


The future of artificial intelligence—two experts disagree

July 17 2017, by Peter Stratton And Michael Milford



Credit: AI-generated image ([disclaimer](#))

Artificial intelligence (AI) promises to revolutionise our lives, drive our cars, diagnose our health problems, and lead us into a new future where thinking machines do things that we're yet to imagine.

Or does it? Not everyone agrees.

Even billionaire entrepreneur Elon Musk, who admits he has access to some of the most cutting-edge AI, [said recently](#) that without some regulation "AI is a fundamental risk to the existence of human civilization".

So what is the future of AI? Michael Milford and Peter Stratton are both heavily involved in AI research and they have different views on how it will impact on our lives in the future.

How widespread is artificial intelligence today?

Michael:

Answering this question depends on what you consider to be "artificial intelligence".

Basic machine learning algorithms underpin many technologies that we interact with in our everyday lives - voice recognition, face recognition - but are application-specific and can only do one very specific defined task (and not always well).

More capable AI - what we might consider as being somewhat smart - is only now becoming widespread in areas such as [online retail and marketing](#), [smartphones](#), [assistive car systems](#) and service robots such as [robotic vacuum cleaners](#).

Peter:

The most obvious and useful examples of current AI are the [speech recognition on your phone](#), and search engines such as [Google](#). There is also IBM's Watson, which in 2011 [beat human champion players](#) at the US TV game show Jeopardy, and is now being trialled in business and healthcare.

Most recently, Google's DeepMind AI called AlphaGo [beat the world champion Go player](#), surprising a lot of people – especially since Go is an extremely complex game, way surpassing chess.

What major advances in AI will we see over the next 10 years?

Peter:

Many auto manufacturers and research institutions are competing to create practical driverless cars for general road use. While currently these cars can drive themselves for much of the time, many challenges remain in dealing with bad weather (heavy rain, fog and snow) and random real-world events such as roadworks, accidents and other blockages.

These incidents often require some degree of human judgement, common sense and even calculated risk to successfully navigate through. We are still a long way from fully autonomous vehicles that don't need a licensed driver ready to take control in an instant.

The same can be said for all the AI that we will see over the coming 10-20 years, such as online [virtual personal assistants](#), [accountants](#), [legal and financial advisers](#), [doctors](#) and even physical [shop-bots](#), [museum guides](#), [cleaners](#) and [security guards](#).

They will be advanced tools that are very useful in specific situations, but they will never fully replace people because they will have little common sense (probably none, in fact).

Michael:

We will definitely see a range of steady, incremental improvements in everyday AI. [Online product recommendations will get better](#), your phone or car will [understand your voice increasingly well](#) and your [vacuum cleaner robot won't get stuck as often](#).

It's likely that we'll see some major advances beyond today's technology in some but not all of the following areas: self-driving cars, [healthcare](#), utilities ([electricity](#), water, and so on) management, legal, and service areas such as cleaning robots.

I disagree on self-driving cars - there's no real reason why there won't be fully autonomous controlled ride-sharing fleets in the affluent centres of cities, and this is indeed the strategy of companies such as [NuTonomy](#), working in Singapore and Boston.

What approaches will lead to the biggest improvements in AI?

Michael:

Major advances will come from two sources.

First, there is a long runway of steady incremental improvements left in many areas of conventional AI - large, complex neural networks and algorithms. These systems will continue to improve steadily as more training data becomes available and as scientists perfect them.

The second area will likely be biological inspiration. Scientists are only [just starting to tap into the knowledge about how brain networks work](#), and it's likely they will copy or adapt what we know about animal and human brains to make current deep learning networks far more capable.

Peter:

Old-fashioned AI, which was based on pure logic and computer programs that tried to get machines to behave intelligently, basically failed to do anything that humans are good at and computers are not (speech and image recognition, playing complex strategic games, for example).

What's quite clear now is that our best-performing AI is based on how we think the brain works.

But our current brain-based AI (called Deep Artificial Neural Networks) is still light years away from emulating an actual brain. Enhanced AI capabilities in the future will come from developing better theories of how the brain works.

The fundamental science needed to cultivate these theories will probably come from publicly funded research institutions, which will then be spun off into commercial start-up companies, and then quickly acquired by interested large corporations if they look like they might be successful.

How will artificial intelligence affect society and jobs?

