AI & SOCIETY
https://doi.org/10.1007/s00146-023-01706-0

CURMUDGEON CORNER

The neural democratisation of AI

Michael Bain[1] · **Allan McCay**[2]

Received: 14 May 2023 / Accepted: 30 May 2023

1 Introduction

Developments in generative AI, in particular the widespread deployment of OpenAI's chatGPT and Google's BARD, have drawn attention to the power of a small number of key commercial players in the development of powerful AI systems. There is of course a sense in which some of these products have democratised AI by bringing the capacity to access and employ AI systems to large numbers of organisations and individuals around the world. However, there is also an important sense in which AI has not been democratised.

While there may be signs that the power to *adapt* high quality AI systems has spread a little due to the increasing availability of open-source large language models, such as LLaMA from Meta, the power to *create* them is still not very dispersed and remains concentrated in the hands of a few. A disconcertingly small number of actors control the direction in which powerful AI systems are developing, which is troubling given the very significant impact they may have on our lives.

However, developments in neurotechnology might affect this concentration. Whilst there is a widespread discussion of the ethical, social and legal implications of AI, and also now consideration of how neurotechnology might have important consequences, less has been said about how these two related and overlapping fields might *interact* with each other in ways that might also have significant implications.

In what follows it will be suggested that advances in neurotechnology might lead to developments in AI that might in turn present a challenge to big tech and lead to a diffusion of power. In other words, *there might be neural democratisation of the power to create high quality AI systems* which are easy to use, accurate and reliable.

Before explaining how things might move towards neural democratisation we might step back in history and say something about how the most useful forms of AI came to be concentrated in the hands of a few large companies.

2 The AI landscape

Recent hype in AI is based on real progress on machine learning algorithms known as neural networks. But this is a quantitative evolution rather than a qualitative paradigm change: the core procedure to train neural networks has been around in some form for over five decades. The advances have been in computer hardware and software, and the availability of data to train networks, driven by the vast new industries of computer gaming and the Internet. Over the past decade there has been unprecedented investment in neural network algorithms by "Big Tech" with the aim of extracting more value from the massive quantities of data uploaded to the web.

This new power to create AI based on machine learning is, however, *limited by access to the technology, the data and people* with the requisite knowledge and skills. It is perhaps noteworthy that two[1] of the Turing award (the "Nobel prize for computing") winning trio of deep learning researchers were hired by Google and Facebook, very well-capitalised companies whose business model is advertising based on user data.

✉ Michael Bain
m.bain@unsw.edu.au

Allan McCay
allan.mccay@sydney.edu.au

[1] School of Computer Science and Engineering, University of New South Wales, Sydney, NSW 2052, Australia

[2] Sydney Law School, University of Sydney, Sydney, NSW 2006, Australia

[1] Although Geoff Hinton recently announced his retirement from Google.

Published online: 27 June 2023

3 Neurotech towards AI democratisation

What about AI for smaller companies? The difficulty for such organisations is the demanding economics of today's neural networks. Details are not often revealed, but producing large-scale high-quality neural network-based AI from scratch can require a multi-million dollar investment, far beyond the reach of many small companies. This is because the path of improvement towards better performing AI systems has, for the most part, been to scale them up. This means companies train larger networks to increase the "intelligence" of AI. The result is a race for ever more data, larger models and supercomputers. But continuing to scale up indefinitely the current approaches at the current rates will not be an option. So what new paths to improvement could there be?

At the heart of the achievements of current neural networks is a conundrum, because they rely on replicating in computers billions of copies of a simplified mathematical neural model bearing little resemblance to the real biological neurons used for intelligence of the natural variety throughout the animal kingdom.

These elegant structures, found from the simplest worm to the brain of a Nobelist, operate quite differently from those in today's commercial AI. Strikingly modest in their need for power and data, biological neurons enable organisms simple and complex to learn quickly about their environment (if they did not, survival prospects would be dim).

Neuroscience is the source of the "neural" metaphor for neural networks so, as our knowledge of the brain continues to advance, will we continue to return to neuroscience for further AI inspiration? We might well do.

The human brain has about one thousand trillion connections, three orders of magnitude larger than today's largest neural networks, but runs on twenty watts of power. The possibility that it is simply the additional size that is responsible for the ways in which human intelligence still exceeds the limits of what current neural networks can do cannot be discounted. And if true it is likely to challenge the thesis being advanced here, that of AI democratisation, since we will need to train networks that are at least one thousand times larger.

