. DAVID WEINBERGER
MACHINE LEARNING WILL CHANGE THE POWER OF DYNAMICS IN EDUCATION!

excited about machine learning is what I hope will be the effect of its model of how the world works on how we think about how the world works, in a couple of ways. One is that all of the outcomes of machine learning are probabilistic, and this can give a confidence level. You can often set confidence levels for the project in which you are engaging with machine learning. And the idea of confidence levels as metadata that we attach to our assertions is something I hope we can learn from machine learning. Machine learning will get us used to this idea – when it pronounces that there is a 0.76 confidence level that it's going to rain tomorrow, or that you have the risk of some disease or whatever. Getting used to hearing assertions with confidence levels attached to them, I think, is a tremendously important lesson for students to learn – for the entire culture to learn, but for students to understand as well. It fundamentally changes the power dynamic, for one thing. One of the ways to change the power dynamic is by helping the continuing move away from teachers' authority – with teachers conveying absolute knowledge, which, you know, nobody believes in at this point, but still there's a sense of that, just from the body language of a classroom – to more of a sense of 'we're in this together, we have reasons to believe this or that at some level of confidence, with some set of reasons.'

"But – and this is actually the second thing that I hope you learn from machine learning's model of the world – we are in an incredibly complex world, a chaotic world, that is so far beyond the tiny speck of matter that we call our brain and our capacity to understand it, that the best we can do is to work together to understand that the world overwhelms us, that we never achieve complete certainty, but that we can together still make our way through this world. And if either or both of those characteristics of machine learning's model of the world, the models that it builds for itself – amazingly complex and detailed models in which the contingency of one piece of data on another may be difficult to find in itself, and the outcomes may rest upon tens of thousands of variables that are interacting in ways that surpass human understanding. We get that sense and the sense that all that we do comes with some level of confidence or lack of confidence. And I think the nature of the educational project changes for the better, because I think that machine learning's model of the world is actually more accurate, truer, than the one that we humans tend to come up with."

THE BEST WE CAN DO IS TO WORK TOGETHER TO UNDERSTAND THAT THE WORLD OVERWHELMS US, THAT WE NEVER ACHIEVE COMPLETE CERTAINTY, BUT THAT WE CAN TOGETHER STILL MAKE OUR WAY THROUGH THIS WORLD.

AI, WHERE TO?
TALKING WITH INDUSTRY LEADERS

DR. ABED ASI

DATA SCIENTIST, MICROSOFT ISRAEL

What do you do at Microsoft? "We are developing a platform for analyzing video using AI technologies. Using Microsoft video indexer, we extract insights like faces, sentiment, brands, and many more. There are several target business scenarios that such a platform can empower. One of them is search engines, but it also can empower recommendation engines, moderation, accessibility, and several other business scenarios."

When using AI and big data to try to understand the world around us, is there risk of bias? "Oh, yeah, this is a critical point. And I think what we, as technologists, should give to the educators a platform, a technology. Educators are the ones who should prioritize their requirements from technology, and we as technologists should enable this through technology, and we should not fix these priorities for educators."

So you're saying that engineers and scientists are not the people to deal with the ethical issues in software, but their clients are? "Exactly. It's not healthy that one side takes responsibility for all the ethical issues related to AI. Conferences like [Shaping the Future 5] are very important for brainstorming, for bringing sensitivities from all of the sides who are developing and are involved in AI-based models."

AI, WHERE TO?
TALKING WITH INDUSTRY LEADERS

PROF. HAGIT MESSER

PROFESSOR OF ELECTRICAL ENGINEERING AT TEL AVIV UNIVERSITY

How do you explain the latest wave in the debate of the ethical issues around AI? "Technology has always been a way to help human beings enhance their abilities and skills. However, while traditionally it supports physical skills or routine intellectual skills (e.g., computing), AI technology enables us to delegate moral decisions to machines. This challenges legal issues, ethical questions, as well a philosophical question about the boundaries between human beings and machines."

What would you say are, if any, the main ethical issues that should be raised about AI use for educational purposes, and why? "One of the main ethical issues of AI is the inherent bias in its decisions. People also are biased, but unlike machines they have the ability to fight against it by education and values."

Is there a concern? If yes, who should be concerned? "I am not concerned about the technology, but about its use. Even the term 'artificial intelligence' is misleading and dangerous. Machines are not intelligent. They are able to do things, including doing complex tasks and making complex decisions, but not as human beings. Cars are used for transportation, and they take us from one place to another fast and safely, but cars are not 'artificial legs.' By referring to machines as 'artificial humans,' people misuse and misinterpret this technology. As cars need gas or other source of energy, because otherwise they won't run, AI needs data, which is a resource and should be treated as such. The use of this resource should be regulated."

CINDY CHOW

EXECUTIVE DIRECTOR, ALIBABA ENTREP FUND

What does your fund invest in? "We invest in startups with a Hong Kong nexus, so it can either be a homegrown startup from Hong Kong, or startups from all over the world that find Hong Kong a good place to be a launchpad for them to expand in Asia."

How do you see your responsibility on how AI is developed and used? "I guess like with everything invented for the benefit of humans, we have to use it wisely. Just like Dr. Waismann [at Shaping the Future 5] mentioned about the use of plastic, so I guess AI, as a technology, will definitely be helping people to learn better, or to be able to bring out the potential of human beings, but we have to think of a way to make sure that it will not be abused when we use it."

What unique, interesting AI technologies would you be interested to invest in? "Something that can, especially for Asian children, revolutionize how we learn in school because in Asia, or more specifically in Hong Kong and China, we are more into tests and exams. So something that can really change how we learn and inspire more children to think more creatively is what we'll be very interested to get to know and invest in."

Have you seen any interesting companies in this area in Israel? "I did, and I'm still trying to see what are the benefits of bringing them over to Hong Kong and also to China. I hope that they can open branches and then have some meaningful operations in Hong Kong, because the key thing I want to do is to bring in talents to Hong Kong. So if they can have a separate development team to localize the technology, then it will be beneficial to the Hong Kong ecosystem because, of course, they will also try to be able to share the knowledge with the local community, and that's something I really want to work towards."

YARIV ADAN

LEADING THE PRODUCT AND ENGINEERING TEAM THAT IS BUILDING THE GOOGLE ASSISTANT, GOOGLE ZURICH

What specifically is your team building? "Our team is responsible for many of the capabilities of intelligence of the Google Assistant. Being at Google you never build all the intelligence, you're not the only one, and you're doing a little bit more, always. My team focuses mostly on some of what we call the capabilities of intelligence of the Google Assistant, which means understanding natural language, understanding context, being able to be proactive, being able to interpret visual input."