Peter:

Most jobs won't be under threat for a long time, probably several generations. Real people are needed to actually make any significant decisions because AI currently has no common sense.

Instead of replacing jobs, our overall quality of life will go up. For example, right now few people can afford a personal assistant, or a full-

time life coach. In the near future, we'll [all have \(a virtual\) one!](#)

Our [virtual doctor](#) will be working for us daily, monitoring our health and making exercise and lifestyle suggestions.

Our houses and workplaces might be cleaner, but we will still need people to clean the spots the robots miss. We'll also need people to deploy, retrieve and maintain all the robots.

Our goods will be cheaper due to [reduced transport costs](#), but we'll still need human drivers to cover all the situations the self-drivers can't.

All this doesn't even mention the [whole new entertainment technologies](#) and industries that will spring up to capture our increased disposable income and to cash-in on our improved quality of life.

So yes, jobs will change, but there will still be plenty of them.

Michael:

It's likely that a significant fraction of jobs will be under threat over the coming decade. It's important to note that this won't necessarily be divided by blue-collar versus white-collar, but rather by [which occupations are easily automatable](#).

It's [unlikely that an effective plumber robot will be built](#) in the near future, but [aspects of the so far undisrupted construction industry may change radically](#).

Some people say machines will never have the emotional capabilities of humans. Whether that is true or not, many jobs will be under threat with even the most rudimentary levels of emotional understanding and interaction.

Don't think about the complex, nuanced interaction you had with your psychologist; instead think about the one with that disinterested, uncaring part-time hospitality worker. The bar for disruption is not as high as many think.

That leaves the question of what happens then. There are two scenarios - the first being that, like in the past, new types of jobs are generated by the technological revolution.

The other is that humanity gradually transitions into a Utopian society where scientific, artistic and sporting pursuits are pursued at leisure. The short to medium-term reality is probably somewhere in between.

Will Skynet/the machines take over and enslave humanity?

Michael:

It's unlikely in the near future but possible. The real danger is the unpredictability. Skynet-like killer cyborgs as featured in the [Terminator film series](#) are unlikely because that development cycle takes a while, and we have multiple opportunities to stop development.

But AI could destroy or damage humanity in other unpredictable ways. For example, when big companies like Google Deepmind start [entering into healthcare](#), it's likely that they will improve patient outcomes through a combination of big data and intelligent systems.

One of the temptations or pressures will be to deploy these extremely complex systems before we completely understand every possible ramification. Imagine the pressure if there is good evidence it will save thousands of lives per year.

As we well know, we have a [long history of negative unintended consequences](#) with new technology that we didn't fully understand.

In a far-fetched but not impossible healthcare scenario, deploying AI may lead to catastrophic outcomes - a world-wide AI network deciding in ways invisible to us human observers to kill us all off to optimise some misguided performance goal.

The challenge is that with newly developing technologies, there is an illusion of 100% control, which doesn't really exist.

Peter:

All our current AI, and any that we can possibly create in the foreseeable future, are just tools – developed for specific jobs and totally useless outside of the exact duties they were designed for. They don't have thoughts or feelings. These AIs are just as likely to try to take over the world as your Xbox or your toaster.

One day, I believe, we will build machines that [rival us in intelligence](#), and these machines will have their own thoughts and possibly learn in an unconstrained way. This sounds scary. But humans are dangerous for exactly the reasons that the machines won't be.

Humans evolved in a constant struggle for life and death, which made us innately competitive and potentially treacherous. When we build the machines, we can instead build them with any underlying motivation that we would like.

For example, we could build an intelligent machine whose only desire is to dismantle itself. Or, we could build in a hidden remote-controlled off switch that is completely separate from any of the machine's own circuits, and an auto-shutdown reflex if the machine somehow ever

notices it.

All these safeguards will be trivial to implement. So there is simply no way that we could accidentally build a machine that then tries to wipe out the human race.

Of course, because humans themselves are dangerous, someone could build a machine that doesn't have these safeguards and use it for nefarious purposes. But we have that same problem now with nuclear weapons.

In the future, just as now, we have to hope that we are simply smart enough to use our technology wisely.

This article was originally published on [The Conversation](#). Read the [original article](#).

Provided by The Conversation

Citation: The future of artificial intelligence—two experts disagree (2017, July 17) retrieved 2 May 2024 from <https://techxplore.com/news/2017-07-future-artificial-intelligencetwo-experts.html>

<p>This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.</p>
--