Robots, Jobs, Taxes, and Responsibilities

Luciano Floridi[1]

AI and robots continue to make news. Alarmist headlines used to be about some kind of Terminator developing in the future to dominate and enslave us, like an inferior species. They are now about tireless machines that, like enslaved persons, will make us redundant, replacing and outperforming us more efficiently and cheaply than we can ever be. This master-slave dialectics is not science fiction. On the 16th of February 2017, the plenary session of the European Parliament voted in favour[1] of a resolution[2] to create a new ethical-legal framework according to which robots may qualify as "electronic persons". The Commission does not have to follow the Parliament's recommendations but, if it refuses, it will have to explain why. The following day, on the 17th of February, in an interview with *Quartz*,[3] Bill Gates, Microsoft co-founder, suggested that there should be a tax on robots.[4]

Regulating robots is a very reasonable idea. Today, we live *onlife*, spending increasing amount of time inside the infosphere. In this digital ocean, robots are the real natives: we scuba dive, they are like fish. So robots of all kinds are going to multiply and proliferate, making the infosphere even more their own space. These smart, autonomous, and social agents perform an increasing number of tasks better than we can. Some of them are already among us. Others are discernible on the horizon, while later generations are still unforeseeable. The solutions that have already arrived come in soft forms, such as apps, webbots, algorithms, and software of all kinds; and hard forms, such as robots, house appliances, personal assistants, smart watches, and other gadgets. In health care, for example, robots and AI solutions are joining nurses, doctors, social workers, technicians, and experts, such as radiologists, by helping perform functions that, just a few years ago, were considered off-

---

✉ Luciano Floridi
luciano.floridi@oii.ox.ac.uk

[1] Oxford Internet Institute, University of Oxford, 1 St Giles, Oxford OX1 3JS, UK

limits for technological disruption: cataloguing images, suggesting diagnosis, monitoring and even moving patients, interpreting radiographies, controlling insulin pumps, extracting new medical information from huge data sets, and so forth. Many trivial, routine tasks will be performed automatically either by AI or by people aided by AI. This is good news. We need AI to deal with increasing levels of complexity and difficulty. With an analogy, we need to remember that the best chess player is neither a human nor a computer, but a human using a computer.

While we can only guess at the scale of the coming disruption, everybody expects it to be profound. Any job in which people serve as menial interfaces—e.g. adjusting the dose of a medication for a patient—is now at risk. Yet new jobs will appear because we will need to manage and coordinate AI solutions. For example, someone will need to ensure that the data collected by insulin pumps and by smart apps are properly combined in order to improve the health care provided and the technologies of the future. What is more, many tasks will not be cost-effective for AI applications. The world never changes at the same pace. In some places, nurses will be irreplaceable for many routine tasks while in others they may coordinate and direct semi-autonomous robots through smart tablets and apps. And some old jobs will survive, even when a machine is doing most of the work: a doctor who delegates some routine tasks to a smart digital assistant will simply have more time to focus on other things, such as prevention. Jobs that were economically not viable until yesterday will become available. Finally, other tasks will be delegated back to us—the patients—to perform them as users, such as testing for blood pressure, something trivial and routine in many countries but still impossible in others.

Another source of uncertainty concerns the point at which AI will no longer be controlled only by a guild of scientists, technicians, and managers. Still relying on the health care example, what will happen when AI becomes "democratised" and a "digital doctor" is available to millions of people on their smartphones or some other device? As Elena Bonfiglioli and Mathias Ekman recently wrote[5]: "As you think innovation in health, you want to think about how to scale the adoption of systems of intelligence making them accessible in more intuitive ways. The vision of AI as "conversations" will empower intelligent health experiences that mirror the way people collaborate and interact with one another, and the way machines proactively understand our intent. […] Systems of intelligence will endemically transform the way we innovate for improved cancer outcomes, the way we optimise clinical and operational processes, and the way we think and do prevention. So, what if people across the healthcare continuum could collaborate and use machine learning to come up with ways to catch cancer earlier and improve outcomes for patients?".

We should investigate how we are going to socialise such systems of intelligence and how we shall best adopt them and adapt to them, from an ethical perspective, because many solutions are far from inevitable, and some may be preferable to others and should be privileged. There is no dystopian science-fiction scenario. *Brave New World* is not coming to life, and the Terminator is not lurking just beyond the horizon, either. There is a good chance that Satya Nadella, Microsoft CEO, may be right when he remarked: "humans and machines will work together – not against one another.

---

[5] Bonfiglioli Elena and Ekman Mathias. 2016. "Innovation race for improved cancer outcomes" https://enterprise.microsoft.com/en-us/industries/health/innovation-race-for-improved-cancer-outcomes/

Computers may win at games, but imagine what's possible when human and machine work together to solve society's greatest challenges like beating disease, ignorance, and poverty."[6] But there are of course risks and challenges in how we shall develop and socialise AI systems and we should tackle them now, to ensure that individual and social benefits are maximised. Quoting Nadella once more: "The most critical next step in our pursuit of A.I. is to agree on an ethical and empathic framework for its design".