What interesting, unique, groundbreaking capabilities does AI enable? "I think it's a paradigm shift in how humans-machines interact. We're basically shifting from an era where humans had to learn UI – people went to courses to learn Windows, Docs, these kinds of things. We're moving towards an era where computers are learning our UI, they're adapting to us. This basically unlocks the number of people that would suddenly be able to see it in the use cases that we'll be able to use, so I expect that this will cause a huge change in how we're going to use computers."

So we're not going to see keyboards and screens in the near future. "I'm not sure. As you know, radio did not replace everything that was there before it, DVD did not replace radio, the telephone did not replace all the methods before it, the internet did not cancel the real world, mobile did not completely replace desktop, and so forth. I think we're adding but differently. Every time we see such a revolution, the main and new usages are re-defined. I strongly believe that in this kind of natural interaction – of you interacting with the machine the same way you interact with other people, when its presence is part of the ambient – you don't perceive the device, it's just there. I think this will become the main use case."

What you're saying is that not only will we not have to type, we won't even necessarily have to talk – sometimes the machine will know what we need and what we want and give it to us. "Yes, I expect that to happen soon – you ask about the weather, and perhaps traffic information in the morning; in the evening maybe there is certain music and media that you like to consume, types of news and maybe some stock-related stuff; when you go in your car or on your bike or on the train, if you have another routine. Some of these are fairly easy to recognize but soon they will also help you with recommendations on personal, leisure, or work-related issues. Yeah, I think that's not anymore science fiction at this point".

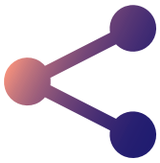

DR. RAMI PUZIS, CYBER SECURITY RESEARCH CENTER AND SOFTWARE AND INFORMATION SYSTEMS ENGINEERING DEPARTMENT, BEN-GURION UNIVERSITY OF THE NEGEV

NAVIGATING THE SOCIAL NETWORKS' STORMY WATERS IN THE ERA OF AI

AI or, more accurately, machine learning, is one of the truly ubiquitous technologies. While driving a modern car, using a smartphone, or searching the internet we rely on machine learning, most of the time without even noticing it. Even the fifth generation of cellular networks (5G) heavily depends on machine learning for continuous fault-free operation. Machine learning is not black magic; it includes techniques and methods of applying statistical analysis to derive insightful predictions.

One of the domains where machine learning is extremely prevalent is online social networks. Consider, for example, the friends suggestion mechanism in Facebook or in other social networking platforms. Statistical link prediction

models, learned from the vast amount of data available to the owner of the social networking platform, predict which social network profiles belong to people who are most likely to be your friends in real life. How can anyone know whether or not John Doe is my friend in real life? Well, if we have many common friends on Facebook, live in the same town, studied in the same college, or just both of us were tagged on the same photos on multiple occasions, then most probably we are acquainted. These are just few out of the many clues used by machine learning algorithms to suggest friends.

Researchers also used link prediction to identify intruders or so-called "unwanted friends" on Facebook. It turned out to be an effective tool for detection of pedophiles and criminals that use social networks to connect with their victims.

MACHINE LEARNING IS NOT BLACK MAGIC; IT INCLUDES TECHNIQUES AND METHODS OF APPLYING STATISTICAL ANALYSIS TO DERIVE INSIGHTFUL PREDICTIONS

Similar techniques were used in the past to detect automated accounts in social networks that spread spam or promote content. Armies of these so-called social-bots are there for hire by anyone who wishes to pay $15 for 5,000 "likes." The social capital industry has grown tremendously, now offering posts, comments, friends, and followers in a variety of social networking platforms and even **24/7 customer support**. Type "buy likes" or "buy retweets" in google to get the right impression. The price for such promotion services depends on the ease with which the social networking platforms detect and ban such in genuine behavior. If the fake accounts that offer social capital for sale can fool the machine learning algorithms trying to detect them then their price goes up.

How do they do it? How can one create hundreds or thousands of social-bots that look like human accounts or generate comments that seem like genuine human response? Of course, using machine learning! Generative Adversarial Networks (GAN) are one example of neural network-based approaches that can generate photos and descriptions for the profile page of the fake account that seem pretty genuine to a human as well as to other machine learning algorithms. Link prediction or similar machine learning approaches can be used to optimize the wiring of the fake accounts within the social network. Machine learning operated chat-bots are used to post comments. Essentially, modern online social networks are war arenas where artificial intelligence fights artificial intelligence.

The stakes go up when techniques similar to those described above are used by political organizations, terrorists, and governments. Whether spreading propaganda or fake news, statistical analysis helps to reach the right audience and create a larger impact. Machine learning also helps hacker groups to collect intelligence about the target organizations and plan and execute the attacks. Sophisticated automated or human-powered social network accounts are wired into the social network of the organization's employees to collect private information and disseminate malware. On the opposite side of the barricade, machine learning algorithms are used to detect fake news and fake accounts, manage social network honeypots, identify the attacker's intention, and more.

So where are we, the humans, in this never-ending war of artificial intelligence against artificial intelligence? We are right in the eye of the storm. On one hand we try to keep safe by identifying fake profiles and fake news, and on the other hand we are a captive audience to the attention-grabbing game of the platforms which want us to do something we did not intend or believe in, something that is not (entirely) true. Fake news and fake social network accounts are there to stay. Machine learning algorithms rely on our behavior to learn and interpret their inputs. We should help the algorithm to interpret our behavior correctly. We can do it by ignoring hoaxes even if they are funny or align with our own vision of the world; carefully inspect friend requests, approving only those coming from people we personally know; and of course never (take candies from strangers) click on links or open files coming from unreliable source. modern online social networks are war arenas where artificial intelligence fights artificial intelligence

Photo by tim de groot on Unsplash

PROF. RENEE HOBBS, DIRECTOR OF MEDIA EDUCATION LAB, UNIVERSITY RHODE ISLAND

MACHINE LEARNING COULD ENABLE A NEW UNDERSTANDING OF HUMAN LEARNING

A conversation with Prof. Renee Hobbs, an internationally recognized authority on media literacy education. Through community and global service and as a researcher, teacher, advocate and media professional, Hobbs has worked to advance the quality of digital and media literacy education in the United States and around the world. She is the founder and director of the Media Education Lab, whose mission is to improve the quality of media literacy education through research and community service.

In a conversation with EdTech Mindset, Prof. Hobbs expressed her concerns about decision makers' choice of promising personalized learning at the cost of students' communication and collaboration skills. Prof. Hobbs' significant knowledge and expertise helps us understand the effects of the latest technological developments on students' basic literacy skills and competencies and bring out the opportunities and challenges we should be aware of.

