Horizon Scanning Series

The Effective and Ethical Development of Artificial Intelligence: An Opportunity to Improve Our Wellbeing

Environmental Costs (and Benefits) of AI

This input paper was prepared by Professor John Quiggin (University of Queensland)

Suggested Citation

Quiggin, J (2018). Environmental Costs (and Benefits) of AI. Input paper for the Horizon Scanning Project "The Effective and Ethical Development of Artificial Intelligence: An Opportunity to Improve Our Wellbeing" on behalf of the Australian Council of Learned Academies, www.acola.org.

Environmental costs (and benefits) of AI

John Quiggin, University of Queensland

Recent developments in artificial intelligence (AI) have raised the prospect of a fundamental transformation of economic and social activity. Some analysts have claimed that jobs of all kinds will be replaced by AI (Brynjolfsson and McAfee 2012). Others have suggested that, on the contrary, the pace of technological progress is slowing (Gordon 2016).

Given the possibility of a radical transformation, it is important to consider the possible environmental consequences, beneficial and otherwise. The purpose of these notes is to suggest directions for relevant research. The starting point is that the greatest potential for transformation arises from distributed intelligence rather than from traditional versions of AI.

The paper is organized as follows. Section 1 argues that AI, as traditionally conceived, has consistently fallen short of the expectations it has raised. Section 2 describes the rise of distributed intelligence. Section 3 examines blockchain, an important instance of how distributed intelligence can be environmentally damaging. Section 4 examines benefits and costs the potential of autonomous vehicles. Section 5 examines the Internet of Things. Section 6 argues that the potential benefits of distributed intelligence can only be realised through comprehensive carbon pricing. Section 7 offers concluding comments.

1. The limited success of traditional AI

The traditional approach to artificial intelligence, exemplified by the work of Marvin Minsky at MIT (Minsky 1961), involves the use of computer programs to replicate or improve on various forms of human intelligence, such as game playing, problem solving, language translation and pattern recognition. The hope was that the principles used in specific domains could be extended to produce a general purpose computer with intelligence comparable to that of human beings.

This program has produced notable success in specific domains, most notably in games of skill such as Chess and Go, where computers are now able to defeat the best human players. More recent, and socially significant, examples include the development of self-driving cars and improvements in voice and face recognition.

Despite these and other specific achievements, the broader goals of traditional AI have not been realised, and substantial progress remains elusive. The techniques used in the cases mentioned above rely on the use of massive computing power to conceptually straightforward pieces of

reasoning. In this they resemble earlier applications of computing power to such problems as the calculation of rocket trajectories.

A notable area of disappointment for the traditional AI program was that of 'expert systems', developed in the 1980s and 1990s with the promise of systematising the judgement of experts in specific domains such as medicine and engineering. The original idea of expert systems was effectively abandoned during the "AI Winter" of the 1990s and has never been revived.

The 'expert system' research program contributed to some noteworthy advances in statistical methods, such as the development of Bayesian network theory (Pearl 1985). The most positive outcome has been that, in a number of fields, such as the prediction of election outcomes, evidence-based research methods and statistical analysis have been shown to outperform 'expert' judgements based on 'seat-of-the-pants' punditry.

2. Distributed intelligence

The traditional AI approach was based on the idea of embodying intelligence in a standalone device, such as a mainframe computer. The rise of the Internet has transformed the way in which we use and interact with information technology, effectively breaking down the distinction between computing and communications.

By contrast with earlier large networks, which relied on powerful central computers and switchboards to connect and manage telephone equipment or 'dumb' terminals, the Internet was the first major instance of *distributed (artificial) intelligence*. The connections that make up the Internet depend on a simple suite protocols, the most important of which are the Transmission Control Protocol (TCP) and the Internet Protocol (IP). The Internet protocol suite provides end-to-end data communication specifying how data should be packetized, addressed, transmitted, routed, and received.

The power of the Internet relied (and still relies) on the capacities that can be implemented on the computers and other devices that can be connected to it. The Internet serves as both the archetype and the enabling technology for a large class of applications of distributed intelligence.

The distribution of intelligence over large numbers of computers (millions of potentially even billions) has the potential to transform resource use, for good or ill.

3. Bitcoin and blockchain

The blockchain technology underlying cryptocurrencies such as Bitcoin is an instance of distributed intelligence leading to severely adverse environmental consequences. Bitcoin

maintains a distributed ledger of all transactions. Updating the ledger requires the provision of 'proof of work', that is, the solution of a complex but uninteresting computational problem.

