

Open Voices

AI thinks like a corporation—and that's worrying

Artificial intelligence was born of organisational decision-making and state power; it needs human ethics, says Jonnie Penn of the University of Cambridge



Open Future

Nov 26th 2018 | by BY JONNIE PENN


Artificial intelligence is everywhere but it is considered in a wholly ahistorical way. To understand the impact AI will have on our lives, it is vital to appreciate the context in which the field was established. After all, statistics and state control have evolved hand in hand for hundreds of years.

Consider computing. Its origins have been traced not only to analytic philosophy, pure mathematics and Alan Turing, but perhaps surprisingly, to the history of public administration. In “The Government Machine: A Revolutionary History of the Computer” from 2003, Jon Agar of University College London charts the development of the British civil service as it ballooned from 16,000 employees in 1797 to 460,000 by 1999. He noticed an uncanny similarity between the functionality of a human bureaucracy and that of the digital electronic computer. (He confessed that he could not tell whether this observation was trivial or profound.)

Get our daily newsletter

Upgrade your inbox and get our Daily Dispatch and Editor's Picks.





Both systems processed large quantities of information using a hierarchy of pre-set but adaptable rules. Yet one predated the other. This suggested a telling link between the organisation of human social structures and the digital tools designed to serve them. Mr Agar draws a link to the very origins of computing: Charles Babbage's Difference Engine in the 1820s in Britain. It had been subsidised by the government, on the expectation that it would serve its sponsor. Babbage's designs, Mr Agar observes, must be seen as "materialisations of state activity".

This relationship between computing systems and human organisational structures echoes through the history of AI. In the 1930s and 1940s, Herbert Simon (pictured below), a political scientist from the University of Chicago who later taught at Carnegie Mellon University, set out to develop a "scientific" account of administrative organisation. Simon had trained under Rudolf Carnap, a member of the Vienna Circle of logical positivists. This informed his belief that existing theories lacked empiricism. His doctoral dissertation in 1947 became "Administrative Behaviour", a book that provided a framework through which all activity in an organisation could be understood using a matrix of decision-making.

Simon says

He went on to make huge contributions in a host of fields—not just political science and economics, but computer science and artificial intelligence. He coined the term "satisficing" (to accept the good rather than strive for the optimal) and developed the idea of "bounded rationality" for which he won a Nobel prize in economics in 1978. But back in the 1950s, Simon was a consultant at the RAND Corporation, an influential think-tank supported by America's Air Force.

At RAND, Simon and two colleagues—Allan Newell, a young mathematician, and J. Clifford Shaw, a former insurance actuary—tried to model human problem-solving in terms that a computer could put into operation. To do so, Simon borrowed elements from the framework that he had developed in "Administrative Behaviour". To make a computer "think" as a human, Simon made it think like a corporation.


The product of the trio's labour was a virtual machine called the Logic Theorist, heralded as the first working prototype of artificial intelligence. Printouts of the Theorist in operation turned heads at the 1956 Dartmouth Summer Research Project on Artificial Intelligence, which gave the field its name and initial membership. In notes from the Dartmouth conference, one participant wrote that the Theorist helped to solve the dreaded "demo to sponsor" problem. This was essential, because the foundation funding AI was sceptical that the research area was worthwhile.

“ To make a computer ‘think’ as a human, Simon made it think like a corporation

How did Simon see his contribution? A year after the Dartmouth conference, he and Newell presented their results as “Heuristic Problem Solving: The Next Advance in Operations Research”. The clue is in the title: “operations research” emerged in Britain during the second world war to apply scientific principles and statistics to optimise military activities, and later, for corporate uses. AI meant business.

In a speech to operations-research practitioners in London in 1957, Simon identified Frederick Taylor, the father of the scientific-management movement, and Charles Babbage, as intellectual predecessors. “Physicists and electrical engineers had little to do with the invention of the digital computer,” Simon said. “The real inventor was the economist Adam Smith.” He explained the connections: Gaspard de Prony, a French civil engineer, set out to “manufacture” logarithms using techniques drawn from Smith’s “The Wealth of Nations”. Babbage, inspired by Prony, converted this insight into mechanical hardware. In the mid-1950s, Simon transmuted it into software code.

The tradition lives on. Many contemporary AI systems do not so much mimic human thinking as they do the less imaginative minds of bureaucratic institutions; our machine-learning techniques are often programmed to achieve superhuman scale, speed and accuracy at the expense of human-level originality, ambition or morals.



Capitalism in the code

These streams of AI history—corporate decision-making, state power and the application of statistics to war—have not survived in the public understanding of AI.

Instead, news of technical breakthroughs or pundits voicing fears are accompanied with imagery, if not of a heavily-armed Terminator, then of the brain, a robot, neon-colored microchips or absurd mathematical equations. Each is a not-so-subtle appeal to the authority of the natural sciences or computer science over that of, say, the “soft” sciences, to borrow Simon’s terminology, of political science, management science or even economics, the field for which he trundled off to Stockholm to collect his Nobel prize.

Perhaps as a result of this misguided impression, public debates continue today about what value, if any, the social sciences could bring to artificial-intelligence research. In Simon’s view, AI itself was born in social science.

David Runciman, a political scientist at the University of Cambridge, has argued that to understand AI, we must first understand how it operates within the capitalist system in which it is embedded. “Corporations are another form of artificial thinking-machine in that they are designed to be capable of taking decisions for themselves,” he explains.

“Many of the fears that people now have about the coming age of intelligent robots are the same ones they have had about corporations for hundreds of years,” says Mr Runciman. The worry is, these are systems we “never really learned how to control.”

