International Journal on Neuropsychology and Behavioural Sciences

Artificial Intelligence: Threat or Promise?

Opinion

Volume 5 Issue 3- 2024

Author Details

*Keith Oatley**

Department of Applied Psychology & Human Development, University of Toronto, Canada

*Corresponding author

Keith Oatley, Department of Applied Psychology & Human Development, University of Toronto, Canada

Article History

Received: August 19, 2024 Accepted: August 27, 2024 Published: August 27, 2024

Opinion

At the beginning of this century very few people had heard of artificial intelligence. Now many know something about it. With the coming of artificially intelligent language processors such as ChatGPT which or who can write about anything, a lot of people think that a threat may be approaching. Are these new kinds of artificially intelligent beings more intelligent than us? Are they going to take over the planet? In The Mind's New Science [1]. Gardner argued that the foundation of cognitive science, including cognitive psychology, began at a conference in 1956 at which presentations were given by Noam Chomsky, Allan Newell with Herbert Simon, and George Miller. The idea at the centre of this was that human brains work like computers. Researchers were creating programs that were doing something like thinking. They did so by getting their systems to proceed by a kind of logic: "if this is the case then do that." The consensus was that this way of understanding psychology was better than any previous basis.

At about the same time as Gardner's book came out, a small group of researchers thought that this way of understanding human minds was wrong. Instead, this group argued that, instead of rules and logic, a different approach could use structures like human neurons, activations of which occur by means of the strengths of the connections among them. In this approach, computers would work like human brains. Here, including the diagram on this page, is how I put it in Our Minds, Our Selves.

Imagine that in a network of the kind pictured, the pieces of information given to the input neurons (at the bottom of the diagram) are digits (ten of them rather than just the three represented in the figure 1), and that the network's task is to classify each as odd or even. Imagine digits being offered to the system, one by one, as activations of one of the ten input neurons (pp. 113-114).

Figure 1:

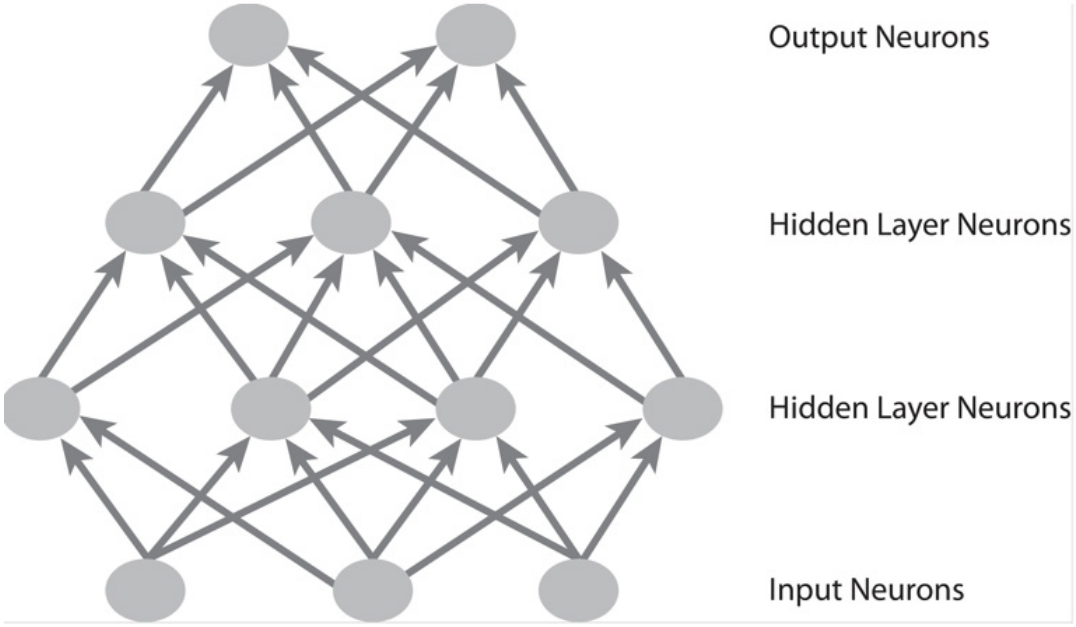

If a digit 2 is input, let's say by activating the neuron in the middle of the bottom row, the programmers will have arranged that the top right neuron is also activated, meaning that this is an even number. The connections between the middle neuron at the bottom and the one at the top right are strengthened. For odd numbers, the output is the top left neuron. Then comes back-propagation in which the programmers send iterations from the output neurons at the top in the opposite direction from those of the arrows in the diagram, strengthening the best pathways for even digits, then for the odd ones, and minimising errors of classification. The system learns by procedures that are like the behaviorist idea of reinforcement: forward and backward passes occur with gradually improving results that need an outside agent to identify the correct outcomes.

