



INDUSTRIES
OF THE
FUTURE

Artificial Intelligence and Machine Learning: Opportunity and Challenge

Discussion paper

Prepared by the Lisbon Council for the High-Level Roundtable on Artificial Intelligence and Machine Learning

British physicist Stephen Hawking says “the development of full artificial intelligence could spell the end of the human race.”¹ Elon Musk, creator of Tesla Motors, adds that artificial intelligence could be like “summoning the demon” that threatens all existence, especially if actions are not taken to design systems that can remain under human control.²

But is it true? Is artificial intelligence really the beginning of the end of civilisation as we know it? Or are these views merely the hype-filled concerns of innovative minds raised on plentiful science fiction, where the evil robot run amok has been enshrined in such cultural classics as *Terminator*, *2001: A Space Odyssey*, Kurt Vonnegut’s 1952 classic *Player Piano* and even Mary Shelley’s *Frankenstein*?

In the end, artificial intelligence is both more and less than it is cracked up to be. On the “less” side, there is ample evidence that – in contrast to what famed writer Ray Kurzweil has written – the “singularity” may not be so near after all.³

To be sure, machines can gather, analyse and store facts on a scale that human minds don’t come close to replicating. The processing power of computers is also growing exponentially, and the underlying base of machine-readable raw data

‘The rate of technological innovation is accelerating, doubling every decade.’ – Ray Kurzweil

-
- 1 See Interview with Stephen Hawking, BBC, 02 December 2014. Interviewed by Rory Cellan-Jones. Prof Hawking’s views are also explored in some depth in Nick Bostrom, *Superintelligence* (Oxford: Oxford University Press, 2016).
 - 2 Mr Musk made this comment in a series of interviews and speeches in Autumn, 2014. See, for example, Elon Musk, “[One-on-One with Elon Musk](#),” *Massachusetts Institute of Technology Centennial Symposium*, 24 October 2014.
 - 3 Ray Kurzweil, *The Singularity is Near* (New York: Viking Penguin, 2005). According to Wikipedia, “singularity” – also known as “technological singularity” – is the hypothesis that the invention of artificial superintelligence will abruptly trigger runaway technological growth, resulting in unfathomable changes to human civilization. According to this hypothesis, an upgradable intelligent agent (such as a computer running software-based artificial general intelligence) would enter a “runaway reaction” of self-improvement cycles, with each new and more intelligent generation appearing more and more rapidly, causing an intelligence explosion and resulting in a powerful superintelligence that would, qualitatively, far surpass all human intelligence. For more, [see the entry on technological singularity, Wikipedia](#).

about our health, transport and hobbies is expanding fast as well. But machines are still subject to running on algorithms that humans write for them, including the ones that give them some capacity to learn. A computer, for example, can be programmed to beat the leading Jeopardy players, or to play chess better than the world's best. But that same machine would have a hard time writing a paper like this, or managing a government department, or painting a painting, or chairing a high-level roundtable like the one where this paper will be discussed. They can be programmed to do these things, but who, ultimately, is in charge of that? A computer is still a machine even as its pattern-recognising capabilities are advancing to uber-human levels and the Internet makes

'More than 140 startups with artificial intelligence as a core part of their product raised €904 million in funding in the third quarter of 2016, making it the second-highest quarter for AI-driven funding, trailing only the second quarter of 2016.' – CB Insights

the sum of all human intelligence searchable with a single click of a mouse. The key point is, you still need a human to make full sense of the search results, as we will see below. And just who is sitting there clicking the mouse?

But there is ample evidence that, within this caveat, machines and learning-enabled computers could be poised to deliver dramatic lifestyle improvements to the economies, cultures, societies and individuals that are smart enough to embrace them. Although these developments seem destined to take place

alongside human intelligence, not to replace it. And, while there is evidence that robots will be able to do some jobs that people do today – including some routine administrative tasks, like data entry, and some non-routine tasks, like translating – the reality is that the successful economy of tomorrow will simply see more people deployed in a different way. There will be fewer and fewer manual jobs – especially in advanced economies. But there will be plenty of new jobs that will require working and interacting with the machines, using their analysis as a basis for sounder judgment, working to build on the unique insight that high-level processing can bring, and building and maintaining the systems which will deliver the new insights and store the everyday life-rooted data which these machines will mine.

