Journal Pre-proof

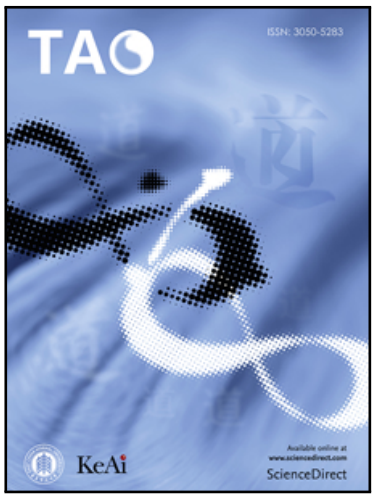

2062: Where is AI taking us?

Toby Walsh

PII: S3050-5283(25)00018-8
DOI: https://doi.org/10.1016/j.tao.2025.100020
Reference: TAO 100020

To appear in: *TAO*

Received date: 11 June 2025
Revised date: 16 July 2025
Accepted date: 31 July 2025

Please cite this article as: Toby Walsh , 2062: Where is AI taking us?, *TAO* (2025), doi: https://doi.org/10.1016/j.tao.2025.100020

Perspective

2062：Where is AI taking us?

Toby Walsh

The Department of Computer Science and Engineering, the University of New South Wales

High St Kensington, NSW 2052, Australia

E-mail address: tw@cse.unsw.edu.au

Received 11 June 2025

Revised 16 July 2025

Accepted 31 July 2025

Abstract: This perspective paper explores the trajectory of Artificial Intelligence (AI) toward the year 2062, the median date experts predict machines will match human cognitive capabilities. Drawing upon Neil Postman's framework for technological change, the author analyzes the potential impacts of AI. The discussion posits that AI operates as an ecological rather than merely additive force, fundamentally altering the human environment, much as the advent of electricity did. The analysis applies five core lessons to the AI context: the inherent trade-offs of technology, illustrated by how navigation aids may erode spatial memory; the unequal distribution of benefits and costs, evidenced by biases in predictive policing; the epistemological shift suggesting that "thinking is computational"; the unpredictability of systemic change; and the rejection of technological inevitability. The paper concludes that while AI promises immense economic utility—projected to contribute $15 trillion to the global economy by 2030—navigating its rise requires critical human agency to maximize positive outcomes while mitigating ethical risks and cognitive atrophy.

Keywords: Artificial intelligence, Technological change, Neil Postman, AI Ethics

Thank you very much. Distinguished guests, ladies and gentlemen, it is a pleasure to talk to you today. I want to take this opportunity to pick up some of the ideas—some of the themes—that Prof. Raj Reddy introduced this morning again.

First, why did I title this talk "2062"？The number, which is also the title of a book I recently

published, refers to the year 2062 and was derived from a survey of 300 of my colleagues, i.e., other artificial intelligence experts worldwide [1]. The survey comprised a simple question: When will computers match all human cognitive capabilities? Now, I have to be honest with you: they gave me a lot of different answers that varied tremendously. However, the median year that they predicted was 2062. I think it's interesting that none of them predicted it would take merely five or ten years. Equally surprising, given the scale of the problem, is that very few of them predicted that it would take more than 100 years for computers to attain human-level capacity.

I think it is easy to be seduced by the success of large language models like ChatGPT—to believe that we are perhaps closer to the end goal of matching and possibly exceeding human intelligence in machines. As Prof. Raj Reddy suspects, computers will not stop evolving once they reach human intelligence; they will soon exceed it. As the President Qian of ECNU (East China Normal University) said this morning, there is much more to intelligence than just language. Hence, while we have language models that are pretty capable of language, they still have many significant limitations in their ability to reason and think. So, we perhaps still have some decades to go before they match us in those capacities. However, when their capabilities match and eventually exceed ours, it will be remarkable. These developments will enable transformations in many aspects of our lives.

Andrew Ng, whom many of you know, was the chief scientist at *Baidu* and co-founder of Coursera. He has compared it, as many people have, to the transformation that electricity brought to our lives. Look around this room now and imagine the possibility of living without the electricity that provides the power for—and the data in—so many of our devices. As Andrew suggests, it is hard to think of any part of our lives that AI won't touch [2]. By 2062, the world we live in with AI will be vastly different.

Of course, some of those returns are immediately visible, even with the limited AI that we have today. We are already seeing significant impacts and returns on our investments in what we can do with that limited intelligence. A study was conducted by PwC in 2017, before the pandemic. The numbers they produced are perhaps a bit off now; nevertheless, they give us an idea of the scale of the opportunity. The study estimated that AI would add approximately $15 trillion to the world's

economy in inflation-adjusted terms by 2030 [3]. Now, I think most of us cannot believe in trillions. Most of us do not even think in billions. I sometimes struggle to even think in terms of *millions*. Again, the estimate is 15 trillion dollars in inflation-adjusted terms. To put that into context, the combined current GDP of China and India amounts to approximately $ 15 trillion. It's quite a significant number and represents a substantial transformation. Quite, indeed. In addition, the world's economy has discovered a new technology that is projected to add $ 15 trillion to the combined GDP of China and India. There are a few other things—in fact, I can think of no others—that will have such a significant impact in the next couple of decades. And, of course, given the scale of that impact, people are not getting nervous about the consequences.