Geoffrey Hinton, known as the godfather of artificial intelligence, moved on from this and, along with students and colleagues, invented what they call deep learning [2] which did away with processes that involved outside agents such as human programmers having to identify which digits were even and which were odd. The new system was able to learn by itself, making its own generalisations, for instance within a set of pictorial images. For the visual task, the system constructs what we might think of as intuitions: that images of one kind have certain features in common, and images of another kind have different features in common. Hinton worked on this with a set of several million images provided by computer scientist Fei-Fei Li: the Image Network. The artificially intelligent system was able to classify them. Then, without doing anything to the network as such, human labels could be offered for different sets: this set of images is of cats and that set is of dogs. Here's something Li wrote in her recently published book; The Worlds I See [3].

I believe our civilization stands on the cusp of a technological revolution with the power to reshape life as we know it. It must respect the collective dignity of a global community. And it must always remember its origins: the restless imagination of an otherwise unremarkable species of hominid. This revolution must, therefore, be unequivocally human centered (p. 8).

Among human centred endeavours, artificial intelligence systems are, for instance, better than human doctors at using X-rays and other imaging devices to diagnose patients' diseases. At the same time, artificially intelligent military robots are likely to be better at fighting than humans. Because they are more clever than us and because they no longer need us to program them, is it possible as Hinton worries that they may be able to control everything in the world. Maybe they'll just dispense with humans. At the end of April 2023, reversing the direction of his research and wanting to warn the world about this, Hinton resigned as a vice-president of Google. So, intense has the debate become of catastrophe or possibility that there's a long article about Hinton (his research and his life) by Joshua Rothman in the 20th November 2023 issue of The New Yorker [4].

When we think about this, we may not realise that we ourselves are robots. We tend to think that when we have children, we pass our genes on to them. As Keith Stanovich pointed out in The Robot's Rebellion, it's the other way round. They are not our genes. Instead, they have programmed us to be their vehicles. Genes are made of strands of DNA [5]. They are the only immortals on earth, you might even think of them as gods. Their vehicles include plants as well as creatures that move about. The job for which we have all been programmed is to reproduce so that the DNA strands can continue. In evolution, as proposed by Charles Darwin, all biological species have been produced by means of three interacting processes [6]. The first, Darwin called superabundance. In reproduction each biological being produces not just one other to replace itself. It produces many more. The second is variation. The offspring have differences. Third is selection. Some of the offspring do not fit well in the niche into which they are cast. They perish without reproducing. Others fit better, and they do reproduce. They do so by being programmed by strands of DNA.

This process is a mode of exploration. What is not always pointed out is that it is random. Some offspring do fit and others do not. Evolution keeps going, often producing unfit beings over and over. We may even wonder how fit we humans are, with wars, racism, inequalities, and reluctance to look after the planet. Now, for the first time it may be possible to add an element of guidance to this process. We humans could perhaps design artificially intelligent robots with a particular purpose. What could that purpose be? An excellent answer is that proposed by Michael Tomasello, see for instance Becoming human: A theory of ontogeny. In what may be the most important line of psychological research of this century, he and his collaborators have shown that, as compared with human infants, chimpanzees of any age can do as well at physical tasks, such as individually looking for something that is not immediately visible [7]. They are, however, far less good (actually rather hopeless) at any task that requires cooperation [8] Tomasello therefore proposed that a fundamental move has occurred in recent human evolution. It's the ability to make arrangements and cooperate with each other. This occurred, Tomasello proposed, in two stages. He called the first stage joint intentionality, and the second collective intentionality. As to the first, chimpanzees roam through their territories in groups. When they find a source of food, they grab what they can, with the alpha male getting first pick. Each individual then goes and eats alone. By contrast, humans forage together, gathering food into one of the earliest pieces of technology the bag then taking it to where the group is, to share it and eat together. The second stage, collective intentionality, has led to human cultures.

The task for the future of super-intelligent robots, therefore, is for humans to guide their design in this kind of direction: cooperatively caring for themselves and caring also for us [9].

References

1. Gardner H (1985) The mind's new science: A history of the cognitive revolution. New York: Basic Books.
2. LeCun Y, Bengio Y, Hinton GE (2015) Deep learning. Nature 521: 436-444.
3. Li FF (2023) The worlds I see: Curiosity, exploration, and discovery at the dawn of AI. New York: Flatiron Books.
4. Rothman J (2023) Metamorphosis. The New Yorker 28-39.
5. Stanovich KE (2004) The robot's rebellion: Finding meaning in the age of Darwin. Chicago: University of Chicago Press.
6. Darwin C (1859) On the origin of species by means of natural selection. London: Murray.
7. Tomasello M (2019) Becoming human: A theory of ontogeny. Cambridge, MA: Harvard University Press.
8. Herrmann E, Call J, Hernandez Lloreda, MV, Hare B, et al. (2007) Humans have evolved specialized skills of social cognition: The cultural intelligence hypothesis. Science 317(5843): 1360-1366.
9. Oatley K (2018) Our minds, our selves: A brief history of psychology. Princeton, NJ: Princeton University Press.