



Imperial College  
London

# THE FORUM

Business as usual  
or the end of work?  
How AI will affect the  
jobs of the future



## Artificial intelligence is set to change the way we work dramatically, and soon. It will both destroy jobs and create them.

A 2015 [PriceWaterhouseCooper report](#) found that 6% of British people were in jobs that didn't exist in 1990 – and existing ones changed, as tasks within them became automated and AI tools to improve human capabilities became available. Similarly, some will be lost. A [recent report](#) for the House of Lords suggested that between 10% and 50% of jobs in the UK and other advanced economies will be at risk from automation in the next 10 to 20 years. A [similar report](#) by the Royal Society and The British Academy says that between 10% and 30% of British jobs are “highly automatable”.



How much this will be because of “artificial intelligence” is up for debate, not least because exactly what counts as “artificial intelligence” changes from year to year and person to person. The AI pioneer John McCarthy once said that “As soon as it works, no one calls it AI anymore.” A [1958 paper](#) argued that “if one could devise a successful chess machine, one would seem to have penetrated to the core of human intellectual endeavour”; but humans have been outclassed by chess computers for two decades, and grandmaster-level chess programs run on smartphones. Not all of us would describe the chess apps on our phone as “AI”.

But those debates are semantic: the line between old-fashioned computing and hypermodern AI is as irrelevant as it is blurry. Increasing computing power and sophistication, and the use of learning algorithms trained on ever-larger data sets, will allow machines to make ever more intelligent decisions, and to perform more tasks that previously only humans could do.

### What is AI?

The standard undergraduate textbook for AI, *Artificial Intelligence: A Modern Approach* by Russell and Norvig, says that the main goal of modern AI research is that of building a machine that is a rational agent: a machine that, given a goal, will make decisions that are most likely to bring about that goal.

One common approach for building such a machine is **machine learning** – creating agents that can improve their performance on future tasks by observing the world. For instance, say Russell and Norvig, a machine-learning agent that sees many images that it is told contain buses may be able to recognise them.

Machine learning can be based on **neural networks**: a system of interconnected nodes, loosely based on the human brain. Each “neuron” has a certain number of both outputs to and inputs from other neurons. Its outputs fire when the inputs reach a certain threshold. The “weight”, or importance, of a connection between two neurons increases the more often it fires, allowing the system to “learn”. **Deep learning** is a more complex form of neural network, with many layers of “neurons”.

Machine learning enables AI to learn how to carry out **specific functions**. These specific competencies do not amount to the broad abilities demonstrated by people. Human-level intelligence or **general AI** is still some time from being developed, and it is not clear when or whether it will be possible.



## Automation

While AI is new, the automation of work has been going on for centuries. **Mark Kennedy**, a professor of technology and business at Imperial College London, points out that in Britain and America in the 18th century, between 66% and 80% of the population worked in agriculture. “Now, it’s around 2% to 3%,” he says. The difference is that machines – at first ploughs, later tractors and combine harvesters – allow a far smaller number of people to do the same amount of work. Similar things have been happening ever since: robots have allowed fewer factory-workers to do the same jobs, ATMs and counting machines have reduced the need for bank tellers. Parts that are routine and repetitive – counting money, turning screws – have become automated, while the parts of the job that are less routine and involve creativity or human interaction have not. Similar patterns are expected in the future.

“In Britain and America in the 18th century, between 66% and 80% of the population worked in agriculture. Now, it’s around 2% to 3%.”

Mark Kennedy

“The whole thing of ‘robots for jobs’ is an old framing of a problem that’s moved on”

Mark Kennedy

Kurt Vonnegut’s 1952 novel **Player Piano** described a dystopian future where near-total automation leaves the vast majority of the population economically obsolete, purposeless and dependent on pay from government-provided busywork.

Those predictions are still made today. “It’s a bit binary,” says **Dr Sandra Wachter**, a research fellow in AI ethics and law at the Oxford Internet Institute. “You have one camp that says AI will automate all of us out of a job and no one is safe, and the other people who say this is going to be amazing, no one has to work anymore, and it’s going to be a leisurely society where everyone can focus on the things they really care about. I think both extremes are wrong, and the truth is somewhere in between.”

