THE IMPACT OF ARTIFICIAL INTELLIGENCE ON EMPLOYMENT

Georgios Petropoulos[1]

Technological development, and in particular digitalisation, has major implications for labour markets. Assessing its impact will be crucial for developing policies that promote efficient labour markets for the benefit of workers, employers and societies as a whole.

Rapid technological progress and innovation can threaten employment. Such a concern is not new but dates back at least to the 1930s, when John Maynard Keynes postulated his 'technological unemployment theory' – technological change causes loss of jobs (Keynes 1937).

Technological innovations can affect employment in two main ways:

- by directly displacing workers from tasks they were previously performing (displacement effect)
- by increasing the demand for labour in industries or jobs that arise or develop due to technological progress (productivity effect).

Autor, Levy and Murnane (2003) stress that technology can replace human labour in routine tasks, whether manual or cognitive, but (as yet) cannot replace human labour in non-routine tasks. Goos and Manning (2007) argue that the impact of technology leads to rising relative demand in well-paid skilled jobs, which typically require non-routine cognitive skills, and rising relative demand in low-paid, least-skilled jobs, which typically require non-routine manual skills.

At the same time, demand for 'middling' jobs, which have typically required routine manual and cognitive skills, will fall. The authors call this process job polarisation. Acemoglu and Autor (2011) found similar results for the US, while Darvas and Wolff (2016) report such developments for a selection of EU countries: France, Germany, Italy, Spain, Sweden and the UK. In all these countries, the number of high-education jobs such as managers, engineers and health professionals is growing, while the number of middle-education jobs (clerks, machine operators, assemblers) is declining. By contrast, the number of low-education service occupations, such as shop workers, which are non-standard and difficult to replace by automation, is growing. A key conclusion is that technology was incorporated into the subset of core job tasks previously performed by middle-skill workers, causing substantial change. The quality of human capital also plays a crucial role. The ability of individuals to use the technological advances for the benefit of their work requires developing particular digital skills through well-designed policies. This underlines the importance of using appropriate instruments to ensure that workers are well prepared to harness the disruptive forces of digital technologies.

In the last decade platforms emerged that contributed to increased connectivity between individuals. For example, using this connectivity, peer providers of durable goods and services can trade online with individuals using collaborative economy platforms. A key common characteristic of collaborative economy models – despite a great deal of variety – is that they provide an economic opportunity for individuals and small enterprises to trade their under-used assets with other individuals through intermediaries that match supply and demand in an efficient way with the help of information technologies. In many cases, this opportunity to individual suppliers is only provided through collaborative platforms, as the supply of goods and services through other channels is subject to licensing and other regulatory barriers. Automation in shopping through ecommerce is another example, with the sector experiencing annual growth of 22% in Europe.[2] The benefits of information technologies increase demand for online retail goods and this in turn leads to an increased overall employment in retail.

However, looking ahead, a new wave of automation and advanced machine-learning techniques is on its way, in which intelligent machines will be increasingly capable of carrying out high-skill and possibly non-routine tasks. Moving from the efficiency gains in online trading to the extensive use of artificial intelligent systems in our industrial production, concerns about the potential displacement of labour emerge. The real question then becomes: which of the two labour market effects – displacement or productivity – will dominate in the artificial intelligence (AI) era?[3]

A first approach to answer this question is to examine the impact of technological breakthroughs on labour markets in previous industrial revolutions (Soete this volume). For example, the introduction of automobiles in daily life led to a decline in horse-related jobs, but new industries also emerged, with a net positive impact on employment. The automobile industry itself grew fast, creating many new jobs, but other sectors also grew because of the growing number of vehicles on the roads, and many new jobs in the motel and fast-food industries arose to serve motorists and truck drivers.

The Economist (2016) reports further case studies that show similar patterns. In general, past industrial revolutions suggest that in the short run the displacement effect may dominate. But in the longer run, when markets and society are fully adapted to major automation shocks, the productivity effect can dominate and have a positive impact on employment.