AI & HUMANS

"You asked what metaphors could best be used to explain the relationship between AI and humans. And I thought a lot about that, of course. Typically, we think of AI as being a tutor, a kind of customized teacher who offers you just what you need. Some people have called AI a kind of crutch, like an assistant, like a walker that supports your learning and enables you to maybe execute some skills with assistance and support. These days, I'm starting to think about AI as a bit of a maze. As you know, a maze is a form of play, it's a way to guide users and delight users and entertain users. And it's also a structure that controls behavior. So for me, AI is a kind of a maze that might be used for guidance, might be used for entertainment, might be used for learning, but we should never forget that it's created as a way to guide or control behavior.

THE SAME WAY THAT HUMAN FREEDOM IS ENABLED THROUGH LEARNING, MACHINES THAT ARE FREE TO LEARN INEVITABLY RAISE THE ISSUE OF MACHINES' FREEDOM TO DO EVIL OR TO DO HARM AS WELL AS TO DO GOOD

MACHINES TAKING DECISIONS

In relation to the idea that machines can truly make decisions, as shown by the AlphaGo program, I think it is important to distinguish between the first generation of AI, which was controlled by humans, and the second form of AI, or machine learning (ML), where machines are unsupervised – they engage in learning independently, on their own. That (second form) offers a lot of opportunities for education to customize learning and to provide targeted instruction. Machine learning can be really powerful. It also might be useful in helping us develop strategies for managing information overload. It might be helpful for machine learning to help us manage the increasing volume of information and emails and entertainment choices that confront us every day. It may also further erode people's ability to engage in certain cognitive tasks, or it might increase our dependency on the assistance provided by machines. So machine learning is powerful, it has some potential, great affordances, and certainly some liabilities. I am looking forward to machine learning that supports my ability to manage information overload, because I think that's one of the biggest problems we have right now in our society.

AI & LEARNING

Your interest in what AI can teach us about learning and whether there is an aspect of human learning that cannot be automated by machines – these are great questions. I've been thinking about what AI actually teaches us. Well, one thing we understand is that AI and machine learning can be used for friendly or aggressive purposes. It can be used as our friend and it can be used as our enemy. I just saw **The Incredibles 2**, that great Pixar movie. It reminded me that in 2004, when **The Incredibles** came out, the first version, the enemy was Omnidroid – it was a bad machine that learned and developed on its own. In some ways, this idea of machine learning enabling machines to work unsupervised leads us to starting to understand the idea of what it means to be independent or truly free. The same way that human freedom is enabled through learning, machines that are free to learn inevitably raise the issue of machines' freedom to do evil or to do harm as well as to do good. Now, one of the ways that machines learn is by organizing and classifying information. I think that machines are helping us understand the importance of organizing and classifying information. As machines learn to read pictures, and to identify the content of pictures, this helps us understand about how definitions and classification shape the way human learning works. And machine learning is also teaching us a lot about feedback which, after all, is what learn-

ing is. We change our behavior based on feedback. We adapt to feedback. In fact, it's possible that machine learning can give us a whole new way of thinking about feedback. And that might help us rethink how we understand the concept of assessment, rather than something that is a value judgment, that is a defining characteristic of a person's identity. Is it a mark of shame? No! Assessment is a form of feedback and feedback promotes learning. I always say to my students, when you get a lot of feedback, that's a great compliment. The instructor is giving you a compliment when you get a lot of feedback, because that's helping you learn more. I think one of the things AI teaches us that's of most concern to me is this idea that human behavior is predictable and can be predicted. Machine learning certainly helps us understand the many ways in which human behavior is predictable, but that phenomenon has some dangerous implications for education and for society. Predicting behavior based on algorithms could exacerbate the problem of ability tracking. And it has larger negative implications for the role of social determinism and free will, in shaping how we understand the potentiality of a human person. And so I'm concerned about the way that machine learning may get us to think in mechanistic and narrow ways about human potential, and I think that's a concern that we should always keep in the forefront of our minds.

A GAME CHANGER OR AN AMPLIFIER?

Is AI a game changer or just an amplifier? I believe it is an amplifier; it will amplify a focus on assessment, and it will focus on and amplify tracking behavior, monitoring and surveilling behavior. Those aspects of education will become increasingly more important as a result of AI. And that, of course, gives me concern because of my own interest in developing learners' autonomy, their freedom to make their

MACHINE LEARNING CAN GIVE US A WHOLE NEW WAY OF THINKING ABOUT FEEDBACK, AND THAT MIGHT HELP US RETHINK HOW WE UNDERSTAND THE CONCEPT OF ASSESSMENT

Prof. Renee Hobbs at Shaping the Future 4

own judgments, their own interpretations of media messages and of informational content, and their creative abilities to synthesize information in new and original ways. After all, that's where innovation comes from. So I'm concerned primarily about the amplifying function of assessment and tracking or surveillance because of the way it will become a business model that will be compelling to entrepreneurs; because big data is increasing and will continue to increase in value, and approaches to teaching and learning that rely on developing empathy, social emotional skills, and even creativity don't have the same market value as the easily commodifiable function of big data.

PERSONALIZED LEARNING

One of the things I'm most concerned about, and I write about in an upcoming issue of the **Journal of Adolescent & Adult Literacy** with my colleague, Julie Coiro, is how digital learning might be redefined in ways that emphasize the primacy of personalization at the cost of emphasizing critical media literacy, the development of students' communication and collaboration skills and the development of youth voice and empowerment. So my concern about how digital learning might get redefined is that IT administrators, when they think about using technology and education, instantly proceed to a focus on using machines to personalize learning, instead of thinking about how machines can help amplify the student voice or promote the development of effective citizens in democratic societies, or even develop students' creativity and global understanding. That's a big concern for me. After all, virtual exchanges and digital learning technologies have the opportunity to increase our ability to genuinely understand people far removed from us in space. And yet, we're not talking about that, we're not talking about how digital learning can promote global understanding. That seems to me a terrible gap and omission that future educational technologists need to address.

DEVELOPING EMPATHY, SOCIAL EMOTIONAL SKILLS, AND EVEN CREATIVITY DON'T HAVE THE SAME MARKET VALUE AS THE EASILY COMMODIFIABLE FUNCTION OF BIG DATA.