Take the medical profession.⁴ Already, computers can diagnose cancer or potential disease from scans with much greater precision than a human. But what happens next? A cancer patient is not going to want to talk to an AI-enabled assistant about his treatment – he will want a real doctor there to tell him what to do next. The point is, the advance will come not when the machine replaces the doctor (which by dint of not being human it cannot effectively do), but when the machine and the doctor find grounds for working together, each borrowing from and drawing on the other's capabilities.⁵ Not coincidentally, the roll-out on a wider scale of artificial intelligence will lead to a higher, better level of service, which the world's most advanced economies are highly advised to embrace.

4 See Robert D. Atkinson, "[It's Going to Kill Us!](#) and Other Myths About the Future of Artificial Intelligence," *Information Technology and Innovation Foundation*, 06 June 2016.

5 To be fair, this vision is close to the version of "singularity" proposed by Ray Kurzweil in his seminal book.

There is ample evidence that human beings will not only retain a role in this new economy, but will continue to have the crucial one. Put simply, machines can't think; they can only calculate. And left to their own devices, machines can just as easily lead us to the wrong calculus as to the right one. A perfect example is the dramatic situation in the recent United States presidential election. Setting aside the question of the outcome, there is ample evidence that social media – using news algorithms – played a possibly decisive role in the spread of false news and demonstrably fake stories. To be sure, the system had been rigged by humans, many of whom were writing the fake news which subsequently spread.⁶ But the computers weren't able to spot the fault. And the result is a travesty. A recent study shows that in the crucial final three months of the U.S. presidential election, fake news sites generated more user engagement than the top news stories from major sites such as *The New York Times*, *Washington Post*, *Huffington Post* and *NBC News*.⁷ The fake sites drew 8,711,000 shares in the final three months of the election period; the best-performing news stories from the nation's 19 leading news agencies drew only 7,367,000 shares in that time.

'The future is already here; it's just not very evenly distributed.' – William Gibson

While leading technology and social-media companies are still preparing their response, the lesson is reasonably clear. Artificial intelligence makes us very powerful.⁸ But we still need humans to ensure that power is used for social utility, and not just to magnify the message of deceit (news forgers) over the will of the well-intentioned (voters). In the past, editors provided that function; and, while social media can rightly claim that it has led to a “golden age” of news consumption (with more stories from more points of view available to more people than ever before), it now

'We do not refer to bulldozers or tractors as superhuman because they can lift 100 times more weight than a human. They are tools that serve our needs. The same is true of artificial intelligence.' – Robert D. Atkinson

seems likely that some form of human editor or editing will need to be involved in the process again.⁹ United States Supreme Court Justice Potter Stewart once famously observed, “I can't define pornography, but I know it when I see it.” The point is we still need human beings to make judgments in key instances. And the best system is not one where a machine decides for a human, but where artificial intelligence is used to enhance the very real human processes, where the core function of our

society continues to reside. This goes for interference in algorithms, as well. The truth may lie somewhere in Potter Stewart's statement. A line may be hard to define, but ultimately it is not so difficult for human intelligence to draw.¹⁰ We do recognise pornography when we see it, even if we can't define it to a level where the law has an easy time regulating it.

6 Andrew Higgins, Mike McIntire and Gabriel J.X. Dance, [“Inside a Fake News Sausage Factory: ‘This is All About Income.’”](#) *The New York Times*, 25 November 2016.

7 Craig Silverman, [“This Analysis Shows How Fake Election News Stories Outperformed Real News on Facebook,”](#) *Buzzfeed*, 16 November 2016.

