

Artificial intelligence and the future of work: Will our jobs be taken by machines?

Sofia Olhede and Patrick Wolfe discuss the current state of data-driven automation and its implications for jobs

Browse any publication these days—from major daily newspapers to government white papers to reports from learned societies—and you will be told that artificial intelligence (AI) and, more generally, automation are putting a significant proportion of jobs at risk, as much as 47%, according to one estimate of the US workforce.

For some jobs, the gradual AI takeover feels almost evolutionary. Over past decades, machines have assumed responsibility for more and more of the repetitive, mundane tasks traditionally carried out by manual workers. Think production lines, and sorting and packing. Here, machines clearly hold an unparalleled advantage over humans: they do not so easily wear out, nor do they suffer from boredom.

But advances in data science and AI are now putting other jobs in reach of machines. Work that has relied primarily on human intelligence in the past is now at risk of being automated in the future, and this feels more like revolution than evolution.

AI support, or AI takeover?

Broadly speaking, there are two types of AI. “Narrow AI” aims to excel at a particular set of tasks. Playing the game of Go, for example, is a demonstration of narrow AI: rules and sets of possibilities are unchanging, and current AI has proved adept at learning strategies at a rate sufficient to overcome millennia of collected human wisdom about the game.

“General AI”, by contrast, should be able to interact with and respond to multiple environments, and hence to learn and complete multiple tasks—which is what we expect of human intelligence. It is here that AI advances are likely to have most impact on the skilled knowledge worker, but general AI remains a somewhat distant prospect, with an unclear timeline.

In the meantime, narrow AI is broadening its abilities. Previously in this magazine we have discussed the growing use of automation in fields like health, finance and law. Invariably, this involves programming and training an algorithm to make decisions - or to recommend decisions - based on the analysis of data; something that statisticians are often called on to do. Indeed, attempts to automate at least parts of the statistical profession are ongoing; the “Automatic Statistician” project says “there are relatively few people trained in the statistical and machine learning methods required to test hypotheses, make predictions, and otherwise create interpretable knowledge” from data, and so hopes to create an AI to do just that.

The Automatic Statistician stands a decent chance of succeeding, at least in a fairly narrow domain. Like almost every job, the work of the statistician involves some rote tasks that can be more efficiently handled by machine. If an AI can be programmed to take on these basic, repetitive tasks, why not let it? This would, after all, free the hands of workers to concentrate on higher-value tasks. The end result should hopefully be increased productivity.

In this scenario, the statistician does not lose their job, they merely have an AI support worker to handle some tasks. This distinction, between “jobs” and “tasks”, is one made by Arntz, Gregory and Zierahn in a 2016 OECD report. Their report highlights the challenges of rapid technological change, especially on how it will impact training for those entering the job market over the next decade. But one of the key arguments they make is that threat of jobs lost to automation is reduced if *tasks*, rather than *jobs*, are the focus of study.

It was a 2013 paper by Frey and Osborne which first suggested that 47% of US workers were employed in jobs that could be performed by algorithms within the next 10 to 20 years. However, Arntz *et al.* write that: “Since occupations usually consist of performing a bundle of tasks not all of which may be easily automatable ... , the potential for automating entire occupations and workplaces may be much lower than suggested”. Their own analysis suggests that only 9% of US jobs are automatable.

Social implications

Whether 9% or 47% of jobs are lost, countries around the world will be left with millions of people who need retraining. And the labour market for future generations will clearly be very different. Taken together, these developments could have huge societal impacts.

In the short term, governments may have to spend more on social security to protect affected workers. Investment in adult education may be needed to ensure those without a job can find their way back into

employment. In the medium- to long-term, education will have to be rethought from the early years on up. For example, in a world in which AI-driven machines are our co-workers, it might make sense for children to be brought up speaking the language of code and algorithms.

We cannot be sure what the future holds, but history can help us prepare. We can look to the after-effects of previous technological advances and industrial revolutions for guidance. The details may be different, but the broad outcomes are likely to be the same. There will be winners and losers. Some will get richer, others poorer. Certain towns and cities will lose jobs, others will see major gains.

Some countries are already wrestling with these sorts of issues. In Finland, for example, there are those advocating for a universal salary to ameliorate the effects of potential future job losses and avoid the prospect of widening inequalities. But, of course, different countries are at different levels of AI development, as was made clear from recent discussions about national advancements in the AI race. This will inevitably mean that that the AI evolution - or revolution - will arrive at different times in different places.