

EDUCATION: FUTURE FRONTIERS

OCCASIONAL PAPER SERIES

Educating for a Digital Future: The Challenge

MARC TUCKER | PRESIDENT OF THE NATIONAL CENTRE ON EDUCATION AND THE ECONOMY



ABOUT THE AUTHOR

Marc Tucker is the President of the National Centre on Education and the Economy, Washington, DC USA, and Visiting Distinguished Fellow at the Harvard Graduate School of Education.



EDUCATION: FUTURE FRONTIERS is an initiative of the NSW Department of Education exploring the implications of developments in AI and automation for education. As part of the Education: Future Frontiers Occasional Paper series, the Department has commissioned essays by distinguished authors to stimulate debate and discussion about AI, education and 21st century skill needs. The views expressed in these essays are solely those of the authors.

I've been asked to write a paper on the challenges and opportunities of an artificial intelligence future for education and learning in school and beyond, and to tell you, the reader, what skills will be needed by your students when they enter the adult workforce and the role of education in fostering those skills. But responding to that request in a serious way presupposes that we know—or at least I know—what kinds of challenges an artificial intelligence future will pose and what opportunities such a future will unveil.

What I know is that these questions are currently the subject of a spirited debate, a debate that is based on well-informed visions of an artificial intelligence future that range from the utterly dystopian to the unreservedly utopian. Whatever I might offer by way of suggestions for educators cannot, if they are to be useful, embrace the full range of dystopian to utopian images of the future. And so I begin by trying to help you understand developments in this very fast moving arena, sharing my own interpretation of the range of possible futures and then, and only then, telling you what I think the implications are for education policy.

The subject, however, is not one subject but many, all of which are very complex and all of which are evolving very quickly. The best I can do is skip lightly over the surface. I have for that reason ended this essay with a reading list, both to give you an idea of the sources I have consulted and to invite you to come to your own conclusions based on your own readings of these sources and the ones that are added every day to this literature.

For many educators, the definitive book on this subject is *The New Division of Labor: How Computers are Changing the World*, by Frank Levy and Richard Murnane. Published in 2012, the book begins by pointing out that there have been repeated apocalyptic warnings about computers putting people out of work, but that future has not yet materialised. Levy and Murnane conclude on the basis of a wide-ranging review of the literature and a thorough analysis that it won't—or at least need not—happen this time either.

WELL-INFORMED VISIONS OF AN ARTIFICIAL INTELLIGENCE FUTURE RANGE FROM THE UTTERLY DYSTOPIAN TO THE UNRESERVEDLY UTOPIAN.

They tell us that intelligent machines are exceptionally good at executing algorithms conceived of as “routines”, which makes them better than humans at a wide range of low and medium skill tasks that essentially involve routine work. But, they say, as such jobs are taken over by the machines, putting people who only have what the educators think of as the old “basic skills” out of work, other jobs—much better paying jobs—are springing up, jobs entailing extensive problem-solving, expert thinking and complex forms of interpersonal communication.

These authors were not polyannas. They were worried that national education systems might not be able to provide vast numbers of people who now get only the

basic skills when they enter the workforce with the much more advanced skills they would need for the jobs that would become available. If that did not happen, if educators could not produce a transformation in the skill endowment of national populations, then the job market would polarise, incomes would polarise and the resulting political tension could threaten our democracies. I came independently to much the same conclusions long ago and have been preaching that gospel for years.

Along the way, Levy and Murnane provide us examples of tasks that workers do that AI will enable machines to do well. They also give us examples of the kinds of tasks that the machines cannot do and will not be able to do for the foreseeable future. A prime example of the latter was driving a car. Their book was published in 2012. Only five years later, Google's cars had been driving themselves down California highways for two years. It is not just the example that is out of date. The whole analysis may be out of date.

Steven Pinker's *How the Mind Works* helped me understand how we got to 2017. Though it was published in 1997, it is still is the best book on its subject. Pinker set out to write a book for specialists that would advance the field while at the same time writing a book for well-educated generalists to introduce them to the field, and he succeeds. Pinker describes psychology as a discipline that for decade after decade did something that might be compared to trying to understand how a steam engine works without ever taking one apart. He is not kind to behaviourism, clinical psychology or any of the precursors to cognitive science, all of which still have an enormous influence on the thinking of educators all over the world.