8 Rob Atkinson compares artificial intelligence to a bulldozer. “We do not refer to bulldozers or tractors as superhuman because they can lift 100 times more weight than a human. They are tools that serve our needs and the same is true of AI.” See Atkinson, op. cit.

9 The “golden age” quotation is from Reuters Institute for the Study of Journalism, [Digital News Report 2016](#) (Oxford: Reuters Institute, 2016).

10 See especially the recent discussion at Facebook over the iconic Vietnam war photograph of a naked child fleeing a napalm attack. Facebook algorithms – using the only criteria they knew – had blocked the photo because it fell afoul of nudity restrictions. Ultimately, the policy was overturned by a human decision. Mark Scott and Mike Isaac, [“Facebook Restores Iconic Vietnam War Photo It Censored for Nudity,”](#) *The New York Times*, 09 September 2016.

Ultimately, a machine cannot make a qualitative decision. A judge or jury can. But the key point is their reasoning can be enhanced and made better through better access to legal documents, case precedent and other machine-assisted learnings. We still need judges. But perhaps what we really need are more judges who can use and work with the new tools for better social outcomes. This has been a key finding in ongoing Lisbon Council research in several areas, including eGovernment, where we have found that the new technology offers prospects of dramatic service improvement to human-led agencies clever enough to integrate the new medium into their work flow.¹¹ It is a catalyst for more and better services.

So what then are the implications for Europe? We see three key points:

1. Machines and machine learning will raise productivity, and productivity growth, which in the words of economist Bart van Ark, is the only source of sustainable long-term growth in an advanced economy.¹² Economies that turn their back on this important new economic input – or regulate it in ways that slow down progress, or prevent citizens in one country from getting the gains available in others – will suffer the burden of relatively lower living standards. Oddly, this outcome doesn't seem as immediately undesirable to many in the age of "post-truth" politics. The British recently took decisions which economists almost unanimously believe will lead to lower living standards for their island. In that climate, we must make sure that we do two things: 1) we must be clear and unequivocal about what will and will not lead to higher living standards, and 2) we must make the point stick in public discourse, using every means at our disposal to counter fake arguments and false prophets. This point has implications for the debate on artificial intelligence. We need to make sure the facts are fairly and accurately presented. And we must avoid scare mongering, especially in areas where there is so much social good to be had and so little scare to monger.

'We must avoid scare mongering, especially in areas where there is so much social good to be had and so little scare to monger.'

'It is on the basis of better, deeper analytics that so much is possible, including the improvements in cancer diagnostics.'

2. The key input for artificial intelligence is data. It is on the basis of better, deeper analytics that so much is possible, including the improvements in cancer diagnostics mentioned above. But European data policy is unevenly applied and sometimes overly restrictive. The European Court of Justice recently threw European data-protection policy wide open, ruling against the European Commission's authority to regulate in this field and challenging local data protection supervisors to challenge the European policy as often as they saw possible. What's more, the forthcoming General Data Protection Regulation

is falling victim to "gold plating," with more and more local requirements being written into a regulation that was intended to be European. If unchecked, this will lead to diverging data

¹¹ Sergey Filippov, [Government of the Future: How Digital Technology Will Change the Way We Live, Work and Govern](#) (Brussels and London: the Lisbon Council and Nesta, 2015).

¹² See, *inter alia*, Bart van Ark, [Productivity and Digitalisation in Europe: Paving the Road to Faster Growth](#) (Brussels: The Lisbon Council, 2014).

protection standards throughout the EU despite the regulation's promise of establishing one set of rules. The result is a chaotic minefield of problems in an area that – most economists agree – will be crucial in the next phases of advanced-economy development. Effective artificial intelligence needs lots of data to mine and analyse. And Europe urgently needs a data policy that serves two key functions: 1) it must protect consumers and enforce privacy standards according to social and legal conventions, and 2) it should enable innovation in this key field. Those objectives ought to be attainable with the right balance and the right attitude to the problem. But it's not at all clear that the existing legal assumptions and attitudes behind these regulations are the right ones. One answer might be to reverse engineer a solution: what are the challenges of a data-driven economy, including the prospects for lifestyle enhancement and the need for credible privacy protection? How can we conceive of an effective European economy – and work our way back to the regulations we do (or don't) need today?