But how reliable is this approach? Researchers from the McKinsey Global Institute estimate that the disruption of society caused by AI is happening 10 times faster and at 300 times the scale of the industrial revolution of the late 18th and early 19th centuries, and is therefore having roughly 3,000 times the impact (Dobbs, Manyika and Woetzel 2015).

Moreover, the main engine of technological progress in the AI era is the continuous development of deep machine-learning techniques that use the function and complexity of the human brain as a model for design (see Petropoulos 2017b); for relevant definitions and analysis see Box 1. Machines are trained to be intelligent, which can have additional implications for the workforce.

BOX 1. AN INTRODUCTION TO MACHINE LEARNING

Machine learning enables computer programs to acquire knowledge and skills, and even improve their own performance. Big data provides the raw material for machine learning, and offers examples that computer programs can use for 'practise' in order to learn, exercise and ultimately perform their assigned tasks more efficiently.

The idea of intelligent machines arose in the early 20th century. From the beginning, the idea of 'human-like' intelligence was key. Following Vannevar Bush's seminal work from 1945, where he proposed "a system which amplifies people's own knowledge and understanding", Alan Turing asked "Can a machine think?" In his famous 1950 imitation game, Turing proposed a test of a machine's ability to exhibit intelligent behaviour equivalent to that of a human.

In principle, machine learning follows Turing's recommendation of teaching a machine to perform specific tasks as if it were a child. By building a machine with sufficient computational resources, offering training examples from real world data and by designing specific algorithms and tools that define a learning process, rather than specific data manipulations, machines can improve their performance through learning by doing, inferring patterns and checking hypotheses.

At the core of this learning process are artificial neural networks, inspired by the networks of neurons in the human brain. A simple artificial neural network is organised in layers. Data is introduced to the network through an input layer. Then come the hidden multiple layers in which information is processed and finally an output layer where results are released. Each neuron within the network is connected to many others, as both inputs and outputs, but the connections are not equal. They are weighted such that a neuron's different outward connections fire at different levels of input activation. A network with many hidden layers can combine, sort or divide signals by applying different weights to them and passing the result to the next layer. The number of hidden layers demonstrates the ability of the network to detect increasingly subtle features of the input data. The training of the network takes place by adjusting neurons' connection weights, so that the network gives the desired response when presented with particular inputs.

The goal of the neural network is to solve problems in the same way that a hypothesised human brain would, albeit without any 'conscious' codified awareness of the rules and patterns that have been inferred from the data. Modern neural network projects typically work with a few thousand to a few million neural units and millions of connections. They are called deep because of the multiple intermediate hidden layers they have. However, deep neural networks are still several orders of magnitude less complex than the human brain and closer to the computing power of a worm.

Deep neural networks have proven very effective. There are several examples of games and competitions in which machines can now beat humans. By now, machines have topped the best humans at most games traditionally held up as measures of human intellect, including chess (recall for example the 1997 game between IBM's Deep Blue and the champion Garry Kasparov), Scrabble, Othello and Jeopardy! Even in more complex games, machines seem to be quickly improving their performance through their learning process. In March 2016, the AlphaGo computer program from the AI start up DeepMind beat Lee Sedol at a five-game match of Go – the oldest board game, invented in China more than 2,500 years ago. However, many of these machines are programmed to perform specific tasks, narrowing the scope of their operation. Humans remain superior in performing general tasks and using experience acquired in one task to deliver another.

A second approach would be to assess the risk of occupations and tasks to be automated in the next decades because of AI systems. Here the literature has focused on the feasibility of automating existing jobs given current and presumed technological advances (Arnold et al. this volume). Frey and Osborne (2013, 2017) famously claimed that 47% of US occupations were at risk of being automated "over some unspecified number of years, perhaps a decade or two" (Frey and Osborne 2017, 265). Bowles (2014) repeated these calculations for the European labour market, and found that on average 54% of EU jobs are at risk of computerisation. By contrast, Arntz, Gregory and Zierahn (2016, 2017) argue that a major limitation of Frey and Osborne is that they focus on deriving predictions over occupations as being threatened by automation rather than tasks. Their criticism is that in this way Frey and Osborne overestimate the automation risks. By using information on task content of jobs at the individual level they conclude that only 9% of US jobs are potentially automatable.