Add machine learning to artificial intelligence and robotics, mix these ingredients with the Internet, the Web, smart phones and apps, cloud computing, big data, and the Internet of Things, and it becomes obvious that there is no time to waste. We are laying down the foundations of the mature information societies of the near future, so we need new ethical solutions for the infosphere, to determine which forms of artificial agency and interactions we like to see flourishing in it. Against this background, one can look at the normative initiative taken by the European Parliament or the debate that has followed Gate's suggestion with a mixed sense of excitement for the aspiration but disappointment for the implementation. For there is too much fantasy about speculative scenarios and too little imagination in designing realistic solutions that could work well. Consider two key issues: *jobs* and *responsibilities.*

Robots replace human workers. Retraining unemployed people was never easy, but it is more challenging now that technological disruption is spreading so rapidly, widely, and unpredictably. Today, a bus driver replaced by a driverless bus is unlikely to become a web master, not least because even that job is at risk of automation. There will be many new forms of employment in other corners of the infosphere. Think of how many people have opened virtual shops on eBay. But these will require new and different skills. So more education and a universal basic income may be needed to mitigate the impact of robotics on the labour market, while ensuring a more equitable redistribution of its economic benefits. This means that society will need more resources. Unfortunately, robots do not pay taxes. And it is unlikely that more profitable companies may pay enough more taxes to compensate for the loss of revenues. So robots cause a higher demand for taxpayers' money but also a lower supply of it. The problem is exacerbated by the fact that people with low income purchase cheap goods, those produced more efficiently by increasingly roboticised processes. How can one get out of this tailspin? The report correctly identifies the problem. But its original recommendation[7] of a robotax on companies that employ robots may be unfeasible—for what exactly counts as a robot, if you need to pay a tax on it?—and counterproductive, for a robotax would disincentive innovation. The final text[8] approved by the European Parliament shuns the recommendation but does not offer an alternative solution to the revenue problem.

Consider next the allocation of responsibilities. If a robot breaks the window of my neighbour, who is responsible? The company who produced it, the shop who sold it, I the owner, or the robot itself, if the robot has become completely autonomous through a learning process and is now capable of intelligent-looking actions? In this case too, the

[6] Nadella Satya. 2016. "The Partnership of the Future", *Slate*, http://www.slate.com/blogs/future_tense/2017/01/11/did_a_federal_surveillance_court_really_reject_an_application_to_monitor.html

[7] http://www.europarl.europa.eu/sides/getDoc.do?type=REPORT&mode=XML&reference=A8-2017-0005&language=EN

[8] http://www.europarl.europa.eu/sides/getDoc.do?pubRef=-//EP//NONSGML+TA+P8-TA-2017-0051+0+DOC+PDF+V0//EN

report identifies the issue. It rightly recommends forms of risk management (insurance and compensation). But it also suggests the creation of a "specific legal status" for more advanced robots, as "electronic persons responsible for making good any damage they may cause". This has been approved in the final document. As a result, we may see a future in which companies do not pay a robotax and are not even liable for some kinds of robots. This is probably a mistake. There is no need to adopt science fiction solutions to solve practical problems of legal liability with which jurisprudence has been dealing successfully for a long time. If robots become one day as good as human agents—think of Droids in Star Wars—we may adapt rules as old as Roman law, according to which the owner of an enslaved person was responsible for any damage caused by that person (*respondeat superior*). As the Romans already knew, attributing some kind of legal personality to robots would deresponsabilise those who should control them. Not to speak of the counterintuitive attribution of rights. For example, do robots as "electronic persons" have the right to own the data they produce (machine-generate data)? Should they be "liberated"? It may be fun to speculate about such questions, but it is also distracting and irresponsible, given the pressing issues we have at hand. The point is not to decide whether robots will qualify some day as a kind of persons, but to realise that we are stuck within the wrong conceptual framework. The digital is forcing us to rethink new solutions for new forms of agency. While doing so we must keep in mind that the debate is not about robots but about us, who will have to live with them, and about the kind of infosphere and societies we want to create. We need less science fiction and more philosophy.[9]

References

Floridi, L. (Ed.) (2008). *Philosophy of computing and information: 5 questions*. Automatic Press/VIP.
Floridi, L. (2012). *The road to the philosophy of information, Luciano Floridi's philosophy of technology* (pp. 245–271). New York: Springer.
Floridi, L. (2013). Distributed morality in an information society. *Science and Engineering Ethics, 19*(3), 727–743.
Pagallo, U. (2013). *The laws of robots: crimes, contracts, and torts*. Dordrecht: Springer.

---

[9] On the debate about a normative framework for robotics see (Pagallo 2013). In terms of philosophy of information see (Floridi 2008, 2012, 2013).