After the 2010 BP oil spill, for example, which killed 11 people and devastated the Gulf of Mexico, no one went to jail. The threat that Mr Runciman cautions against is that AI techniques, like playbooks for escaping corporate liability, will be used with impunity.

Today, pioneering researchers such as Julia Angwin, Virginia Eubanks and Cathy O’Neil reveal how various algorithmic systems calcify oppression, erode human dignity and undermine basic democratic mechanisms like accountability when engineered irresponsibly. Harm need not be deliberate; biased data-sets used to train predictive models also wreak havoc. It may be, given the costly labour required to identify and address these harms, that something akin to “ethics as a service” will emerge as a new cottage industry. Ms O’Neil, for example, now runs her own service that audits algorithms.

In the 1950s, after having coined the term “artificial intelligence” for the Dartmouth conference, John McCarthy, one of the field’s early pioneers, wrote in his notes: “Once one system of epistemology is programmed and works, no other will be taken seriously unless it also leads to intelligent programmes.” By this view, DeepMind’s initial slogan, “Solve intelligence. Use that to solve everything else”, looks quasi-imperial.

McCarthy’s suggestion was that influence, not authority, could decide the scientific consensus in his field. DeepMind doesn’t have to “solve” intelligence (assuming such a thing is even possible) it just needs to outshine the competition. That the company’s new slogan is, “Solve Intelligence. Use it to make the world a better place,” suggests that it too is aware of the need for diplomacy in this era’s AI-powered vision of totality.

🔴 **Many fears of intelligent robots are the same as ones held about corporations for hundreds of years. We never learned to control them.**

Stephen Cave, director of the Leverhulme Centre for the Future of Intelligence, has shown that the definition of intelligence has been used throughout history as a tool for domination. Aristotle appealed to the “natural law” of social hierarchy to explain why women, slaves and animals were to be subjugated by intellectual men. To reckon with this legacy of violence, the politics of corporate and computational agency must contend with profound questions arising from scholarship on race, gender, sexuality and colonialism, among other areas of identity.

A central promise of AI is that it enables large-scale automated categorisation. Machine learning, for instance, can be used to tell a cancerous mole from a benign one. This “promise” becomes a menace when directed at the complexities of everyday life. Careless labels can oppress and do harm when they assert false authority. In protest at inadequate labels that are used to “know” the world, many young people today proudly defy unwelcome categorisations, be they traditional gender binaries or sexual binaries.

Machines who think again

It may come as a surprise that there is a lack scholarship on the social, material and political histories of the origins of artificial intelligence. Indeed, a great deal has been written about the history of AI—by Simon in 1996 and Newell in 2000, among others. Most of these histories, however, follow a narrow mould, seeing it “mainly in intellectual terms,” in the words of Paul Edwards, a historian of information technologies.

The two quasi-official histories of AI are each a history of ideas: Pamela McCorduck’s “Machines Who Think”, which “forged the template for most subsequent histories” after its initial publication in 1979; and Daniel Crevier’s “AI: The Tumultuous History”, published 1993. Both books relied primarily on in-depth interviews with key researchers to construct their

narratives.

“ The politics of corporate and computational identity become the politics of the new era

Neither, perhaps as a result, sought to understand AI in its broader context, embedded in the rise of operations research, “big science,” the actuarial sciences, and American military funding as it has evolved since the second world war. Expunged from these histories, AI can appear divorced of its historical and political context.

Without this context, AI can also appear divorced from the knowledge systems that created it. In his 1957 talk to operations-research professionals, Simon celebrated the diversity of his field's past. He described contributions from French weavers and the mechanics of the Jacquard loom, as well as from Smith, de Prony, Babbage and his peers in the “soft” sciences, as adding up to a “debt” that remained to be repaid.

That new knowledge could come about so unexpectedly, and from so many places, was what excited Simon about his work—and can stimulate us to think similarly today. Modern AI can do more than mirror the organisational dogma that characterised its birth, it can also reflect our humanity.


Jonnie Penn ([@jonnienenn](#)) is a doctoral candidate studying artificial intelligence at the University of Cambridge, and is an affiliate at Harvard University's Berkman Klein Centre for Internet & Society. He is the project development lead on “AI: History” at the Leverhulme Center for the Future of Intelligence at Cambridge. He has previously held fellowships at Google and the MIT Media Lab.

Open Future

Nov 26th 2018 | by BY JONNIE PENN

Reuse this content

About *The Economist*



Not so silky

Italy's plan to join China's Belt and Road Initiative ruffles feathers

At home and abroad

Management algorithms

Uber drivers demand their data

Brextension time

Theresa May asks to extend Article 50 to push her deal through Parliament

Orbán decay

The European centre-right suspends Viktor Orbán's Fidesz

CHARLEMAGNE'S NOTEBOOK

Classified ads

Subscribe

Group subscriptions

Contact us

Help

Keep updated



Sign up to get more from *The Economist*

Get 5 free articles per month, daily newsletters and more.

Email address

Sign up

About *The Economist*

[Advertise](#)

[Reprints](#)

[Careers](#)

[Media Centre](#)

Published since September 1843 to take part in
*“a severe contest between intelligence, which presses forward,
and an unworthy, timid ignorance obstructing our progress.”*

[Terms of Use](#)

[Privacy](#)

[Cookie Policy](#)

[Manage Cookies](#)

[Accessibility](#)

[Modern Slavery Statement](#)

Copyright © The Economist Newspaper Limited 2019. All rights reserved.