These studies can be viewed as feasibility tests on the potential impact of AI and focus on the displacement effect of automation. Assessing the impact of the productivity effect – the potential for new machines to increase employment – is much more challenging. Bessen's (2017) empirical research found that computer technology is associated with job growth that is particularly observable in non-manufacturing industries. At the same time there are potential sector spillover effects: as Acemoglu and Restrepo (2016) illustrate in their theoretical model, the aggregate labour market impacts of new technologies depend not only on the industries in which they operate, but also on adjustment in other parts of the economy. For example, other sectors and occupations might expand to absorb the labour freed from the tasks that are now performed by machines.

That would require adopting an equilibrium approach because what is technologically feasible does not necessarily correspond to the equilibrium impact of automation on employment and wages. For example, we need to take into account that firms' market strategies and investments are endogenous to technology shocks: Even if the presumed technological advances materialise, there is no guarantee that firms would choose to automate; that would depend on the costs of substituting machines for labour and how much wages change in response to this threat.

That brings us to the third approach of assessing the impact of AI on employment. A common characteristic of most of research papers that are moving towards this equilibrium approach is that all focus on one automated technology, the industrial robots and their impact on employment. This is because of the existence of good quality data on the penetration of industrial robots in the main industries in major economies around the world.

An industrial robot is defined as "an automatically controlled, reprogrammable, multipurpose manipulator programmable in three or more axes, which can be either fixed in place or mobile for use in industrial automation applications" (International Federation of Robotics 2016). Following this definition, a classification test would have required a clear answer to the following three questions:

- Does it have multiple purposes?
- Can it be reprogrammed to perform another task?
- Does it require a human control for performing its task?

While our coffee machine or the elevator at our home building does not pass this classification test, fully autonomous machines that do not need a human operator and that can be programmed to perform several manual tasks such as welding, painting, assembling, handling materials or packaging are classified as industrial robots.

Figure 7.1 presents the number of operational industrial robots per thousands of workers in China, the EU and the US. The EU so far has been the region with the most robots in operation, followed by the US while China is behind.

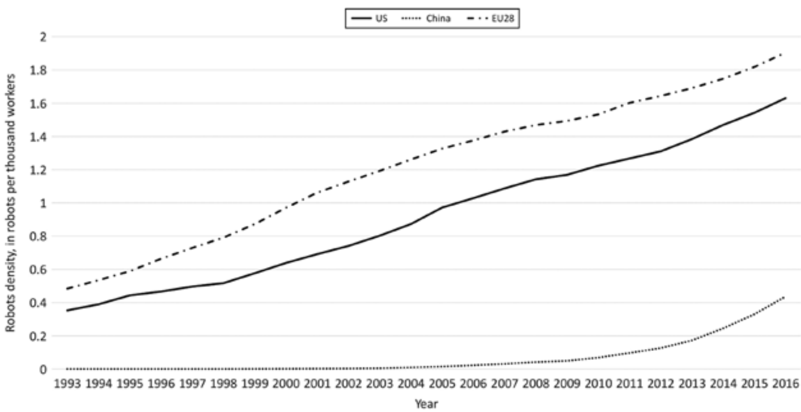

Figure 7.1 Robot density in China, EU and US. *Source*: Data from International Labor Organisation (2017), IFR (2016).

Figure 7.2 shows how the operational industrial robots per thousands of workers are distributed in different sectors in EU countries. So far, the EU automotive industry has introduced by far the most industrial robots in its production process, followed by the plastic and chemicals sector.