CHATBOTS

One thing that's really clear is that not only are machines learning to talk to us, but we are adapting to talk to them and like them, and that's something media literacy educators want us to reflect on – what's happening in our own behavior that is enabling us to change to become more like machines. And that's something to be concerned about. It's something to reflect on as we think about how our own behaviors are implicated in this digital landscape. I think one of the things that's most interesting about chatbots, that leaves me a little bit on the fence, is what we're learning about the importance of exposure to dialogue and conversation during early childhood. You know, the data on that is unmistakable. There are big gaps between children who grow up in well-educated families where, on average, between the ages of zero and three a young child will hear 2,000 or more words per hour, as opposed to a kid growing up in poverty, who might only hear 600 words per hour. Longitudinal research has shown that the amount of exposure to turn taking and conversation in early childhood is the single best predictor of reading skills at grade three. So encouraging language, turn-taking language, language that invites children to elaborate and to explain and to reason and to discuss – that is essentially the foundation for the development of literacy competencies. I'm not sure if machines are up to the task of providing that kind of support for learning, but I like the idea of strategies that help to reduce the gap between children who are growing up where there's lots of exposure to that rich language and children who grow up in families where there's not so much exposure to that. Perhaps chatbots can provide a way to reduce the inequality that is now rampant in our communities, in our countries and around the world.

Photo by matthew henry on Unsplash

CYBERSECURITY AS A GATEWAY TO USER SAFETY

A CONVERSATION WITH **ORLI GAN**, HEAD OF THREAT PREVENTION PRODUCTS AT CHECKPOINT

While we ponder the potential dangers of artificial intelligence, AI is already selflessly protecting us from concrete dangers, such as cyberattacks. But make no mistake: AI isn't merely making cybersecurity easier, quicker, and better – it's what renders cybersecurity possible at all, explains **Orli Gan**, head of Threat Prevention Products at Checkpoint: "AI is a means to being able to address the scope of the problem in a way that human beings simply cannot. The necessary scale that is required to truly follow and combat modern-day threats would require an amount of manual labor and analytics that is simply not achievable by any vendor or government." Rest assured, AI isn't out to get your menial or flashy jobs.

"Our use of AI technologies is confined to promoting better, more efficient cybersecurity," promises Gan, echoing the debate about ethics and bias issues that have been marring AI's image in recent years – a Google photo analysis tool labeled black people as gorillas, an Amazon recruiting tool was biased against women, a criminal risk assessment AI was favorable to white people, and we've yet to find out how the different autonomous car makers solve the "trolley problem," to name a few. However, Gan is confident that her company is immune to those aspects: "Detection accuracy is the key factor to the notion of practical prevention, i.e. the ability to employ cyber-defense technology in prevent mode, such that attacks are blocked at the gate rather than mitigated or remediated after the effect. This type

of usage is typically not as susceptible to bias or unethical use, simply by its nature. For a given protected infrastructure, the definition of an adversary is clear and non-ambiguous, so our challenge is focused primarily on reaching accuracy in our detection, rather than on determining if an activity is ethical or bias-free. With market reach across the globe, and with presence at every part of the IT infrastructure, our learning data sets are as versatile as they can be, and offer us a trustworthy source for training our algorithms".

Gan suggests that often AI is given the power of making decisions in real time:

"It has the ability to look at the data and then reach some conclusions, sometimes on its own and at other times in conjunction with other, non-artificial intelligence engines. But very often, it makes decisions on its own." Whether this could be a recipe for disaster is an issue we have to be aware of. "Artificial intelligence, as we all know, in its current incarnation, is very prone to errors, meaning it can offer false positives as well as misdetection. So if you rely solely on artificial intelligence, chances of getting it wrong could be high. And of course, unlike in image categorization where it's no biggie if I mis-categorized a certain image, or if I had to tell my Alexa something twice instead of once in order for it to understand, with cybersecurity, missing an attack, letting it go through or even categorizing something as an attack when in fact it isn't, can have rather detrimental implications for the organization".

AI, IN ITS CURRENT INCARNATION, IS PRONE TO ERRORS, MEANING IT CAN OFFER FALSE POSITIVES AS WELL AS MISDETECTION. SO IF YOU RELY SOLELY ON ARTIFICIAL INTELLIGENCE, CHANCES OF GETTING IT WRONG COULD BE HIGH.

So Checkpoint doesn't let AI run around with scissors, unattended. "Our studies demonstrated that AI systems cannot be blindly trusted," says Gan. "We are still at a point where human supervision is required, and the best results are achieved when juxtaposing several technologies, AI and traditional, together to reach higher levels of accuracy. We have also learned that field expertise is very much a necessity. Engines that claim to be general-purpose perform very poorly when applied towards solving cybersecurity problems, and tweaks offered by people with knowledge of the domain have led to major improvements in the overall performance."

The tech community needs to safeguard the technology, Gan says: "AI is still in its infancy. Like many different technologies, when they're being introduced people don't tend to think about the potential threats that they pose. And we have an opportunity here, when this is at a very early stage, to insist on cybersecurity being very much a part of this artificial intelligence revolution. If somebody was able to somehow poison the data from which the algorithm learns, it could influence the decision making in a way that could benefit this bad actor and hurt everyone else using it. These are very real risks that we need to address from the get-go, and not come back to later on when they've already matured. It's a little more difficult to fix it as an afterthought than build it in from day one."

Malicious hackers naturally also adopt AI. "A future in which AI is battling another AI is not far-fetched. Although keeping in mind the various motivations of the attackers – be they earning crime money or inflicting damage on the other side – the methodology may be vastly different and represent different uses of technologies and adaptations," Gan says. "Of greater concern may be future attacks targeting the AI algorithms themselves. The future cyber wars may very well be all about modifying the expected behavior of an otherwise-trusted AI engine, which offers attackers opportunities to generate bias or an alternate verdict in a way that benefits them. It may be very difficult to protect against such attacks, or even to identify their presence."

MAIA ARON, JOURNALIST

GETTING REAL: A LOOK AT AI IMPLEMENTATION IN INTERNATIONAL CLASSROOMS

While much of "Shaping the Future 5" conference was devoted to what AI can do, one panel focused on what it is doing – or not doing – in international classrooms today.

The three panelists, all experienced in largescale educational implementation, are chiefly concerned with students who are being educated during the very earliest stages of AI. Nevertheless, they will graduate into a world that will almost certainly be dominated by it.

What is the best way to prepare them, particularly when many are learning in disadvantaged environments that lack basic EdTech infrastructure?

Joyse John, Director of Education at UK-based Nesta, a foundation that has focused on introducing innovation for the past 20 years, said Nesta is just beginning to look at emerging trends in AI education and to invest in related companies.

"We've looked at everything from flipped learning, to digital education, to artificial intelligence, and we recently just started doing some exploration, just to understand all the AI tools out there," she said.