Pinker points out that the people who pioneered artificial intelligence were rarely psychologists and the psychologists, until recently, took very little interest in thinking machines. But, early on, the artificial intelligence

community concluded that they could only make progress by conceiving of intelligence as a form of computation, the kind of computation that underlies information processing. The key to the success of cognitive science in unlocking the way the mind works is that it, too, defines intelligence as a process of computation. In the computer, the information processing algorithms are implemented in silicon, in the mind by cells and electric currents. There are limitations and possibilities in both mediums that are very different from each other, but cognitive scientists and artificial intelligence researchers are essentially studying the same thing: the algorithms that account for intelligence and intelligent behaviour.

The early version of artificial intelligence assumed that intelligence is what happens when humans invoke mental procedures in the form of algorithms that follow deductive logic. My 1987 dictionary defines an algorithm as "a set of rules for solving a problem in a finite number of steps, as for finding the greatest common divisor." The same dictionary defines intelligence as the "capacity for learning, reasoning, understanding and similar forms of mental activity."

You noticed, of course, that there is a world of difference between these two definitions. The dictionary's definition of algorithm invokes the image of a deductive process that converts a set of inputs into a predetermined output using a set of tools that follow an inexorable logic. The definition of intelligence goes far beyond that to include learning, reasoning and understanding. The difference between the world that Levy and Murnane were looking at and the world in which Google's cars were driving themselves down Highway 101 in California is the difference between those two definitions.

Cognitive science and the artificial intelligence community both drew heavily on the computational theories of information worked out during and after the Second World War by Claude Shannon and other pioneers.

But it was not until these two fields started to draw on each other, as each advanced, that artificial intelligence and cognitive science both really accelerated in a kind of intellectual symbiosis.

While that was going on, Moore's Law, predicting a doubling in computer speed and capacity every two years, was doing its work. Computers were becoming more powerful on a logarithmic curve and the development of global networks began to provide those computers access to unimaginable amounts of data. This was a formula for impressive developmental growth.

In the first instance, these technological developments made computers conceived of as powerhouses of deductive logic much more powerful than they had been previously. The IBM computer that beat the world's leading chess champion did it by computing all possible moves faster and more accurately than any human can. You might think of that as brute force computing.

But, at the time, shrewd observers noted that the same machine could not perform many of the cognitive functions that a normal three-month old child could do easily. Nor could it demonstrate any of what most of us think of as common sense. It had no idea what human emotions were much less identify them in action, have them or respond to them. It could find and regurgitate information that was given to it, but had no idea how to formulate a problem nor was it able to learn how to do something it had not already been taught to do. This is the world that Levy and Murnane were writing about.

It turns out that playing chess is a very bounded problem, one very suited to deductive logic and sheer computing power; but one cannot assume that a machine that can beat the world's chess champion is an intellectual giant. All in all, a three year old is much smarter. But a few years after Levy and Murnane wrote their book, a Google machine won a game from an expert player of the Chinese game of Go. There are almost an infinite number of possible moves in that game. It cannot be won in the same way as a chess game can be won. Go players win by a kind of intuition based on pattern recognition. It is a very human kind of cognition, the kind we developed to assess a very complex situation almost instantly on the savannah quickly enough to avoid getting killed there 200,000 years ago.

IT WAS NOT UNTIL
COGNITIVE SCIENCE
AND AI STARTED TO
DRAW ON EACH
OTHER, AS EACH
ADVANCED, THAT THESE
FIELDS BOTH REALLY
ACCELERATED IN A
KIND OF INTELLECTUAL
SYMBIOSIS.

By that time, Levy and Murnane had been proven right... at warp speed. Waiters and waitresses were being put out of work by i-Pads stuck on dining tables that enabled the customer to place an order and pay her bill. Grocery clerks were being replaced by machines that automated the check-out lane and took automatic inventory. Miners were being replaced by automatic mining machinery which not only did the mining, but took the ore to the surface, loaded it on driverless trucks, offloaded it on to automated trains and then automatically put the ore on the ships to take it to China. Automated equipment had long since replaced the gas station attendant. Robots were being ordered by the millions to replace the Chinese workers who had been making the laptops, smartphones and inkjet printers sent from the coastal provinces all over the world. These developments were not only idling literate but only moderately skilled people by the millions in the developed world, but were also removing rungs from the ladder the people in the developing world had been using to join the developed world. In the United States, manufacturing accounted for as much of the gross national product as it had 30 years earlier, but accounted for a much smaller fraction of total employment. Machines were rapidly replacing humans on the factory floor.