As Bitcoin has become more valuable, the complexity of the problems required to 'mine' new coins has increased and so has the electricity required to drive the computers needed to solve them. On recent estimates, Bitcoin mining now accounts for 0.5 per cent of total world consumption of electricity.

Despite this massive consumption of scarce resources and the associated emissions of carbon dioxide, there is no evidence that Bitcoin serves any useful social purpose or has any underlying value. The calculations used to generate Bitcoin serve no purpose.

Moreover, the idea that Bitcoin and other cryptocurrencies will serve as a replacement for conventional currencies has been thoroughly refuted by experience. The number of merchants who accept Bitcoin is minimal, and is declining over time. It is now suggested that rather than being a currency, Bitcoin is a 'store of value', a striking claim given that Bitcoin has no actual value, while causing immense environmental damage.

Bitcoin and other cryptocurrencies do not fit comfortably into the traditional conception of AI. However, they have enough in common with AI to be of interest. The catastrophic environmental effects of Bitcoin could feasibly arise with general AI technologies. Bitcoin thus represents an illustration of the potential environmental risks of distributed artificial intelligence.

4. Autonomous vehicles

The recent development of autonomous vehicles provides a point of intersection between traditional and distributed AI. Autonomous vehicles have been developed to the point where they can manage complex navigational exercises without external support. On the other hand, it seems unlikely they will ever manage urban environments safely enough to unless they communicate with other vehicles and draw on information from sensors in the traffic network.

Autonomous vehicles have the potential for environmental benefits, but also for great environmental harm. Among the environmental benefits are more efficient organisation of traffic and reductions in the stock of vehicles required to provide transportation services.

On the other hand, less favourable scenarios can also be imagined. For example, autonomous vehicles might be instructed to drive their owners into the city, then circle around continuously to avoid parking charges.

5. Internet of Things

Future developments of the same kind may arise with the 'Internet of Things', that is, the proliferation of devices of all kinds connected to the Internet (Mattern and Floerkemeier 2010). As with autonomous vehicles, the Internet of Things has the potential to allow for more efficient use of energy and resources. On the other hand, there is also the potential to exploit opportunities for arbitrage.

As an example of potential benefits, Internet-connected devices could be set to use electricity only when there is an excess supply of renewable electricity. On the other hand, if electricity generated from coal-fired power stations late at night is cheap, such devices could exploit these sources.

6. The role of prices

The complexity of distributed intelligence creates great difficulties for solution involving regulation of energy use. The only comprehensive solution is one based on pricing of carbon emissions from all sources (Pezzey, Jotzo and Quiggin 2008).

In the presence of a comprehensive price, every autonomous vehicle and every appliance connected to the Internet of Things operates under incentives to match the benefits of emissions to the costs of carbon emissions.

7. Concluding comments

Much remains unclear about the future of AI. It remains to be seen whether the promise of a radical transformation of economic and social life will be realised. If such a transformation takes place, social institutions need to adjust to ensure that the potential benefits will be realised.

References

Brynjolfsson, E. and McAfee, A. (2012) *Race Against the Machine: How the Digital Revolution is Accelerating Innovation, Driving Productivity, and Irreversibly Transforming Employment and the Economy,* Digital Frontier Press,

Gordon, R.J. (2016) *The Rise and Fall of American Growth: The U.s. Standard of Living Since the Civil War,* Princeton University Press,

Mattern, F. and Floerkemeier, C. (2010), 'From the Internet of Computers to the Internet of Things', pp. 242-259 in *From active data management to event-based systems and more*, Springer,

Minsky, M.L. (1961), 'Steps Toward Artificial Intelligence', http://books.google.com.au/books?id=9_3xOgAACAAJ&hl=&source=gbs_api.

Pearl, J. (1985), 'Bayesian Networks: A Model of Self-Activated Memory for Evidential Reasoning', Proceedings of the 7th Conference of the Cognitive Science Society, University of California, Irvine, CA. pp. 329–334. Retrieved 2009-05-01,

Pezzey, J., Jotzo, F. and Quiggin, J. (2008), 'Fiddling while carbon burns: why climate policy needs pervasive emission pricing as well as technology promotion', *Australian Journal of Agricultural and Resource Economics*, http://anzsee.org/anzsee2007papers/Abstracts/Jotzo.Frank.pdf.