In addition, the foundation recently completed a wide-ranging research project with a group of partners that investigates the impact AI will have on specific jobs.

The research forecasts that by the year 2035 – a point when many of today's students will be entering the workforce - 10% of current jobs will see an increase in demand, 20% will see a decrease in demand, and 70% will have an uncertain future – they may or may not exist. Thus, her focus is to

identify and teach the skills that will match future needs and the personal qualities that will prepare students to adapt. The goal is to give them the best chance of having higher quality jobs in the future.

Regarding investment, she noted that Nesta is supporting companies that include Codebook, with expertise in personalization and adaptive learning, and Third Space Learning, with expertise in online math tutoring with real-time feedback that helps teachers improve their practice.

In addition, she noted Nesta's work in system-wide innovation, for example, using AI to predict which schools are at risk of failing and support them before they do.

She urged the audience to use an approach that begins in the classroom:

"Rather than starting with the technology and then looking for a problem, let's look for the problems that learners, teachers and the system face and try to solve that, because we definitely can do that," she said.

AI EDUCATION IS AT A VERY EARLY STAGE AND IS LIMITED IN WHAT IT CAN DO.

Carolina Jeux Conde, CEO of Digital Education Telefonica, concentrates on education and training in South American and European communities that lack the funds, infrastructure and training to introduce AI – but whose students will enter an AI-dominated world, nevertheless. Her company works with governments, companies and K-12 end-users. It has a platform that trains some 700,000 teachers in new methodologies and a MOOC platform for the Latin American world.

"To be realistic, we're very far from being able to include AI solutions into the educational systems in the sectors of the countries where we are," she said. "We are still talking about adaptive learning, blockchain, lots of technology is appearing, but the truth is that many governments don't have their content digitized yet and the teachers have a whole lot of needs and training just to launch a project-based learning service in their classrooms."

The solution?

"We are trying to teach the skills that AI can't do," she said, as well as trying to identify and implement the skills that are the most likely to be necessary.

"We're trying to give them confidence that it doesn't matter if they have to keep reinventing themselves," she said.

To that point, she noted that AT&T recently determined that more than 50% of their 250,000 employees are working in jobs that will be obsolete in five years. In response, the company launched a major re-skilling program in order to retain as many current employees as possible going forward.

Daisuke Asano, director of the Educational Servicing Industry at Japan's Ministry of Economy, Trade and Industry, is further along, despite the need to overcome a very traditional school environment.

In most Japanese schools, he explained, teachers write textbook-based information on a blackboard and students are tested by reciting the information. Teachers tend to be highly resistant to change and, because system-wide EdTech has not been introduced, most students don't have personal computers at school. In addition, schools are not connected to the Internet.

Nevertheless, his group initiated The Learning Innovation Project this year to change the role of the teacher and the system of learning in the classroom. Most notably, the Project's 30 EdTech pilot programs include a highly promising AI math application at a Tokyo junior high school.

The AI project requires students to learn math independently on laptops while teachers monitor, intervene and support their learning in real time.

"The teacher, looking at the real time data, understands which point, which chapter, the student is doing, how deeply the student understands, and what is the setback point for the student," he explained.

"This application will change the way of teaching and learning and communication in the classroom."

Moderator Prof. Karine Nahon, president of the Israel Internet Association, summed up with the major ethical questions posed by the introduction of AI to the educational environment, i.e. what should be regulated and who should do the regulating?

This is particularly significant since AI will be used to determine value and resource allocation. How can humans track mistakes, and who would take responsibility for correcting mistakes?

Daisuke Asano was alone among the speakers in reasoning that regulation does not need to be addressed at the current time. He argued that AI education is at a very early stage and is limited in what it can do. Therefore, he said, the government's main role should be to support and implement EdTech and AI, plus promote equality and accessibility of its content.

"It's too early for the government to regulate," he said. "The government should support it, not regulate it."

Q&A

WITH **BOB ROSENSCHEIN**

PIONEER OF INTERNET ENTREPRENEURIAL SOLUTIONS, FOUNDER, CHAIRMAN AND CEO OF ANSWERS.COM

Q There are many answers out there, but we all know that to get to the right answer, we need a very intelligent question. So my first question to you is, what kind of intelligence were you relying on when you were founding Answers.com?

B.R. "It went through several shifts – we call it a pivot. First, we started out with very simple factual encyclopedic information, but the pivot we did that made the most sense was to user-generated questions and answers. So we had a situation where people would put in any question – we can talk about the different kinds of questions – and other people could not only answer the question, but edited each other. We were the Wikipedia of questions and answers. The other thing we learned how to do was basically judge the quality of the questions and the quality of the answers through the user-generated process. There are many kinds of questions. There are simple factual questions – two plus two, how many pounds in a kilogram, what's the weather in Tel Aviv tomorrow. These are easier to answer. The harder ones require more judgment or multi-stage or more complex information. And we found the most interesting ones were the ones that blended opinion, background, depth, and length. So we had to develop

a system, we called it Remus, which basically measured, automatically and through user input, signals of both what we call ADQ, which was answers document quality, and AAK, which was answers author karma. And we came up with an SEO improver, basically, which was able to get the best questions and the best answers the most exposure in the system."

Q Users actually rated both the questions and the answers?

B.R. "There were many ways or many signals to this machine. Basically, we got a lot of traffic from Google. Google in those days believed very much in keeping you on their page for as short as possible. We also had a very special relationship with Google, where for about five years we had a special link in the corner of every Google English results page with a reference to Answers.com. So we got a lot of traffic from them. But what they did is spider and they indexed our pages very successfully. To give you a very simple example: we would get a new automotive question in the order of every 26 seconds on average. Someone would ask, "Where is the air filter in a Mazda 6?" – a very specific question. That question was asked, someone else may answer it, Google indexed it, and anyone in the world who wanted to know the answer to that question would find it. It would be a factual question but a very specific one, and we found that the more questions we had in the system, the more answers we had; the more answers – the more traffic. So it was a virtuous cycle which would give us opportunity. Some of the other questions we had were personal, about belief, relationships. People would ask Answers.com questions they would never ask their mother. I'm not even gonna go into it – "Can you get pregnant from...," it doesn't matter. I'll tell you the saddest question ever asked of our system. We had 17 million questions and answers in our database. The saddest question was someone one day asking, "Why do boys tell you they love you when they don't?" Anyway, many of the questions were much more factual or much more scientific – politics, history, chemistry – you name it, religion even, but you should realize that not all questions have simple factual answers. Some are opinion, some are expository, and we've benefited from all of those because basically we were capturing data, textual data primarily, that would be used to build the engine."