But Erik Brynjolfsson and Andrew McAfee describe another effect of the advance of intelligent machinery that is less well understood by the general public and no less important in their seminal book, *The Second Machine Age: Work, Progress and Prosperity in a Time of Brilliant Technologies*. It has been described as the 'winner take all' phenomenon. These authors use the example of the Eastman Kodak company to make the point. At its height, Kodak employed more than 145,000 people helping others share billions of photos, as well as thousands more in its supply chain. And then it went bankrupt, a victim of the conversion to digital photography. A team of 15 people at Instagram developed an app which was also

used by customers all over the world to share billions of photos. Fifteen months after they founded the company, Instagram was sold for over \$1 billion to Facebook.

Everywhere we look, small groups of very highly educated and trained people are creating applications (think algorithms). The first one costs a great deal of money. But the next copy costs virtually nothing, and the one after that and the one after that cost no more. More often than not, the product can be used all over the world. A small group of people in San Francisco run a worldwide taxi company, putting countless taxi companies out of business. They own no taxis, only the rights to an algorithm. They are now developing taxis that will drive themselves to customers who will call them with their cell phones and pay them with the same cell phones. No taxis, no drivers, no clerks, no dispatchers. Why pay to listen to a local musician when you can hear the world's leading musicians for next to nothing on your smartphone? Why go to the mall when you can sit in the comfort of your own home, comparison shop worldwide and have the product you are looking for at a great price delivered to your door for nothing? Department stores are going bankrupt and malls are closing all over the developed world and the people who used to work in them are being replaced by apps and machines controlled by the companies that got there first. A handful of winners become very, very rich doing this, but a great many people are ending up less well off, on contingent employment or simply unemployed.

The first stage of machine intelligence extended this line of work by incorporating the accumulated craft and intuitive knowledge of renowned experts in a variety of fields into the machine's database, in a process called "knowledge engineering". The knowledge—which could include, for example, the diagnostic knowledge of renowned doctors and medical researchers—was certainly not routine, but putting that knowledge at the disposal of

rural family doctors did not involve machines that could learn something that was not already in their database, nor did it require the machine to demonstrate intuition, distinguish someone who is sad from someone who is happy, have the common sense of a six-month old child or communicate to an artificial leg with electrical signals all the information the brain normally supplies to a real leg required for it to accomplish the incredibly complex movements that all of us make countless times every day.

SOME OF THE BIGGEST INVESTMENT COMPANIES IN THE WORLD ARE REPLACING THEIR VERY HIGHLY PAID ANALYSTS WITH ALGORITHMS THAT SEEM TO BE MAKING INVESTMENT DECISIONS JUST AS SOUND AS THOSE MADE BY THE PEOPLE THEY REPLACED.

The situation is very different now, since Levy and Murnane wrote their book. When researchers played a classical music concert to a group of expert critics recently and asked them which piece they preferred, the majority selected a piece written by a computer that they praised for its emotional power. They were enraged when told that it had been written by a computer. Music companies are now employing computers that analyse popular music to find out what distinguishes the Platinum hits from those that do not do so well; the computers then write original songs that mimic the best. Computer programs are now capable of minutely analysing ordinary human speech to discover the patterns that correspond to various

human personalities and using that information to match people who call in to customer service centres to staffers who will make them feel comfortable. Popular real estate websites feature software that estimates the value of homes both on the market and not on the market using the same factors and values that licensed appraisers use, putting the appraisers out of business. Other programs can discern from the patterns of relationships among the features on people's faces what emotions they are feeling and changing the content of ads in response. There are now programs that will enable soldiers whose limbs have been blown off and replaced by artificial limbs to communicate with and thereby control those limbs with their thoughts alone. Some of the biggest investment companies in the world are replacing their very highly paid analysts with algorithms that seem to be making investment decisions just as sound as those made by the people they replaced. None of this sounds like the routine work described by Levy and Murnane.