'Managed well – and embraced enthusiastically – artificial intelligence holds out the possibility of giving Europeans a way of leveraging their traditional strengths.'

3. Machine and machine-based learning are areas where Europe could conceivably be very strong. In past decades, Americans have held the edge in software, particularly in developing companies capable of taking cutting-edge software and software-based services to market. But Europeans have always excelled in high-end manufacturing – the merger of machines and software – making products that command a higher price and greater share in global markets. Managed well – and embraced enthusiastically – artificial intelligence holds out the possibility of giving Europeans a way of leveraging their traditional strengths. If the new technology can be successfully and seamlessly integrated into high-end products, Europe will have an easier time preserving its global dominance in high-end manufacturing. This is not an idle point, either. Europe needs success in this field. It would be counterproductive if all that European regulators saw here was a threat. Managed well, it's a colossal opportunity.

'We still need judges. But perhaps what we really need are more judges who can use and work with the new tools for better social outcomes.'

And, finally, there is the key problem of diffusion. Most computer scientists are not social scientists. And while some have written long and fascinating books about artificial intelligence – a short bibliography is included in this paper – they often overlook the thorny question of diffusion.

Is it enough that some parts of society – or even other societies – have access to these great advances, while others are so palpably behind? And here may lie the regulator's chief responsibility – to make sure that the benefits of artificial intelligence are available to all. That can and will require an enabling framework, a willingness to let the market develop without intervening to block "problems" that may not exist and an open-mindedness to a future that promises to be so very different from our own.

Key questions for discussion

- What are the prospects and opportunities for artificial intelligence? What great social challenges could be solved? What advances are just around the corner?
- How does the United States regulate this? What are the key artificial intelligence frameworks there? Are they effective? What could Europe learn from them?
- What are the labour market implications of artificial intelligence, especially for less skilled and white-collar workers?
- How can/should our institutes of higher learning be adapted?
- Are we yet spending enough on research and development in this key field? Where would be the most fertile areas for research and/or future public/private initiative?
- What sort of liability regime do we envisage for heavy artificial intelligence-footprint companies like Facebook and Google? Given the extent to which all modern industry relies on machine learning, does the new technology pose questions for existing product liability rules?
- An enormous range of personal data will be collected by artificial intelligence and machine learning as the new technology spreads. How do we ensure personal data is correctly protected according to the new European framework?
- How do we distinguish between private and non-private data? In relation to non-private data, who should own this information? How should ownership be decided between the customers generating it, the companies building the machines on which it will be generated or the organisations analysing and operationalising it?
- What are the biggest risks for cyber-security in relation to artificial intelligence and machine learning?