THE WORLD IS DEVELOPING TO WHERE PEOPLE ARE JUST LOOKING FOR VERY SPECIFIC KNOWLEDGE ON DEMAND

Q Actually crowdsourcing intelligence. So if people would have asked their mother or if they were too embarrassed, rabbi or a doctor, now they would ask Answers.com?

B.R. "Just like Wikipedia, we could not guarantee the accuracy of the answers. In fact, those signals were reinforced by traffic, reinforced by edits, and reinforced by scores. There were a lot of different ways. Our goal was to get the highest quality answers, but it's a difficult area and I would say that we used some machine learning techniques to measure the question quality and the answer quality, but at that stage we weren't generating answers intelligently."

Q You could tell by the way someone phrases a question, types in a question, what aspect of the answer or what version of the answer, that you have in your inventory, they were looking for?

B.R. "No. At that point we were just reading. Sometimes we would have many different versions. So if you had a question such as "When did WWII begin?" that could be phrased in many ways – WWII, Second World War – and there were, in fact, many different aspects. Someone said: "September 1, 1939"; someone else said: "Well, with this invasion." There are many different ways to look at something like that. The goal is getting information. By the way, nowadays Google does a fantastic job of giving a first cut at many of the answers, so not only can you

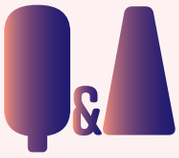

WITH BOB ROSENSCHEIN

ask what is the population of Athens – that's easy – they're actually answering questions like: "I wanted to know yesterday how to turn off shuffle on an album in my Apple Music, and they gave me a step-by-step thing." So they're doing much more. But systems like Quora and others, certainly Stack Overflow in the technical area, are fantastic at gathering that information today. And when we sold Answers.com we were well aware of where we felt the industry was going."

Q So today when people have a question they go on Google, and some people even feel like they have a personal relationship with Google – it can complete their question. So if I'm a hypochondriac, Google always completes my questions with a disease or some catastrophe. What do you feel is the added value of the human-generated Q&A-based system like Answers.com over machine learning?

B.R. "It's symbiotic because Google is looking for information out there on the web. They also supply some of their own very well, but they're looking for information. Other sites are presenting data and text and Google is very good at connecting it. People go directly to Quora and places like Answers.com and certainly Stack Overflow nowadays. Nowadays I might ask a question of Siri or Alexa, but the world is developing to where people are just looking for very specific knowledge on demand."

Q What do you think the future of asking questions to get the right answers is? Where do you think we can go?

B.R. "I think there will always be a need, of course, for getting quick information at your fingertips, when you want it, where you want it, how you want it, etc. So it may not be textual and maybe a little bit more word-guessing in advance. Amazon is coming up with a system now where they're going to send you what they think you want and if you didn't want it you put it in the box and send it back for free. So there are all kinds of different models for anticipating requirements in use, but It's a very exciting time in the industry and we'll see where it goes from here."

MAIA ARON

FROM DIGITAL TO ADAPTIVE: BUILDING THE PILOT'S ENGINE

HOW, EXACTLY, DO YOU BUILD AND IMPLEMENT AN ADAPTIVE LEARNING ENGINE?

Thanks to the CET-Microsoft partnership on the math class pilot program, we have a real-life example. The explanation comes from Daniel Sitton, engineering manager at Next Gen Edu, whose team developed the algorithm:

The pilot's digitized math curriculum was the starting point. Sitton's team collected a massive amount of data on each student's behavior and performance, while maintaining confidentiality of personal identities.

So, for example, the team could see that "Student A" was completing tasks quickly and perfectly, while "Student B" was taking longer, making mistakes and giving up. "Student A" was under-challenged and losing interest, while "Student B" was over-challenged and losing motivation.

An aggregate soon emerged that showed students moving through "zones of learning": those stretching to succeed were in a "challenge zone" and those giving up were in a "panic zone". The adaptive learning engine was designed to interact with students to help them reach their "effective learning zone".

Sitton explained how:

"Given the mass of data, we can create the machine learning models, which enables us to look into the future, meaning, for a student and a task that he has not seen yet, we can make a highly accurate prediction: what is the likelihood for this student, on that specific task, to coplete it on the first attempt; what is the likelihood for him to complete it faster than average, to use hints, or to abandon it, and so forth, and then we can take all those predictions … and determine whether a particular question would be inside a zone or outside a zone for this particular student."

Putting it into practice, the algorithm gives "Student A" challenging tasks that require use of hints and research, such as looking at videos, to reach the "effective learning zone". It gives "Student B" easier questions that build confidence and reinforce knowledge before proceeding to more difficult tasks. Students receive positive feedback as they succeed.

"Effective learning occurs when the content is getting more and more difficult as the student is gaining mastery, so anywhere those two things are well aligned we have effective learning happening, and the goal is to reach that perfect balance," Sitton concluded.