It’s unlikely that many jobs will be outright replaced by robots, says Kennedy. “Mostly gone are the days when jobs were designed around specific functions,” he says. “That was a model from the late industrial era of mass manufacturing, an assembly line where each person turns one screw. But early robots started to Hoover those jobs up and automate them.”

Now, most jobs – whether in manufacturing or other sectors – are “a whole bag of tasks”, he says. “The whole thing of ‘robots for jobs’ is an old framing of a problem that’s moved on,” he says. “Now we talk more about tools for tasks.” That is, augmenting the ability of the human to do their jobs, using AI tools.

## Augmentation

**Nick Jennings**, a professor of artificial intelligence at Imperial College London, says that “while there will be domains and activities which will be automated, I mainly talk about AI working in partnership with humans”. Machines, he says, “are good at particular things, and humans are good at things. When they work together properly that’s good for everyone.”

He gives the example of radiology. “It’s not a surprise that AI outperforms humans at looking at medical images,” he says – a project that his Imperial College colleague **Ben Glocker** is working on. “It’s a skilled task but very narrowly defined, and you can train a machine on many more images – an algorithm can look at more images in 24 hours than a human will see in its lifetime. Plus they don’t suffer from human frailties: they don’t get bored, they don’t get tired.”



Mark Kennedy address a policy workshop on AI and the future of work

Similarly, predictions of either a glorious work-free future or a grim world of joblessness are not new. John Maynard Keynes **thought** that by 2030 we would all work 15-hour weeks; an IBM economist **told Time magazine in 1965** that 20-hour weeks and a “mass leisure class” would be standard “within a century”.



But they still make mistakes. “And the mistakes that humans and machines make are different,” he says, “which is good.” It means that if they work through the same list, if the human misses something the AI is likely to get it, and vice versa. One [study](#) found that a medical-imaging AI was 92% accurate, and a trained human pathologist was 96% accurate, but the two in combination were 99.5% accurate.

The idea of a partnership goes deeper, though. As AI takes up more of the work in image recognition, it frees up the radiologist’s time elsewhere. “If a doctor’s spending less time on diagnosis and looking at images,” says Jennings, “then they can spend more time with the patient. It will be a shift in balance: now a machine does the dull mundane part of my job where I look at hundreds of images every day, while I look at a few and spend the rest of my time on things I’m better suited for.”

Kennedy talks in terms of “augmentation” rather than “automation” of jobs. But, of course, this still means that fewer humans are needed to perform a set amount of work. It may not be that radiologists are no longer necessary, but it may mean that a hospital could employ three instead of 10 and achieve the same outcome.

---

**“Machines are good at particular things, and humans are good at things. When they work together properly that’s good for everyone.”**

Nick Jennings

---

**“At the beginning of the 1900s, 80% of Americans were employed in agriculture,”**

**“Now it’s 2%, but that doesn’t mean that 78% of people are unemployed.”**

Nick Jennings

That’s where job losses are possible. Jennings, though, says that history suggests that it won’t work like that. He points, as Kennedy did, to the drop in agricultural workers. “At the beginning of the 1900s, 80% of Americans were employed in agriculture,” he says. “Now it’s 2%, but that doesn’t mean that 78% of people are unemployed.” Instead, new jobs will arise – there were no software engineers in 1900 – and the skill mix of old jobs will change. Kennedy raises the example of nursing. “The part of the job that’s turning patients over in bed,” he says, “we could get a robot to do that. It wouldn’t be cheap, and it would need to be portable and not hurt people, but we could do it.” Researchers are working on projects much like this, including [Petar Kormushev](#), one of Jennings’s colleagues at Imperial.

“But that’s only one part of the job,” says Jennings. “The other parts are a smile and conversation, helping people with their inability.” These sorts of people skills, which are rarely taught in education and do not attract a premium in the job market, could become more valuable and highly remunerated, he thinks, while more traditional skills could face downward pressure on wages.

