
The 4th Industrial Revolution Series: The 4th Industrial Revolution Series

Author: Barry Dwolatzky

When the short stocky man in the front row started to shout at me, I was horrified. He waved his hands as he growled in a deep voice. He blamed me for “this Fourth Industrial Revolution”. He glared angrily and said, “you and your robots are stealing my job!”. I thought this was all a bit unfair, since I’m certainly not responsible for the Fourth Industrial Revolution - although I might have made some minor contributions here and there. Nor do I own a single robot!

My talk was on “The Fourth Industrial Revolution”, or 4IR, and I’m the kind of person who loves interaction when I talk to an audience. When the man in the front row raised his hand, I paused and smiled, giving him space to interject, which he did, loudly and clearly!

From then on, my presentation moved off in an unplanned direction. Several other people spoke. While they were less confrontational and far more polite, they certainly shared his view. It reminded me of a popular movie that came out in the 1960’s, in the midst of the Cold War, with the title “The Russians are Coming the Russians are Coming”. There was a sense of irrational panic, as the consensus in my audience seemed to be that “The Robots are Coming the Robots are coming”. My talk became a lively debate on the theme of robots and jobs.

An “Industrial Revolution” implies a disruption in the way things are done. In the 18th Century steam power replaced animal and human power. Hand loom operators were replaced by machines. In the early 1900s electricity and mass production led to craft workers becoming all but redundant. Digital computers in the late 20th Century automated repetitive and tedious office tasks, leading to a loss of jobs. Each of these previous Industrial Revolutions resulted, however, in a net gain of jobs, rather than a loss.

So why is it that everyone seems to think that 4IR will lead to huge job losses?

The main reason is that 4IR is partly about giving machines human, or super-human, levels of intelligence. There is the fear that robots with high levels of artificial intelligence (AI) will soon be capable of doing anything that humans can do. They will be faster, cheaper and better than us humans.

How valid is this fear? To answer this question, we need to understand what the limitations of AI are, rather than its achievements.

In the 1980’s I worked at a large corporate research laboratory in Britain. I was part of a team developing AI for a flexible assembly robot. The task we were trying to automate was this: you are given a box of parts and some written instructions. It could be any box of parts with corresponding instructions. Using standard tools, you follow the instructions and assemble a product from the parts in the box. This is a job that a competent human worker, or even a smart teenager, can do quite easily. However, getting an AI-empowered robot to successfully do this in the 1980s proved to be impossible. Nearly 40 years later this task is still too complex for a robot.

The AI we used in the 1980s was called “symbolic AI”, and required us to write hundreds of rules that would allow the robot to infer other rules in finding a solution. The problem we couldn’t solve then was how to give the AI “common sense”. In AI, “common sense” is all of those things that humans “just know”. Even a small child knows that objects fall downwards, that water flows, that if something is hidden behind something else it’s still there. In 1984 Douglas Lenat launched the “CYC Project” with the intention to codify, in symbolic AI rules, EVERYTHING that humans take for granted. For example: every human has a father; every human has a mother; a cup of coffee gets cooler if you leave it standing on a table; etc. Thirty-six years later Lenat and his colleagues are still hard at work. They estimate that they are only 5% finished. It seems like an impossible task!

Today's AI is based on Machine Learning (ML). Instead of writing rules we now create large systems that aim to mimic the neural networks in a brain. We "teach" these networks using huge amounts of training data. These types of systems have proven to be extremely successful in areas like speech recognition and computer vision. Neural networks, however, also fall short when it comes to "common sense". This short-coming leads to otherwise very smart AI systems coming to some unbelievably silly conclusions. Not only that, but it's impossible to explain why it reached such conclusions.

In its famous contest against human competitors in the TV quiz show Jeopardy, IBM's Watson AI system was asked to name a specific city in the USA. It gave the answer "Toronto". Although Watson went on to beat the best human Jeopardy players, its "trainers" were unable to explain how it could make such a silly mistake.

The good news for humans, like the man who confronted me during my talk on 4IR, is that until we can equip our robots with as much common sense as a human, robots won't be taking our jobs in many categories. I do concede, however, that in certain areas, where knowledge is very narrow and specific, humans will be replaced by robots. At the same time other jobs will open up. Maybe Douglas Lenat's team could use a few thousand more humans to codify 'common sense' rules for the CYC Project?

I'm convinced that the Robots are NOT Coming! There are many things they won't be able to do for a long time yet. Finally, If you're interested in learning more about "common sense" in AI, read Melanie Mitchell's recent book "Artificial Intelligence: A Guide for Thinking Humans" (2019).