What should governments be doing about the rise of artificial intelligence?

October 31 2017, by David Glance

Credit: Sora Shimazaki from Pexels

There is little doubt that artificial intelligence (AI) is transforming almost every facet of human life. How far this transformation will go and what the full ramifications for society will be are still unknown but

this hasn't prevented people from making both [optimistic](#) and [dire](#) predictions.

Elon Musk's call for AI [regulation](#) has been matched by equal calls for governments [not to](#).

AI's problem with definition

One of the principle problems with AI has been the [confusion](#) that surrounds what it is exactly, and what it can and can't actually do. The single biggest problem in understanding AI however has been making it clear how current AI techniques (like [deep learning](#)) differ from human intelligence.

Getting to the facts

In order to answer some of these questions, the OECD held a [conference](#) last week on AI. Government and industry representatives, AI academics and others met to review the state of AI and pose the question of what governments could, and should do, in creating policy to take advantage of the benefits of AI whilst minimising the risks.

The first thing that became clear is that the focus of discussion was mainly on machine learning and in particular, [deep learning](#). Deep learning software learns to be able to recognise patterns from data. Google, for example, is using it to recognise [pets](#) by their faces. Another company, [DeepL](#), uses deep learning to do high quality language translation.

Speakers emphasised that deep learning works only because it uses a large amount of data that is processed on powerful computers. It has become successful as a technique because companies have access to

large amounts of data and at the same time, to large amounts of cheap processing power.

The concerns about data used for AI applications

With the use of large amounts of data, questions immediately arise as to from where the data is collected, and what exactly it is being used for.

The use of large amounts of potentially personal data raises privacy concerns, and also concerns about how exactly this data is used in determining outcomes with real world consequences.

In the US for example, deep learning is already being used to calculate the terms of sentencing in court cases. There is no way for anyone to know how the software arrived at a particular decision, especially what factors in the data were the most important in making that determination. In one particular case in the US, machine learning assisted sentencing was subsequently challenged. The challenge failed however because the courts felt that the outputs of the machine learning sentencing system were sufficiently transparent, and further details of how the system worked shouldn't be revealed.

The dangers of biased data

AI researcher Joanna Bryson has previously shown that data used to train machine learning contains a range of biases including those around race and gender. This has serious consequences with the decisions that are made when AI systems are trained with this type of data. Biased data will reinforce bias in the decisions of these systems.

Other researchers have shown that it is possible to corrupt data used to train machine learning making it possible to for example, fool cars by

adding silver tape to a road sign, possibly triggering the car to act inappropriately.

Whilst the benefits of machine learning may not be fully realised without access to a great deal of data, there has to be a balance of the risks of concentrating the collection of ever more personal data that is held in the hands of a few companies or governments.

Many participants of the conference viewed data as being the central driver of AI and the area most in most need of government regulation.

Responsibility and liability in AI applications

Another important set of questions arise out of the question of product liability and corporate responsibility. If a self-driving car causes an accident, who should be held to blame? The manufacturer of the car, the software developer of the component AI that failed, or the owner of the vehicle? Again, there has been much discussion on the subject but no real conclusions, although there is an expectation that the liability will fall on the car makers.

It isn't just hype

One concern that has been expressed in the discussion on AI is the possibility that its impact is being over-hyped. The applications that are making the most impact today are examples of pattern recognition and not general intelligence. This ability is still very useful in a range of areas like science, medicine, cybersecurity as well as a wide range of other areas.

So, what should government being doing about AI?

When it comes to what governments should be doing, there was an implied agreement at the conference that they should be enabling AI to be used for their obvious benefits to society. This has to be balanced by minimising the risks of the increased collection of personal data and also the risks of how the AI is actually using that data.

There are many more areas of discussion that become important for governments and the public in considering the role of AI in their societies. What makes this a challenge is that AI touches every aspect of life to a greater or lesser extent. What we still don't know yet is how far the development of AI will go, and ultimately how successful it will be in becoming a generalised, human-like intelligence.

This article was originally published on The Conversation. Read the original article.

Provided by The Conversation