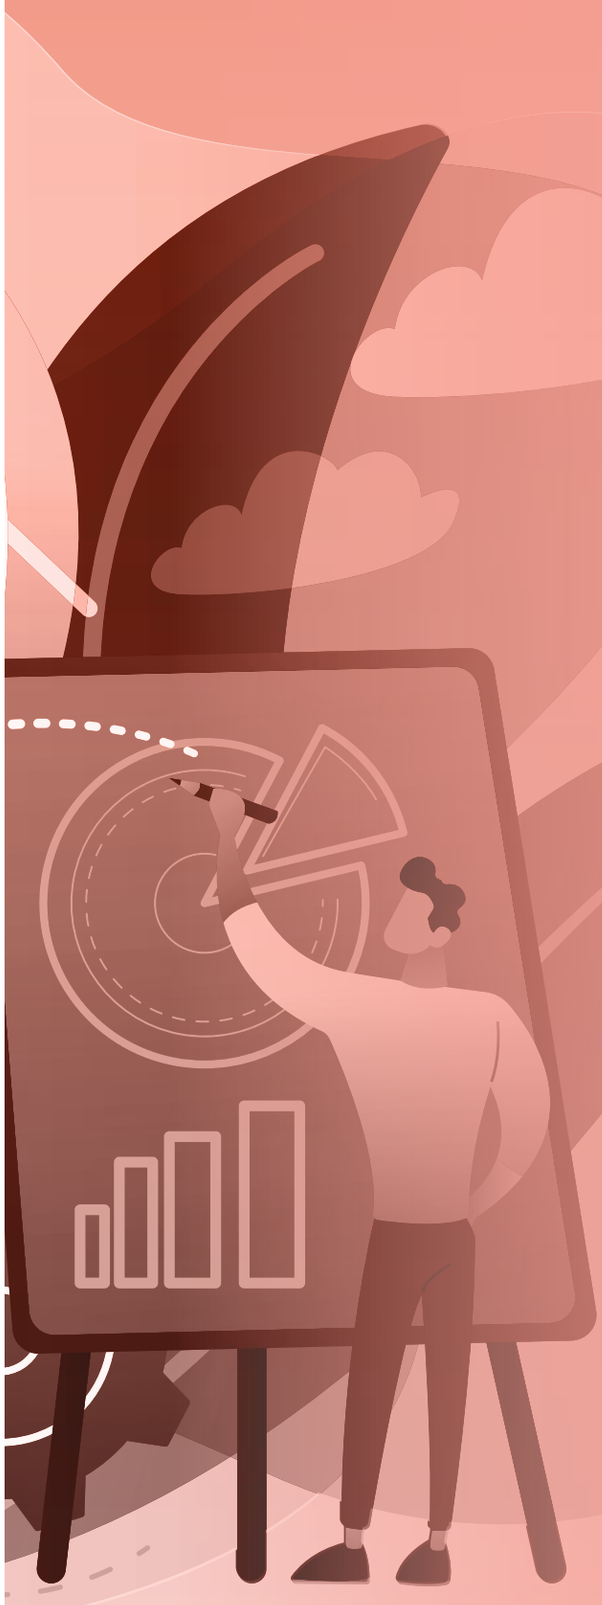

GUY LEVI, CHIEF INNOVATION OFFICER AT THE CENTER FOR EDUCATIONAL TECHNOLOGY

THE PROMISE OF ADAPTIVE LEARNING

The world's leadership is today being called upon to address the growing gap between the needs of the labor market and the skills that students acquire in the education system. This gap is a key issue that has concerned policy makers and many educators in recent years. The growing labor market of the 21st century is characterized by new professions that require a different kind of functioning, new abilities, and different skills from those that were appropriate for the last century. People are connected in new, previously unheard-of ways that influence the development and growth of the global economy. This new situation demands the development of new patterns of teaching and learning and new models to provide a response to these changing needs. The more advanced education systems in the world are setting such goals as education toward learning, exploring ways of learning ideas that are yet unknown, and being ready to solve problems that cannot yet be formulated. Technology has presented us with a new reality that will allow the creation of these new patterns and models. However, technology alone is not sufficient, as Michael Barber argues in his introduction to Fullan & Donnelly's (2013) paper: "…The future will belong not to those who focus on technology alone, but to those who position it within a broader context, and see it as but one component in a broad, comprehensive process of systemic change."

GUY LEVI

THE PROMISE OF ADAPTIVE LEARNING

A key component in this broad systemic change is "personalization," which means the adaptation of the teaching and learning processes to the personalized needs of the student. Hence, the educational and cultural emphases are now on 21st-century skills and competencies and they require new research and developments to meet the new objectives and goals. For example, in mathematics: whereas in the 1970s emphasis was placed on dealing with mathematical procedures efficiently and precisely, today the emphasis has moved to the solution of problems ("problem-solving skills" in 21st-century language), to application, reasoning, creativity, and critical evaluation. These goals, which had previously been almost totally absent from the goals of formal (K-12) education, were made possible through the development of technological tools (which had previously not existed) that could carry out the procedures. This also created an ability to focus the teaching and learning processes on other important ideas.

THE CHALLENGE OF ADAPTIVE LEARNING

Today, with the aid of the new technologies, it is possible to develop learning approaches that also include the use of representations, work on those representations, research into mathematical phenomena through dynamic technological applications, and feedback from the computer through mirroring of the outcome of the student's action ("intellectual mirroring"). The feedback allows students to solve problems, to research and test different alternatives and decide whether they have achieved what they set out to do, and, by testing, to generalize ideas and phenomena. In this regard, feedback is changed from a confirmation of prior knowledge – feed**back** – to the key to new knowledge – feed**forward**. This distinction is of utmost importance because, while feedback focuses on current performance, feedforward looks ahead to the next assignment, i.e. the predominance of formative assessment over that of summative assessment. In addition, the technology allows, on the one hand, the assembly of rich content to develop the required concepts and ideas in the field, together with the disciplinary goals and learning skills, while on the other hand it allows the student's learning abilities to be checked and analyzed using analytical tools applied to "big data," collected and analyzed on an ongoing basis, and, on the basis of matches between them, to construct teaching and learning processes appropriate to each student.

In the K-12 world of digital learning, there are today more and more solutions based on adaptive learning using various models – either products that promote personalization and offer a complete solution for a given syllabus (particularly in mathematics, although not only in that field), or as solutions that can be integrated into existing environments and products, and which offer significant added value to learning and to the transition to personalization:

the growing use of learning analytics (analysis of learning data) which allows the teacher to obtain a report on each student at any time and over time; the conception that learning has to be relevant and of value to the learner and the instructor, and thus should be active, not passive; and finally, the understanding that learning is possible at all times and is not limited to particular times or places – these explain the broad range of offerings and the variety of solutions that adaptive learning is today beginning to make available.

> THE MORE ADVANCED EDUCATION SYSTEMS IN THE WORLD ARE SETTING GOALS AS EXPLORING WAYS OF LEARNING IDEAS THAT ARE YET UNKNOWN, AND BEING READY TO SOLVE PROBLEMS THAT CANNOT YET BE FORMULATED

WHAT IS AN ADAPTIVE LEARNING SYSTEM?

Technological learning systems are considered adaptive when they change dynamically in regard to each student and in response to data collected in the course of the learning itself, to better adjust themselves to that learning. The system makes use of the data accumulating while the student is working, in order to change,

for example, the way in which a concept is presented, the level of difficulty, the sequence of problems or tasks, and the nature of the hints or feedback provided to the student. Thus, the students receive an individualized pace and pedagogic approach and a flexible study track in keeping with their needs, interests, and choices.

TYPES OF ADAPTIVE LEARNING SYSTEMS

There are three types of adaptive processes that can be found in the various systems:

TECHNOLOGICAL LEARNING SYSTEMS ARE CONSIDERED ADAPTIVE WHEN THEY CHANGE DYNAMICALLY IN REGARD TO EACH STUDENT AND IN RESPONSE TO DATA COLLECTED IN THE COURSE OF THE LEARNING ITSELF, TO BETTER ADJUST THEMSELVES TO THAT LEARNING

Model based on diagnostic tests – built on the basis of a test administered at the beginning of learning a topic. The student is tested on the requisite knowledge for learning the topic to come and on the knowledge that he is about to learn. The test indicates to both the student and the teacher what the student knows and what he does not yet know. Teaching continues with teacher adaptation of content from the existing content repertoire. This model is applicable to book-based learning, computerized activities, and other media.