People are developing new ethical frameworks to understand the types of questions and the kinds of impacts that the technology will have. In many cases, we do not need to ask new questions. We can ask the same questions we asked about electricity, the same questions we asked about the steam engine, and those we should have asked about the steam engine. Like any technology, it will perform very similar tasks. That does not mean there are no new things to worry about. There are new things to be concerned about because technology is going to change the speed, scale, and cost at which we can do things. It will break things. Therefore, in new ways, even if the questions we should ask are very much like the questions we should have asked all along, we need to identify those questions: *What are the questions that we should ask? What are the impacts? How do we think about this significant transformation happening in our lives?*

However, there is also a beautiful essay I encourage you to take a look at, written in 1998, before AI became such a significant factor. It was not written as an essay, nor was it about artificial intelligence. The essay speaks about computing in general and was written by Neil Postman, a famous humanist, for a commencement address he gave in 1998. To find copies of this speech in PDF format on the internet, just search for "five things to know about technological changes," which is its title [4]. It is a fantastic speech. Today, I will draw on those five concepts but discuss them in the context of artificial intelligence, as I believe this is an excellent framework for understanding **the potential impacts of AI**. It allows us to thoughtfully consider these impacts and strive to maximize positive outcomes while minimizing negative ones. **The five lessons** I am going to offer you comprise the total of my talk.

The first lesson is very much in the spirit of the Tao of AI, the *Yin* and *Yang.* Is technological change driven by AI a trade-off? We should all ask ourselves, "What does it give us? What does it take from us?" And, most importantly, "What does it take away?" The example I want to give you is perhaps the most common use of AI in your life today. Suppose you get directions on your smartphone using one of the mapping services available today—Apple Maps, Google Maps, or the Chinese equivalent—or you are getting directions on your car's GPS. In that case, you are utilizing a little AI algorithm that was invented in the 1970s. The algorithm has a name: "A* search," invented for SHAKY, the robot at the Stanford Research Institute in the 1970s [5]. This AI has now been repurposed to direct humans. It works out the quickest way to get from A to B, taking into account the real-time traffic conditions that we now have. I spend a lot less time getting lost, and I get to my destination more frequently on time, courtesy of all this AI power and navigation.

In this case, it is clear what it gives us, but what does it take away? In the UK, there is a thing called "the Knowledge." If you want to become a taxi driver in London, you must spend a year-long test learning the streets of London. You have to know how to get from any two places and the shortest route, in terms of distance, between any two locations. It takes a year of your life to learn this. It requires driving around on a moped for a year to learn the streets of London. When you get into a black cab, you can be assured that you will be taken to your destination in the quickest possible way.

Brain studies using fMRI (functional magnetic resonance imaging) have been conducted on people while taking the black cab driving test. The researchers measured the hippocampus, the part of the brain that is responsible for spatial reasoning and spatial memory. The hippocampus in people who use "the Knowledge" gets 15% larger. However, if we are not learning the streets, we are not navigating for ourselves; if we outsource activities to machines, then our brains, our hippocampi, are not getting any bigger. It is actually taking something away from us. We should worry that AI, more generally, may "dumb us down." If we outsource so many of our cognitive tasks to machines, what will happen to us? While we get these wonderful tools that make our lives more convenient, we may ourselves be changed and may be dumbed down. That was the first

lesson: Every technological change is a trade-off. AI is going to be a trade-off, as well as the *Yin* and *Yang*.

The second lesson from Neil Postman's essay is that the benefits and costs are not evenly distributed. There will be winners, and there will be losers, again, very much in the spirit of the Tao of AI. I would like to illustrate this point about AI using examples, such as AI for predictive policing, which is increasingly common in many countries, particularly in the United States. For predictive policing, AI predicts where crime is going to occur. Most countries have limited resources for police to investigate crime. Therefore, using AI in this way sounds like a good idea. Indeed, it is a rather good idea to say, "Actually, we've got lots of data on where crime took place; let's use those predictions." Specifically, by using machine learning to predict where crime will take place, we can direct our limited resources to the areas of greatest need.

This sounds like a perfect idea, although it is fundamentally flawed, like so many things we do in relation to "artificial intelligence," because it is based on historical data and therefore inherits the flaws and falsehoods of the past. Maybe we do not actually know what is true about circumstances "on the ground," so to speak. We do not actually know where some crimes took place. We only have a proxy for where people were arrested and charged for a crime, and "approximate" is not "exactly the same." For example, more police patrols may be active in certain regions, such as poorer neighborhoods and/or black neighborhoods, in cities in the United States, and, therefore, arrest rates for poorer black people may be disproportionately high. In such cases, we are prosecuting poverty and race, not crime. If we are not careful, we are going to perpetuate those biases into the future with these AI tools. There are winners, and there are losers. We have to be well aware that some groups may be losing out.