Perhaps the most interesting recent development is machine learning. The current version of the Oxford English Dictionary does not define algorithm as a set of rules for solving a problem. It defines algorithm as a "process or set of rules to be followed in calculations or other problem solving operations. . . .". What could a problem solving process be if not a process involving following a set of rules, especially if we are speaking of a process that is best described as information processing? What if I told you that what we are speaking of here is decisions made by intelligent machines on the basis of inference rather than deduction, on the basis of probabilities rather than hard facts, on the basis not of what has been programmed into the machine but on the basis of what it decides it has to learn from data it decides to gather? What if I told you that when expert programmers look at the algorithms driving the most

WHILE IT IS STILL TRUE THAT THERE ARE VASTLY MORE CONNECTIONS AVAILABLE IN THE HUMAN BRAIN THAN IN ANY COMPUTER, COMPUTERS ARE MUCH FASTER AND ARE NOW CONNECTED TO A WORLDWIDE MEMORY BANK FAR LARGER THAN ANY HUMAN'S LONG-TERM MEMORY.

advanced machine learning systems, they have no idea how the machines reached the conclusions they reached because there is no train of deductive logic for them to follow. The machines are deciding for themselves what to do and how to do it.

The new generation of machines are eager learners. Give them a goal, a set of algorithms and a mountain of data and they will learn what they need to learn to reach the goal, remarkably quickly. They will develop a theory. It could be wacky. Then test the theory out on the data. It might work a little. It will change the theory a bit. It might work a little better. It will keep doing this over and over again until the theory can not only explain the data it started with but a great deal of new data it gets its hands on. This is the essence of human intelligence. The search for patterns that explain a great mass of seemingly unrelated phenomena is what Einstein was doing in the customs office. It is a long way from brute force calculation.

What has made this possible are enormous advances in information processing speed, the ability to see patterns where before they saw only confusion, the sophistication of the algorithms available to them and access to enormous amounts of data, courtesy of the World Wide Web and the very large data banks being assembled by businesses, government and researchers. While all of that has been going on, other people have been making rapid advances in sensors of all kinds and in the degree to which these intelligent machines are at home in the world, speaking here of the kinds of things that are second nature to a six-month old but have been very hard for intelligent machines. They have not yet made machines with the flexibility and skill of the human hand nor do these machines yet have the common sense that a six-month old has, but remarkable progress is being made and there is no reason to believe that it will not continue.

While it is still true that there are vastly more connections available in the human brain than in any computer, computers are much faster than the connections in the brain and are now connected to a worldwide memory bank far larger than any human's long-term memory. This is a recipe for a subtle, flexible and powerful intelligence. It is no longer a question of what the machines can do; it is a question of what they cannot do, a domain that is getting smaller quickly.

Four years ago, Carl Benedikt Frey and Michael Osborne, a pair of Oxford University researchers, calculated that half of the jobs in the United States economy could be automated by equipment when available. More recently,

the McKinsey Company, a consulting organisation, completed a more sophisticated analysis. Combining a list of the functions that intelligent machines can now accomplish and running that against a detailed description of thousands of different kinds of jobs tracked by the United States Department of Labor, they looked at which parts of those jobs could be done by the machines and which parts could only be done by humans. McKinsey concluded that fewer than five percent of American jobs will be fully eliminated by intelligent machines. Their report envisions a world in which machines and humans do most jobs together, welded at the hip. That is rather more comforting than the Oxford report.

But consider one of the jobs McKinsey analysed, retail sales. One of the functions of a retail sales person is greeting customers, which, according to McKinsey, requires such capacities as “sensory perception,” “social and emotional sensing” and “natural language generation”, which the machine, it says, cannot yet do. But retail malls, as I said above, employing very large numbers of people, are closing all over the United States, as are the giant department stores that used to anchor those malls, because customers prefer to sit in their living rooms ordering the stuff they used to buy in malls from Amazon. Amazon employs far fewer people than worked in the establishments it is replacing. And, even so, Amazon is working hard to replace many of the people in their warehouses with automated equipment. Another version of the Kodak story, but on an even larger scale.

In this case and many others, the McKinsey analysis makes very little sense to me. Ignore for the moment the fact that intelligent machinery is available right now that is quite good at sensory perception, social and emotional sensing and natural language processing. Focus instead on the fact that Amazon did not deconstruct the job of the retail sales clerk and then use machines to do only

the ‘automatable’ parts. They did an end run around the whole retail enterprise, which is precisely what is occurring in one domain after another.