We can’t know what the new jobs will be. “I don’t think we can predict with any certainty the new jobs that the era of AI will bring about,” says Maja Pantic, professor of affective and behavioural computing at Imperial. “The last revolution, the rise of computing and digitalisation, gave rise to jobs that no one previously could have predicted.” The role of “social media manager” would have been impossible to predict even 20 years ago. But there are obvious possible areas where new jobs will arise.

One such area will be helping to train and improve AI systems. “Last year I paid £100,000 to ‘data annotators’,” says [Stephanos Zafeiriou](#), a reader in machine learning and computer vision at Imperial. A data annotator is someone who looks at the data being fed to some machine-learning system and makes it comprehensible for that system, so it can learn from it. In the case of self-driving cars, for instance, the annotators would look at images of roads, “and annotate where the pavements are, where the traffic lights are – draw lines around the roads and the trees and the pedestrians”. Then the system will become more able to recognise them.

Similar work could be done on medical imaging – “they would say this is a lesion, this is non-cancerous” – or face recognition – “annotators drawing around the nose and the lips and giving the age of the person”. Some of these jobs could be done by people without degrees; others, such as the medical imaging example, would require expertise. Similar jobs revolving around helping to train AIs will likely be an important part of the future.



Others might involve building and maintaining the infrastructure for new AI technologies. Pantic suggests that if self-driving cars become standard – which she thinks unlikely, for regulatory reasons more than technical ones – then there will be many jobs in building the charging stations and maintaining the cars themselves. Any other major robotics industries will similarly need the robots and the infrastructure around them building and maintaining.



**Antoine Cully**, a lecturer in robotics and AI at Imperial, said AI has so far created “way more jobs” than it has destroyed, and that in the near future at least that will continue to be the case. “You still need humans in the loop to design and develop these new AI technologies,” he says. He suggests that another big change will be how AI augments jobs, as discussed above, as well as creating new ones. “For instance,” he says, “in the design of video games, it will become easier to generate automatically huge landscapes, or unique characters. Instead of cutting the jobs of graphical designers, it will become their paint brush, enabling them to create deeper and more detailed worlds.

“And medical robotics will not replace surgeons; it will assist them, to minimise the risk of operations. Robots on search and rescue missions after nuclear incidents or earthquakes can enable rescue teams to operate faster and save more people without putting human lives at risk.

“AI and robots will take the most dangerous parts, but humans will still play a central role, and new jobs will be created to ensure the maintenance of these robots and the coordination of human-robot teams.”



*Maja Pantic, Professor of Affective and Behavioural Computing at Imperial*



## Uncertain impact

Just as it is impossible to know exactly what jobs will be created, the overall impact of AI on the number of jobs on the economy is highly uncertain. Some worry that there will be a major net loss. Andrew Yang, an American entrepreneur who is running for the Democratic nomination for president in 2020, **warns** that automation of retail, call centre, fast food and, most importantly in the US, trucking jobs will lead to millions of job losses in the next few decades. Erik Brynjolfsson and Andrew McAfee of MIT **suggest** that the correct analogy is not how previous automation affected human jobs, but how the invention of the internal combustion engine affected the economic use of horses. In 1900, there were more horses in the USA than ever before – around 21 million. By 1960, that figure had dropped by 88% to 3 million. As machines became able to do horses' jobs better, more cheaply, and more reliably than horses themselves, they became economically obsolete.

The argument Brynjolfsson and McAfee – and Yang – make is that AI will, eventually, be able to do a large proportion of jobs better than humans, and so will largely replace humans in the economy. Kennedy and Jennings are doubtful that this will be the case in the long term, Jennings calling it “scaremongering” and Kennedy saying “bollocks”. But, they say, it doesn't mean that there won't be huge effects. This appears to be the consensus view, that we can expect neither disaster nor business as usual: The British Academy and Royal Society report agrees that discussion of AI and the future of work tends to “oscillate between fears of the ‘end of work’ and reassurances that little will change in terms of overall employment”, but that “evidence suggests neither of these extremes is likely”.