There are other possible challenges to the democratisation of AI thesis being advanced here. The strengths of current brain science, roughly speaking, are in understanding processes both at the level of individual neurons and at the level of cognition. Unfortunately, this is of little immediate help. At the level of individual neurons, the limitations of the artificial kind have already been mentioned, but building neural networks out of more biologically realistic ones has not yet proved to be a better approach. Neither, on the other hand, have contributions from cognitive science had much impact on today's neural networks.

However, and very importantly, something that could well lead to the neural democratisation of AI would be a better understanding of how the brain actually organises individual neurons to implement cognition.

One reason for the current gaps in our understanding of the brain is the difficulty of measuring brain activity. However, here we can expect rapid progress to be made as neurotechnology advances. An example is the development of brain-computer interfaces which allow the brain directly to control devices by the activity of human thought. The success of such technology fundamentally depends on the data from brain activity sensors, and this data can be analysed to improve the science of brain structure and function. Furthermore, it is notable that this technology is advancing beyond the medical space, with current applications in a variety of fields including gaming, wellness, and workplace safety, and so the range of different human activities that might produce useful data is increasing.

It is hard to predict the extent to which adoption of neurotechnology into patients, workers and consumers' lives will improve our understanding of the brain. Nevertheless, it seems quite possible that in making the full inner workings of the human brain more transparent this will inspire new approaches to machine learning. There may well be other efforts towards AI democratisation that do not rely on better understanding of the brain. Nonetheless, the possibility that neurotechnology will help in deciphering the algorithms of the brain and lead to better AI is tantalising.[2]

If these ideas enable us to create high-quality intelligent systems with multiple trillions of connections that learn from small amounts of data and run on low power we will have started on the journey to the neural democratisation of AI.

4 Implications

But will this advantage smaller companies? It seems possible that they might be able to use the neurotech-inspired insights to more cheaply produce high quality AI and thereby challenge the existing oligopoly of AI creation.

Of course, it seems possible that existing large AI entities might try to hold on to power and thereby maintain the concentration, perhaps by acquiring smaller rivals as they start

[2] As suggested by leading neuroscientist Professor Rafael Yuste in "Neuro-Rights and New Charts of Digital Rights: A Dialogue Beyond the Limits of the Law". Rafael Yuste, Tomás de la Quadra-Salcedo, and Miguel García Fernández, Indiana Journal of Global Legal Studies, Volume 30, Issue 1, 2023, pp. 15–3.

to become successful. However, there might be a limit to how achievable this is and of course this would be contingent on the extent of the neural democratising effect advanced here. If the effect is strong enough it might become difficult to hold on to power as too many competitors emerge.

But let us now assume that neural democratisation was to emerge as a force in the AI ecosystem. There might be some consequential implications to such an emergence.

One impact might be financial—it might have implications for the valuation of big tech players if they start to have to compete with smaller, but now credible rivals, that have benefitted from the neurotech-inspired advances in AI.

Another impact might be regulatory. Although there is something disconcerting about the existing concentration of power, a neurally democratised AI ecosystem might present more of a regulatory challenge than one in which power is concentrated in the hand of a few large and easily identifiable actors. If progress is made towards more powerful AI systems, perhaps even artificial general intelligence, then presumably issues of AI alignment and AI safety will become more prominent.

A neurally democratised economy might make it harder to for governments to ensure that AI systems are safe and aligned to values which are endorsed by the community, by presenting them with too many regulatory targets. However, although regulators might be presented with a daunting array of producers of powerful AI, any regulatory efforts might at least be less likely to be thwarted by the lobbying efforts of one or a few large and powerful AI companies who are pushing back against regulation. This might make the regulatory project in one way, more feasible even if in another, more daunting.

Advances in neurotechnology might thus lead to the neural democratisation of AI but this development might be a mixed blessing from the perspective of regulators and those who are concerned about the possible harms that AI might produce. It seems possible that in a neurally democratised world, it might not just be the existing big players that wished things had stayed undemocratic.

Curmudgeon Corner Curmudgeon Corner is a short opinionated column on trends in technology, arts, science and society, commenting on issues of concern to the research community and wider society. Whilst the drive for super-human intelligence promotes potential benefits to wider society, it also raises deep concerns of existential risk, thereby highlighting the need for an ongoing conversation between technology and society. At the core of Curmudgeon concern is the question: What is it to be human in the age of the AI machine? -Editor.

Data availability statement The manuscript has no associated data.

Declarations

Conflict of interest On behalf of all authors, the corresponding author states that there is no conflict of interest.