Graetz and Michaels (2015) estimate that between 1990 and 2005 the price of industrial robots in six major developed economies fell by approximately one-half or one-fifth if we adjust for the quality of robots. Moreover, between 1993 and 2007, the stock of robots per million hours worked increased by more than 150%, from 0.58 to 1.48, in 17 countries of the sample, leading to significant productivity gains. The study also finds that in these countries increased use of robots per hour worked from 1993 to 2007 raised the annual

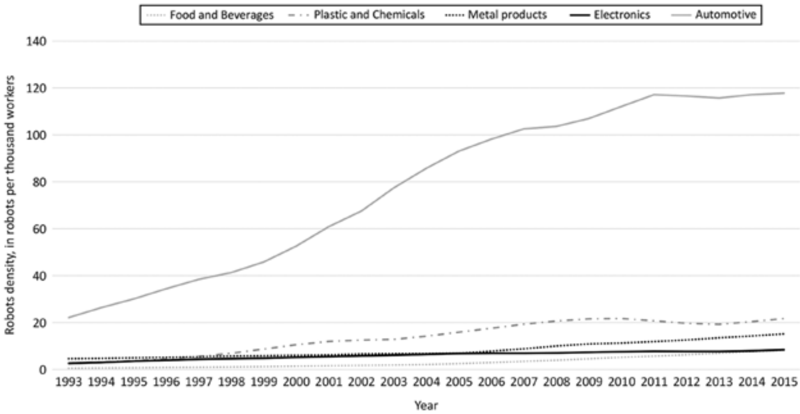

Figure 7.2 Robot density in several industries in Europe. *Source*: Datafrom EUKLEMS (2017), IFR (2016).

growth of labour productivity by about 0.37 percentage points. When considering an industry-country panel specification, they find that robots appear to reduce the share of hours worked by low-skilled workers relative to middle-skilled and high-skilled workers, they do not polarise the labour market, but appear to hurt the relative position of low-skilled workers rather than middle-skilled ones. Nevertheless, the use of robots per hour worked appears to boost total factor productivity and average wages. No significant impact on labour shares is found.

In a more recent study, Acemoglu and Restrepo (2017) used data in the post-1990 era to show that 1 additional robot per 1,000 workers reduces the US employment-to-population ratio by 0.18–0.34% and wages by 0.25–0.5%. When interpreting these results we should not forget that there are still few industrial robots in the US economy; if the spread of robots proceeds over the next two decades as expected by experts such as Brynjolfsson, McAfee and Ford, its aggregate implications for employment will be much larger (Brynjolfsson and McAfee 2012; Ford 2015). The novel element of their study is that they adopt a more regional approach than the industry-country panel approach of Graetz and Michaels (2015). As the labour force competes with robots for production, they exploit the heterogeneity in both local labour distribution across industries and national change in the use of robots to refine their results. They can therefore estimate the impact of industrial robots' penetration in local labour markets. Their negative result suggests that the displacement effect dominates the productivity effect of operation industrial robots. In addition, positive spillover effects are very modest. The employment effects of robots are most pronounced in manufacturing, particularly in industries most exposed to robots; in routine manual, blue-collar, assembly and related occupations; and for workers without a college education.

Dauth et al. (2017) repeat the empirical exercise of Acemoglu and Restrepo (2017) for Germany but they do not find any significant negative impact of robots. While industrial robots have a negative impact on employment in the German manufacturing sector, there is a positive and significant spillover effect as labour in the non-manufacturing sectors increases and overall counterbalances the negative effect.

The focus of these studies is on the impact of industrial robots on employment so far, without making any predictions for the future. These predictions would require the imposition of specific assumptions whose validity cannot be assessed with certainty.

While this allows for a more reliable assessment of the impact, we should keep in mind that the era of AI is in its early stages and the penetration of robots in our economy and industrial production is expected to significantly rise as a consequent of the rapid, ongoing technological progress. This suggests that existing studies using this third approach are able to capture only the onset of the AI era and not its full deployment. If indeed short-run and long-run effects are not in the same direction, these studies may only be able to capture some parts of the short-run effects.

Industrial robots are just one of the AI technologies that have been developed. At the forefront of the fourth industrial revolution will be a connected framework of machines that communicate with each other. Such connectivity is expected to be a major step forward, increasing the efficiency gains in AI markets and services. Completing a full economic framework for the impact of AI on labour markets before these new developments are deployed is a difficult task.