Kennedy draws a comparison to the industrial revolution. It's true, he says, that there was no long-term decline in employment. “Over 250 years you can tell a very positive story,” he says. “But in the mid-19th century, when there were industrial areas growing and rural areas in decline, it was rough. The French had a bloody revolution over it, and the UK could have done if it hadn't happened in France first.”

And the impacts the AI revolution will be, in some ways, more profound. “The industrial revolution is a good comparison,” says **Marcus du Sautoy**, a professor of mathematics and author of **The Creativity Code**, a book about the impact of AI. “But the difference is the speed. The industrial revolution affected one generation to the next: your sons and daughters once would

—  
“But the difference is the speed. The industrial revolution affected one generation to the next: your sons and daughters once would have followed you into your profession, but now they can't.”

Marcus du Sautoy



have followed you into your profession, but now they can't. This time, things are happening so fast that it's not your children, it's you.”

This will have profound implications for public policy. Automation is expected to have the biggest impact on lower-skilled work. Kennedy points out that wage stagnation in the US and Britain has now gone on for a decade, longer than can be explained by the financial crisis, and that many people suspect automation is behind it. Yang, the US presidential candidate, **argues** that “the reason Donald Trump is our president today is we automated away 4 million manufacturing jobs”.

There is also a risk that AI will lead to greater inequality, and that it could further strip away the workers' protections that have grown over the last century or two and that have already been damaged by the gig economy, according to **Christina Hritova**, a digital ethics researcher at the **Alan Turing Institute** for artificial intelligence. Pantic adds that security in jobs is likely a thing of the past – that “more and more we will have multiple jobs. We will develop an expertise, and a company will need this expertise for a certain period, and then the job will be done and another company will need it.” This feels insecure to us, she says, “but that is also because of perception, because our parents had stable jobs for 30 years”. People who grow up with that model of work may not be so concerned about it.





Further, AI has the potential to concentrate expertise more than ever into the hands of a few giant US companies. The House of Lords report **points out** that smaller firms are already lagging behind larger ones technologically. Many small and medium-sized enterprises (SMEs) are “struggling to gain access to large, high-quality datasets, making it extremely difficult for them to compete with the large, mostly US-owned technology companies”, says the report. AI software itself may be prohibitively expensive for SMEs, and larger companies can afford to pay more money for talented staff.

There is also a problem, says Pantic, that AI expertise will tend to concentrate in those same few large US technology companies, just when it is needed by society at large. “AI is coming into all parts of our life,” she says. “But you have very few experts in AI, most of whom are employed by five or six of the usual suspects, so you have a concentration of intellectual capital.

“All the innovation will come from those companies, and other companies will find it harder to rise up because they don’t have the expertise. And universities cannot compete, because these firms will offer researchers twice or five times the salary and not require them to fill out tedious forms or raise funding.” Keeping AI research publicly available, via academia, will be important, she says, so universities need to work hard to make themselves attractive to talented staff.

**The economic issues are more complicated. If people are going to go out of work, whether long-term or short-term, there will be a need to provide social safety nets, all the experts agreed.**

## Ethics

There are other concerns. Hritova raises the problem that the AI algorithms will be trained on existing datasets, and will therefore encode existing biases – such as racism or sexism – and give them a gloss of objectivity. Amazon created an algorithm to help with hiring decisions, and trained it on all the CVs that had been submitted to the company over 10 years. But they **had to scrap it** after realising that it was systematically underrating women for technical jobs – because so few female engineers had been hired by the humans who made the decisions in that time. An algorithm designed to assess the risk that a prisoner will reoffend if they are released on parole **turned out** to be much more likely to release white prisoners than black ones.

Hritova also points out that there is a potential problem when algorithms are used as tools to suggest courses of action – such as deciding which prisoners to parole. “There’s an accountability problem,” she says. “If the AI makes a suggestion, and you follow it, and it was wrong – or you don’t, and it was right – you need to be very clear about who is responsible.” If a worker feels they will be blamed for overriding an AI, they won’t stop it making what they think are bad decisions, and the “partnership” effect will be reduced.