Further reading

- Armstrong, Stuart. *Smarter Than Us: The Rise of Machine Intelligence* (Berkeley, CA: Machine Intelligence Research Institute, 2014)
- Atkinson, Robert. "It's Going to Kill Us!" and Other Myths About the Future of Artificial Intelligence (Washington, DC: Information Technology and Information Foundation, 2016) <http://www2.itif.org/2016-myths-machine-learning.pdf>
- Bostrom, Nick. *Superintelligence: Paths, Dangers, Strategies* (Oxford: Oxford University Press, 2014)
- Brynjolfsson, Erik, and Andrew McAfee. *Race Against the Machine: How the Digital Revolution Is Accelerating Innovation, Driving Productivity, and Irreversibly Transforming Employment and the Economy* (Lexington, MA: Digital Frontier Press, 2011)
- Castro, Daniel, and Joshua New. *The Promise of Artificial Intelligence* (Brussels and Washington, DC: Centre for Data Innovation, 2016) <https://www.datainnovation.org/2016/10/the-promise-of-artificial-intelligence/>
- Denning, Steve. "The "Jobless Future" Is a Myth," *Forbes*, 04 June 2015 <http://www.forbes.com/sites/stevedenning/2015/06/04/the-robots-are-not-coming/>
- Filippov, Sergey. *Government of the Future: How Digital Technology Will Change the Way We Live, Work and Govern* (Brussels and London: The Lisbon Council and Nesta, 2015) <http://lisboncouncil.net/publication/publication/130-government-of-the-future-.html>
- Ford, Martin. *Rise of the Robots: Technology and the Threat of a Jobless Future* (New York, NY: Basic Books, 2015)
- Hawking, Stephen. *Speech to the Leverhulme Centre for the Future of Intelligence*. <https://www.youtube.com/watch?v=Jjn8dsFTf58>
- Kurzweil, Ray. *The Singularity is Near* (New York: Viking Penguin, 2004)
- Levy, Steven. "How Google is Remaking Itself as a 'Machine Learning First' Company," *Backchannel*, 22 June 2016 <https://backchannel.com/how-google-is-remaking-itself-as-a-machine-learning-first-company-ada63defcb70#.zcxfbx4pl>
- National Science and Technology Council. *Preparing for the Future of Artificial Intelligence* (Washington, DC: Executive Office of the President, 2016) https://www.whitehouse.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/preparing_for_the_future_of_ai.pdf
- Ng, Andrew. "What Artificial Intelligence Can and Can't Do Right Now," *Harvard Business Review*, 09 November 2016 <https://hbr.org/2016/11/what-artificial-intelligence-can-and-cant-do-right-now>
- Study Panel. *Artificial Intelligence and Life in 2030* (Stanford: Stanford University, 2016) https://ai100.stanford.edu/sites/default/files/ai100report10032016fnl_singles.pdf
- Vanian, Jonathan. "Google Artificial Intelligence Whiz Describes Our Sci-Fi Future," *Fortune*, 26 November 2016 <http://fortune.com/2016/11/26/google-artificial-intelligence-jeff-dean/>
- World Economic Forum. *What is the Current State of Artificial Intelligence?* <https://www.youtube.com/watch?v=c3YdaWiQuZY>



About Industries of the Future

The Industries of the Future initiative brings unique, multi-disciplinary collaboration to one of the key questions of our time: how will digital technology transform and revitalise industry itself, bringing new models to the fore and posing new challenges for policymakers and society alike? In a multi-stakeholder environment informed by an advanced think tank perspective, participants examine the opportunities and dilemmas of tomorrow – and the steps we should be taking today. The initiative is led by FTI Consulting, Google and the Lisbon Council.



About FTI Consulting

FTI Consulting combines deep financial, forensic, economic, technology and communications experience with extensive industry expertise throughout the world. Its website is www.fticonsulting.com.



About Google

Google's mission is to organise the world's information and make it universally accessible and useful. Its website is www.google.com.



About the Lisbon Council

The Lisbon Council is a Brussels-based think tank and policy network committed to engaging political leaders and the public at large in a constructive dialogue about 21st century economic and social challenges. Its website is www.lisboncouncil.net.

Published under the editorial responsibility of FTI Consulting, Google and the Lisbon Council asbl. The responsible editor is Paul Hofheinz, president, the Lisbon Council asbl

Lead Author:

Paul Hofheinz

Project Team:

Sarah Elnusairi, Sergey Filippov, Sylwia Giepmans-Stępień, Stéphanie Lepczynski, Elizabeth Miles, Victoria de Posson, Marzena Rembowski, Ondrej Socuvka and Lynette Webb

FTI Consulting
Avenue Marnix 23
1000 Brussels
Belgium

Google
Chaussée d'Etterbeek 176-180
1040 Brussels
Belgium

The Lisbon Council asbl
Résidence Palace
Rue de la Loi 155
1040 Brussels, Belgium