The consequences are all around us. Not only are we seeing job categories employing millions of people suffering as a consequence, but it is now clear that those people who have become underemployed or unemployed as a result of the introduction of these technologies are not getting new jobs that will enable them to live as well as they did when they had the old ones. One of the most important reasons that the advanced industrial nations have not seen wage inflation as they have been recovering from the Great Recession is that so many people who used to have full-time well-paying jobs are now willing to take part-time jobs and jobs paying much less than they used to make because they do not have the skills needed to join the ranks of the fortunate few who do the high-paying jobs that are available.

It was not the former production workers at Eastman Kodak who wrote the Instagram apps. The few who do have those remarkable skills are able to command astronomical salaries, benefits and stock options, to say nothing of working conditions that might have been envied by King Tut. But there are very few of them, and their ranks are not increasing at anything near the rates that jobs for those with less esoteric skills are declining. Average productivity is not rising the way economists expected it to because, while a few people are much more productive, many are much less productive.

The results are very sobering. A recent new book, *The Vanishing Class: Prejudice and Power in a Dual Economy*, by Peter Temin, an MIT economist, tells us that a model used by Nobel-winning economist W. Arthur Lewis more than 70 years ago to explain the economics of low-income developing countries perfectly describes

the United States today as a dual economy with islands of rich people who have most of the investable savings surrounded by a much larger group of people just trying to get by. Larger and larger fractions of the working age population in the United States have been dropping out of the workforce, unable to find work at all, so dispirited and depressed that they have become the epicentre of the national epidemic of opioid drug abuse. The American economy is splitting in two pieces. One piece—highly educated and skilled—is benefitting hugely from the new technologies I have been describing—at least so far—and the other, undereducated and less skilled, is being put out of work by them.

The idea that the people I have just described should be thought of as surplus labor and put on a permanent dole has left the realm of the think tanks. Countries and cities in the developed world are now implementing policies based on that idea. The future has arrived. The political tensions that inevitably accompany increasingly polarised incomes and opportunities are now on view on the evening news programs on TVs all over the developed world.

It is not a law of nature that the introduction of new technologies will put a lot of people out of work in the short term, but will then create just as many new jobs that are even better in the long term. What is distinctive about these technologies is that they incorporate the very thing that makes us so different from any other thing animate or inanimate on earth: high intelligence. We have gotten inside the black box of the mind and have been very busy reverse engineering it. It is now becoming clear that intelligent agents already exceed human capacity in some domains of intelligent behaviour. The only question is whether they have the potential to exceed humans in all domains of human intelligence, and, if they do, how long it will take to get there. In this crucial sense, these technologies are unlike anything we have seen before.

Reading all of this material has led me to two conclusions. One is that the first stage of the evolution of these technologies is well advanced in its implementation and is now driving the economic divide I just mentioned. That stage has been characterised by what is becoming a vast extinction in the advanced industrial countries of the kind of jobs requiring basic literacy that the industrial model of public education was designed to prepare most graduates for. If that were the end of the story, the solution would be to redesign our education systems to prepare all of our graduates for the kind of work that our elites have been doing—professional work requiring complex thinking skills, deep knowledge in multiple domains, strong communication skills and social skills, strong values and strong character. That is an enormous task, but one that a growing number of countries are learning how to do.

But that is not the end of the story. I have come to the conclusion that the first stage will be succeeded by a second stage in which the utopian and dystopian possibilities I described earlier loom into view, a world in which intelligent agents take on more and more tasks now done by humans and accomplish many of them more effectively and efficiently than humans can do them, a world in which it becomes harder and harder to distinguish the human from the machine as we find more and more ways to alter our genes and augment not just our motor capabilities but our emotional and intellectual capabilities with intelligent agents.

If the human community continues on its current course, Hariri's vision of the future seems all too probable to me, a future in which a small number of humans manage to become literally immortal and to live on forever a life of immense power and wealth, a larger number may live quite well—though not forever—in the style of Renaissance artists, thinkers and craftspeople serving the ultra wealthy and the vast mass of the people thought of as surplus labor are paid out with a universal basic

HUMAN BEINGS WERE BORN TO WORK. OUR SURVIVAL DEPENDED ON IT AND SO THE WORK WE DO BECAME FOR MANY OF US THE SOURCE OF OUR PRIDE AND OUR IDENTITY. THE IDEA OF A DOLE FOR OUR SURPLUS LABOR FLIES IN THE FACE OF THAT REALITY.

income. It is all too possible that will be a world, again like Renaissance Italy, in which the wealthy clans are constantly duking it out with the other clans, only this time with weapons of unimaginable destructive power.