Society, governments and employers will have to take steps to reduce the impact of these problems. The issues of bias and accountability, says Hritova, will be best served by making AI “explainable”: ensuring that it shows its working, that the reasoning behind its decisions is visible and understandable. That could have knock-on effects: for instance, the Amazon AI may have been biased, but it also revealed bias in their existing hiring practices, which could allow Amazon to correct for it.

The economic issues are more complicated. If people are going to go out of work, whether long-term or short-term, there will be a need to provide social safety nets, all the experts agreed. Kennedy points to the introduction of the Speenhamland system in 1792, as an early example; it provided free food for impoverished workers outside their own parishes. “It jumpstarted the industrial revolution in Britain,” he said, by making it possible for workers to move around for work, although it also trapped people in dependency and was replaced by a more comprehensive system in 1834.

## Policy responses

The equivalent modern response to automation may be a universal basic income, or UBI. Yang's presidential bid is centred around this idea: he proposes giving \$1,000 a month to every adult US citizen, no questions asked, funded by a value-added tax on major tech firms. Wachter says that there have been positive experiences in small trials of UBI in Finland and the Netherlands, but that more trials are important. The Trade Union Congress feels the **idea has merit** but says that there are more practical solutions that more directly address the problems, such as wage stagnation and inequality, associated with automation.

Beyond social security, others point to the need for a huge amount of training and retraining, likely via a large government-funded program. "We have to recognise that people are going to have to constantly retrain in the light of a very fast changing workplace," says Du Sautoy. "Every 10 years you're going to have to retrain. The things you're doing now are going to change." Zafeiriou agrees. "There are a lot of professions and jobs where people don't have advanced degrees," he says. "All these people are going to have to be retrained for new jobs and new opportunities."

There are already moves to support this in the public policy sphere. In 2018 the Department for Education **announced** a National Retraining Scheme, and the Trade Union Congress has **called for** extensive government-funded retraining for mid-career workers whose jobs are at risk.

Whether large-scale retraining, social security measures, or some alternative, large government programmes will have to be funded somehow. Zafeiriou says that because so much of the main tech companies' wealth has been built on user data and taxpayer-funded infrastructure, "these companies have to give something back". He proposes major retraining schemes funded by taxes on the main firms. Wachter highlights the possibility of an automation tax: "If you have fewer people employed then they pay less tax, so that cuts into the government budget, so the idea is if you automate a lot of your workforce, you should pay a tax." Hritova points out that the French finance minister **recently suggested** that a tax on the digital advertising revenue of major international firms could bring in £500 million a year. Both of them, however, warn that over-taxing firms could prevent them from innovating and expanding, and thus cause more harm than good.

## Conclusion

One way or another, the AI jobs revolution is coming. And despite the real concerns, it is, says Wachter, also something to celebrate. "If we look back in 10 years' time, we'll probably think that it was irresponsible of us to let humans do some of the things we do," she says.

Most of all, though, there will be a choice. Automation will continue to increase the amount of work that can be done by a single human. Firms can go one of two ways with that, says Jennings. "Some will use it to reduce their workforce," he says, to do the same work for less money. That will be good for consumers: those services will become cheaper. But others will go the other way. "They'll reinvest. They'll keep their loyal, trained workers, and have them doing more things. They'll make more money, and offer a better service at the same cost." Both have advantages: neither is the sole "correct" path. But if enough firms do the latter, humans can thrive in the jobs of an AI-led future.

### Written for Imperial by Tom Chivers

Tom Chivers is freelance science writer. He was formerly a science writer for BuzzFeed UK and the Telegraph. He is the author of *The AI Does Not Hate You: The Rationalists and Their Quest to Save the World*.



## **The Forum features**

The Forum is Imperial's policy engagement programme. It connects Imperial researchers with policy-makers to discover new thinking on global challenges. Our features provide a shop window into the world leading research taking place at Imperial and provide insight into how it can inform and contribute to public policy debates.

**THE  
FORUM**

### **Connect with The Forum**

[the.forum@imperial.ac.uk](mailto:the.forum@imperial.ac.uk)

[imperial.ac.uk/the-forum](https://imperial.ac.uk/the-forum)

[@imperial\\_forum](https://twitter.com/imperial_forum)