These future-facing studies do not reach a consensus over the potential impact of automation on labour markets. The fact that it is difficult to predict the exact impact of AI makes it complex to frame a policy response. But some society-level reaction is surely needed. It is therefore necessary to initiate an open consultation of all involved parties, to define our approach towards the AI era. This process should have several steps:

1. Ensure that society, and particularly policymakers, politicians and business leaders, understands what AI is and its potential for modern economies.
2. Define a framework of rules for the operation of machines and AI automated systems. These must go far beyond Asimov's famous three laws of robotics. The Civil Law Rules on Robotics proposed by the European Parliament can also motivate social dialogue about issues related to liability, safety, security and privacy in the coming AI era. Tegmark (2017) identifies numerous challenges on these matters, which should be addressed adequately. Adopting clear rules based on a good understanding of this new era could make the transition easier and mitigate potential concerns. However, adopting rules without good understanding and knowledge of how this new technology will be implemented (first step) would be counterproductive.
3. Design and implement those policies that will help us to accommodate new technology possibilities. Education and training programmes should be carefully redesigned so that they provide the right qualifications for workers to interact and work efficiently alongside machines and boost relevant digital skills. This might reduce potential displacement concerns as jobs typically consist of a number of distinct but interrelated tasks. In most cases, only some of these tasks are likely to be suitable for automation. By preparing human labour to interact effectively and efficiently with machines, we can maximise the productivity gains from the

interrelated tasks. That could potentially lead to the development of new jobs or occupations that will result from this cooperation and the advancements of the technology. Initiatives to prepare effectively human labour for this new era will require the close interaction of authorities and institutions with major technological firms which have both the knowhow and the capacity to contribute to the training. Improved instruments for job search assistance and job reallocation could also be beneficial and would mitigate concerns associated with the displacement effect.

However, we should not rush into a response (see Atkinson this volume). The time for policy will come, but at the moment we are still in the early stages of understanding the potential of AI and the various ways it might impact our economy. To deepen this understanding, we should promote further social dialogue among all the involved parties (researchers, policymakers, industry representatives and trade unions, politicians and so on). This is a vital first step to better grasp the challenges and opportunities of this new industrial revolution. And although we should not rush to conclusions, we should not adopt a passive attitude. We must act swiftly to assess and understand the implications of AI. The speed with which technology advances may introduce disruptive forces in the market earlier than some people expect.

NOTES

1. This chapter is an updated version of my article 'Do We Understand the Impact of Artificial Intelligence on Employment?' published by Bruegel (Petropoulos 2017a). The superb research assistance by Nicolas Moës is gratefully acknowledged.
2. See Marcus and Petropoulos (2016) for further statistics and discussion.
3. AI refers to intelligence exhibited by machines. Hence, the AI era refers to that period in time in which machines equipped with deep learning techniques that are based on neural network architecture (see Box 1) will be able to perform tasks that require some form of intelligence, in an automatic way and without requiring human intervention.

REFERENCES

Acemoglu, D. and D. Autor (2011), 'Skills, Tasks and Technologies: Implications for Employment and Earnings', in O. Ashenfelter and D. Card (eds), *Handbook of Labor Economics*, vol. 4, Amsterdam: Elsevier.

Acemoglu, D. and P. Restrepo (2016), *The Race Between Machine and Man: Implications of Technology for Growth, Factor Shares and Employment*, NBER Working Paper 22252, http://www.nber.org/papers/w22252.pdf.

Acemoglu, D. and P. Restrepo (2017), *Robots and Jobs: Evidence from US Labor Markets*, NBER Working Paper 23285, http://www.nber.org/papers/w23285.pdf.

Arntz, M., T. Gregory and U. Zierahn (2016), *The Risk of Automation for Jobs in OECD Countries: a Comparative Analysis*, Social, Employment and Migration Working Paper 189, Paris: Organisation for Economic Co-operation and Development.