That is not a world I want for my grandchildren—that is, after all, whom we are talking about here—even if they are able to become members of one of the first two classes.

The utopians have a point. We may indeed be on the cusp of being able to cut and edit our genes so as to eliminate a vast range of diseases, feed the millions with nutritious foods grown in a way that will not poison the planet, process all our waste to turn it into the resources we need to provide for everyone, and in general, provide a good life to virtually everyone while restoring our home—planet earth—to health.

Doing that would require a human population with great imagination and high skills. More than any technical skill, it would require a very high order of political skills, not just on the part of our political leaders, but on the part of the citizens who vote for them—or fail to do so.

If we succeed in this venture, most people who wanted to do so could lead a life of leisure filled with creative and rewarding activities—social, artistic, intellectual. There could be plenty of what we now think of as work for those who wanted it.

But, to get there, we would have to reconceive how the bounty I just described could be created and distributed. Human beings were born to work. Our survival depended on it and so the work we do became for many of us the source of our pride and our identity. The idea of a dole for our surplus labor flies in the face of that reality. If intelligent machines end up doing most of the work that is needed to provide the stuff and the services we need and want, we will have to reinvent our social and political and economic systems to make the arrow point toward the more utopian visions rather than the more dystopian ones. That cannot be done by a few political leaders acting out of rare foresight on their own. It will have to be done by the people.

I conclude from all this that the prescription I shared above calling for reshaping our education systems so that all students are offered an education previously reserved for an elite is correct but not enough. Yes, many more students will need strong cognitive skills, much deeper

knowledge and much more sophisticated skills, if they are going to be partners to increasingly intelligent agents and not put out of work by them in the near to intermediate term. And they will need to be very strong where the intelligent agents are, at least for the time being, relatively weak—in areas like creativity, imagination, and the whole range of social and emotional and communication skills that will be the necessary complements to intelligent agents.

THE QUESTION OF WHAT IT MEANS TO BE HUMAN HAS NEVER BEEN MORE URGENT.

But that leaves out what I take to be the decisive factor as the second round of the development of intelligent agents gains steam: the question as to what it will mean for our lives together and for what it means to be a human being, for the distribution of opportunity and wealth and fulfillment. It is in this realm that education may turn out to be decisive in determining the future of humanity. The question of what it means to be human has never been more urgent. The need to understand history at a deep level in order to prepare ourselves for the future has never been more urgent. The need to enable students to understand others very different from themselves and to be able to see the world from their point of view is essential if we are going to avoid blowing ourselves up on the way to utopia. The liberal arts are disappearing from colleges and universities in

the United States as students, increasingly anxious about their economic future, focus their time in college and university on their vocational goals. But the liberal arts—reconceived—may be the key to our survival as a species.

It is essential that we reconceive schooling not just in terms of greatly ratcheting up the standards of students' cognitive development and not just adding to that the need to provide in a very deliberate way for the development of students' communication, social and emotional skills and, more broadly, their character, but also to reconceive the curriculum in a way that will prepare students for citizenship in a way and to a degree that is totally new, for a world that will call on them to make unprecedented decisions about the structure of their societies, the structure of their economies, the nature of work and their responsibilities to others in the world that intelligent technology is creating. Above all, a curriculum that is about values, about what it means to be human and what we value about being human. If we fail at this task, it may only be a matter of time before the machines and a very small technological elite are deciding these issues, and we are not likely to be happy with their decisions. ➔

READING LIST

I'D RECOMMEND FOUR BOOKS TO START, READ IN THE FOLLOWING ORDER:

Erik Brynjolfsson and Andrew McAfee, *"The Second Machine Age: Work, Progress and Prosperity in a Time of Brilliant Technologies"* (New York, W.W. Norton, 2014).

Martin Ford, *"The Rise of the Robots: Technology and the Threat of a Jobless Future"* (New York, Basic Books, 2015).

Steven Pinker, *"How the Mind Works"* (New York, W.W. Norton, 1997).

Yuval Noah Harari, *"Homo Deus: A Brief History of Tomorrow"* (New York, Harper-Collins, 2017).