The learning content may also be computerized or not, as the teacher chooses. This approach may be classed as personalization since the teaching is adapted to the variation among the students. The bulk of the work of adaptation, as well as significant portions of the teaching process, remain in the hands of the teacher.

Rule-based model – constructed through the use of "if-then" functions. The student is asked one or more questions. If he answers correctly, he moves to the next activity or content unit; if he errs, he receives a hint or an explanation that is somewhat different from the earlier one, depending on the answer that he has chosen. This model is computerized; it makes use of the knowledge of experts in the knowledge domain, who create branched structures and rules for progress regarding the content being learned. All the progression rules are defined in advance by the content specialists. In this approach, both the adaptation and many parts of the teaching process are carried out using the computer.

Algorithm-based model – a computerized model constructed using mathematical (statistical) functions that analyze the student's performance and collect information on content. The more students who work with the system, the more up to date and precise the data is. As learning progresses, the system learns more and more about the student and about the content and is able to combine what it has learned more effectively. The model is capable not only of assessing what the student has already carried out but also of adding information about what the student knows at a higher level of detail, so as to more precisely adapt content to him. Systems such as this make use of educational data mining, are involved in big data analysis, and create complex algorithms for predicting the probability of a student's success in learning particular content. The learning tracks are built from computerized analysis of the student's performance, and an infinite number of such learning tracks can be created. This adaptive learning system may be facilitator-driven and provide real-time data that the teachers can use and based on which they can act. At the same time, the system can be assessment-driven, with the ability to carry out the adaptation itself, thus allowing the students to progress on their own as they proceed through the course. The future of AI promises new breakthroughs in adaptive learning models which will be integrated also into skills and competency-based education.

Reference

Fullan, M., & Donnelly, K (2013). "Alive in the Swamp, Assessing Digital Innovations in Education." Nesta, NewSchools Venture Fund.

Photo by Shutterstock

DR. ORLY FUHRMAN, RESEARCHER ADAPTIVE LEARNING, CET

MIND READING, READING MINDS

When we say "artificial intelligence in education" what we really mean is "reading students' minds." For the teachers among us, this will not come as a big surprise: standing in front of a classroom of 10, 20, or 30 students, and even more, practically means engaging in continuous mind reading. "Does the puzzled look on that girl's face mean she didn't understand what I just said? Or maybe something happened at recess and she can't stop thinking about it. And what about that boy in the front row, vigorously nodding after each example I give? Maybe this is too easy for him and he's already bored and uninterested. How can I make sure both students **stay** with me and make progress?"

Indeed, peeping into the learner's mind, the infamous "black box," is a long-time fantasy of every curriculum designer and educator. It is important not only for assessment, but also for efficient teaching, one that adapts to the level, pace and inclinations of the class as a whole and of each individual student in real time. Probing the "black box" is of special interest when a complex process like reading is at hand. The challenges a young reader faces span a wide range: from bottom up to top down skills, from low level demands such as translating symbols (letters) into sounds, through deciphering the meaning of whole sentences and passages, sustaining attention, to making inferences, connections and interpretations.

WHEN WE SAY "ARTIFICIAL INTELLIGENCE IN EDUCATION" WHAT WE REALLY MEAN IS "READING STUDENTS' MINDS"

At present, reading comprehension, as in other fields, is commonly assessed by question posing. If we wish to peep into readers' minds, it seems like the perfect solution to patiently wait for them to read, and then ask them a series of reading comprehension questions. Right? Wrong. A post-reading questionnaire is often (justifiably) perceived as a performance test, and becomes the goal of reading for many students. Oftentimes, students will not start with the text at all, but rather read the questions and try to "fish out" or guess the correct answers. Not only does this introduce stress and alter the goal of reading, but, most importantly, it doesn't relate to the reading process itself as a learning opportunity.

Reading a text, whether printed or in digital format, can be a very lonely experience. This is especially true for struggling readers, particularly young ones, who have not yet mastered the inner dialogue, or the inner Q&A, needed to read through a text and truly understand it. They can benefit from the guidance and support of a reading tutor who will sit by their side and ask: "Are you familiar with this word?" or "What just happened?" and "Can you guess what's going to come next?" An experienced reading companion would adapt these questions according to the signals young readers send them: reading aloud and deciphering individual words for the non-fluent reader; asking more scaffolding questions to keep the interest of the inattentive reader; and making connections and asking interpretive questions to broaden the understanding of the more shallow reader. Unfortunately, personal reading companions are not always at hand, and one-on-one teaching sessions are a luxury only to be dreamt of in most educational settings.

How can artificial intelligence come to the rescue? Adaptive learning technology, which promises tailored learning experiences for each student, may provide the answer. Typically, an adaptive learning platform would collect data from a big group of students who complete the same learning program. Based on data analysis of many different dimensions, it would offer an activity to individual students, record their interaction with it, and, similarly to that human reading companion, suggest the next step accordingly. Classic modifications in the student's learning path could be handing easier exercises for struggling students, or offering more scaffolding in same-level activities, alongside more sophisticated challenges that would keep the advanced students interested. Of course, when reading is involved, focusing on the post-reading questionnaires alone would make us miss out on myriad opportunities to understand individual reading challenges and help readers surmount them. AI-based digital reading tutors could cater to different readers' needs in real time, **while** they read. This way, teachers won't have to wait for the text, the learning unit, or the semester to end in order to know where their class is at and offer assistance. Moreover, an adaptive reading unit can decrease frustration and offer solutions and learning opportunities before it is completed; for example, collecting eye or mouse movement data can point to different reading paces, areas of difficulty and even sustained attention patterns. "Live text" that allows for interactions such as glossing unfamiliar words, asking questions or answering embedded ones can give clues as to the right questions to pose to individual students and, of course, which text they will be able to handle next.

Adaptive reading programs can instill readers with self-esteem and empower them to handle texts more skillfully. They can also, in turn, point teachers to the exact areas where individual students need guidance and support, thus meeting class-wide goals while making each student feel like a competent, successful, reader.

Israel 2019 EdTech Week

MindCet Changing Education Mindset

CET • THE CENTER FOR EDUCATIONAL TECHNOLOGY

join this dynamic,
ingenious and groundbreaking event

Israel EdTech Week

September 16-19, 2019

and meet the protagonists of the Startup Nation including the most daring annual EdTech Summit –SHAPING THE FUTURE 6 with world renowned innovators showcasing how education is being transformed

Part of: DLD TEL AVIV INNOVATION FESTIVAL

For more information:
shapingthefuture@cet.ac.il