Arntz, M., T. Gregory and U. Zierahn (2017), 'Revisiting the Risk of Automation', *Economics Letters*, 159: 157–60.

Autor, D., F. Levy and R. J. Murnane (2003), 'The Skill Content of Recent Technological Change: an Empirical Exploration', *Quarterly Journal of Economics*, 118(4): 1279–333, http://www.jstor.org/stable/pdf/25053940.pdf.

Bessen, J. E. (2017), *Automation and Jobs: When Technology Boosts Employment*, Boston University School of Law, Law and Economics Research Paper 17-09.

Bowles, J. (2014), 'The Computerisation of European Jobs', blog, 24 July, Bruegel, http://bruegel.org/2014/07/the-computerisation-of-european-jobs/.

Brynjolfsson, E. and A. McAfee (2012), *Race Against the Machine: How the Digital Revolution is Accelerating Innovation, Driving Productivity, and Irreversibly Transforming Employment and the Economy*, Lexington, Massachusetts: Digital Frontier Press.

Bush V. (July 1945), 'As We May Think', The Atlantic https://www.theatlantic.com/magazine/archive/1945/07/as-we-may-think/303881/

Darvas, Z. and G. B. Wolff (2016), *An Anatomy of Inclusive Growth in Europe*, Brussels: Bruegel, http://bruegel.org/2016/10/an-anatomy-of-inclusive-growth-in-europe/.

Dauth, W., S. Findeisen, J. Südekum and N. Woessner (2017), *German Robots: the Impact of Industrial Robots on Workers*, CEPR Discussion Paper DP12306.

Dobbs, R., J. Manyika and J. Woetzel (2015), *The Four Global Forces Breaking all the Trends*, London, San Francisco, Shanghai: McKinsey Global Institute.

Ford, M. (2015), *The Rise of the Robots: Technology and the Threat of Mass Unemployment*, London: Oneworld Publications.

Frey, C. and M. Osborne (2013), *The Future of Employment: How Susceptible Are Jobs To Computerisation?*, Oxford: Oxford Martin School, https://www.oxfordmartin.ox.ac.uk/downloads/academic/The_Future_of_Employment.pdf.

Frey, C. B. and M. A. Osborne (2017), 'The Future of Employment: How Susceptible Are Jobs To Computerisation?', *Technological Forecasting and Social Change*, 114: 254–80.

Goos, M. and A. Manning (2007), 'Lousy and Lovely Jobs: the Rising Polarization of Work in Britain', *The Review of Economics and Statistics*, 89(1): 118–33.

Graetz, G. and G. Michaels (2015), *Robots at Work*, CEPS Discussion Paper 1335, London: London School of Economics, http://cep.lse.ac.uk/pubs/download/dp1335.pdf.

International Federation of Robotics (2016), *Industrial Robots: Definition and Classification.* Frankfurt am Main: International Federation of Robotics.

Keynes, J. M. (1937), 'The General Theory of Employment', *Quarterly Journal of Economics*, 51(2): 209–23.

Marcus, J. S. and G. Petropoulos (2016), 'Cross-Border Parcel Delivery Prices: Intuitions drawn from the World of Telecommunications', paper for the 27th European Regional Conference of the International Telecommunications Society (ITS).

Petropoulos, G. (2017a), 'Do We Understand the Impact of Artificial Intelligence on Employment?', Bruegel blogpost, http://bruegel.org/2017/04/do-we-understand-the-impact-of-artificial-intelligence-on-employment/

Petropoulos, G. (2017b), 'Machines that Learn to Do, and Do to Learn: What is Artificial Intelligence?', Bruegel, 27 April, http://bruegel.org/2017/04/machines-that-learn-to-do-and-do-to-learn-what-is-artificial-intelligence/.

Tegmark, Max (2017), *Life 3.0. Being Human in the Age of Artificial Intelligence*, New York: Knopf.

The Economist (2016), 'Automation and Anxiety', *The Economist*, 25 June, http://www.economist.com/news/special-report/21700758-will-smarter-machines-cause-mass-unemployment-automation-and-anxiety.