HERE ARE SOME OTHER BOOKS THAT MIGHT INTEREST YOU IF YOU WANT TO GO BEYOND THIS CORE:

Pedro Domingos, *"The Master Algorithm: How the Quest for the Ultimate Machine Will Remake Our World"* (New York, Basic Books, 2015).

Nick Bostrom, *"Superintelligence: Paths, Dangers, Strategies"* (Oxford, Oxford University Press, 2014).

Adam Piore, *"The Body Builders: Inside the Science of the Engineered Human"* (New York, Harper-Collins, 2017).

Richard Yonck, *"The Heart of the Machine: Our Future in a World of Artificial Emotional Intelligence"* (New York, Arcade Publishing, 2017).

Ray Kurzweil, *"The Singularity is Near: When Humans Transcend Biology"* (New York, Penguin, 2005).

ARTICLES, REPORTS AND OTHER BOOKS

Daron Acemoglu and Pascal Restrepo, *"Robots and Jobs: Evidence from U.S. Labor Markets,"* National Bureau of Economic Research, March 7, 2017

Daron Acemoglu and Pascal Restrepo, *"The Race Between Machine and Man: Implications of Technology for Growth, Factor Shares and Employment,"* National Bureau of Economic Research, May 2016

David Autor, *"Why are There Still So Many Jobs? The History and Future of Workplace Automation,"* Journal of Economic Perspectives, Summer 2015

David Autor and David Dorn, *"The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market,"* American Economic Review 2013

Carl Benedikt Frey and Michael Osborne, *"The Future of Employment: How Susceptible Are Jobs to Computerization?"* Oxford University, September 17, 2013

Kai-ming Cheng, *"Advancing 21st Century Competencies in East Asian Education Systems,"* Asia Society, 2017

Committee on Information Technology, Automation and the U.S. Workforce, National Academy of Sciences, *"Information Technology and the U.S. Workforce: Where Are We and Where Do We Go From Here?"* (Washington, DC, National Academy of Sciences, 2017)

Thomas H. Davenport, *"Why Trump Doesn't Tweet About Automation,"* Harvard Business Review, January 12, 2017

Jeff Guo, *"We're so unprepared for the robot apocalypse,"* The Washington Post, March 30, 2017

Dirk Helbing, Bruno Frey, Gerd Gigerenzer, Ernst Hafen, Mishael Hagner, Yvonne Hofstetter, Jeroen van den Hoven, Roberto Zicari and Andrej Zwitter, *"Will Democracy Survive Big Data and Artificial Intelligence?"* Scientific American, February 25, 2017

Raffi Khatchadourian, *"The Domsday Invention: Will artificial intelligence bring us utopia or destruction?"* The New Yorker, November 23, 2015

Paul Krugman, *"Robot Geometry (Very Wonkish),"* New York Times, March 20, 2017

Robin Kwong, Joanna Kao, Claire Manibog and Toyoki Nakanishi, *"Can a Robot Do Your Job,"* Financial Times, April 7, 2017

Gideon Lewis-Krauss, *"The Great A.I. Awakening: How Google used artificial intelligence to transform Google Translate, one of its more popular services—and how machine learning is poised to reinvent computing itself,"* New York Time Magazine, December 14, 2016

Steve Lohr, *"A.I. Is Doing Legal Work. But It Won't Replace Lawyers. Yet."* New York Times, March 19, 2017

James Manyika, Michael Chui, Mehdi Miremadi, Jacques Bughin, Katy George, Paul Willmot and Martin Dewhurst, *"Harnessing Automation for a Future That Works,"* McKinsey&Company, January 2017

James Pellegrino and Margaret Hilton, editors, *"Education for Life and Work: Developing Transferable Knowledge and Skills in the 21st Century,"* National Academy of Sciences, 2012

David Rotman, *"How Technology is Destroying Jobs,"* MIT Technology Review June 12, 2013

Jeff Schwartz, Laurence Collins, Heather Stockton, Darryl Wagner and Brett Walsh, *"The future of work: The augmented workforce 2017 Global Human Capital Trends,"* Deloitte, February 28, 2017

Lawrence H. Summers, *"Lawrence H. Summers on the Economic Challenge of the Future: Jobs,"* The Wall Street Journal, July 7, 2014

Peter Temin, *"The Vanishing Class: Prejudice and Power in a Dual Economy"* (Cambridge, MIT Press, 2017)



© June 2017.