The third of his five lessons is that every technology contains a powerful idea, and that AI also contains a beautiful and immensely powerful one. That is the idea that thinking is computational. Anything that we can do with our brains, the machines—computers—could also possibly do. This idea is dangerously seductive because we tend to then project onto the machine the other traits that make humans special. Computers are not normal beings. They do not suffer. They do not have our consciousness. They are not aware. They are not things that can be punished. Therefore, we

should be cautious when, as Professor Raj Reddy discussed, considering transforming things like warfare by handing over those decisions, those life and death decisions, to machines. Furthermore, it is misguided to claim, as Yuval Harari said, that comparing us to AI will elevate us to the status of gods. We will remain mortal; we will remain human. While our machines may possess certain godlike properties, they, too, are ultimately mortal. Our own mortality remains the most crucial criterion.

The fourth of the five lessons is that technological change is not additive. It is vast, unpredictable, and even exponential. We do live in exponential times. We see this, undoubtedly, with artificial intelligence, which makes it so hard to predict the future. The example I would like to share with you is from 1969. In 1969, the first jumbo jet took off from Boeing's works in Seattle. It was very easy to predict the impact that technology was going to have on the world: It was going to "shrink" the world. It was going to give us the jet age, as it did. It was going to make transportation much cheaper. It was going to lower the barriers between our nations. It was going to introduce international trade and international tourism. Those effects were easy to predict.

It has, however, been forgotten that in 1969, we were at the height of the previous pandemic. It was widely referred to as the Asian flu, although it is unclear whether Asia was truly responsible for the fatality rate. Nevertheless, it was known as the Asian flu pandemic and resulted in the deaths of approximately 1 million older adults worldwide. At the time, few were discussing the actual impact of the jumbo jet, which exacerbated the next pandemic—the one we have just experienced—by rapidly spreading the virus across the globe as air travel rose in the jet age. In nearly every country, we closed our borders, not only to isolate ourselves but also to contain the virus. Predicting such changes is extremely difficult.

Another example I want to talk about includes things like the Facebook algorithm and social media. People working for companies like Facebook wanted to encourage meaningful engagement with the platforms. Again, they did not have this grounded in truth. They only had proxies for likes and clicks, which they encouraged by incentivizing them. Of course, we know now where that took us, and the unexpected consequence of that was that it encouraged polarized speech. It resulted in things like the Brexit referendum and the election of Trump, the vast nastiness that we

now find on social media, and the ability of companies like Cambridge Analytica to manipulate what people want to do and how they vote. Again, it is very hard to predict those unexpected consequences.

The fifth and final lesson I want to share with you, from Neil Postman's essay, is that it is easy to think technology is inevitable. Technology is not, however, part of any natural order, and nothing is inevitable. It is too often idolized. We get to make the choices about where and how technology is introduced into our lives. We usually forget that. I want to remind you that we are talking about "artificial intelligence." When people talk about "artificial intelligence," or "AI," they tend to focus on the second word, "intelligence," because, of course, that is *our* most important characteristic—that we are intelligent. That is, in some sense, it feels like the most important characteristic of the "artificial intelligence" we are building is that it is going to be somewhat intelligent, but you should not forget the other word there, "artificial." I am increasingly of the view that AI will play an incredibly important role. We should remind ourselves that the "intelligence" that we build into machines is very artificial, very different from human intelligence. There is no reason to suppose it is going to be the same as human intelligence.

Meanwhile, other forms of intelligence also exist on our planet. The octopus is my favorite example. Anyone who has worked with octopuses knows how intelligent they are. They are very intelligent and seemingly quite conscious beings. Scientists who work with them say that these animals can recognize human faces and recognize people who enter the room when they are in a tank. Being famous escapologists, they are very smart. One measure of intelligence is the ability to use tools. Octopuses demonstrate this capability; they can be trained to open jars to obtain food. By many criteria, they are considered highly intelligent, but their intelligence is fundamentally different from human intelligence.

Now, of course, all life on Earth, as far as we understand, is related to all other life—all ultimately descended from the same genetic material. However, if you want to trace back the tree of life to find where we are related to octopuses, you have to go back a very long way. Because we are mammals, we belong to the mammal branch of the animal kingdom, while they belong to the vertebrate branch. You have to go back hundreds of millions of years to find where we have a

common ancestor, and that's where life started to become multicellular. There was no intelligence then. There were barely enough cells to make anything interesting. So, whatever intelligence the octopuses have, it evolved completely separately from human intelligence. If you want to understand what alien intelligence looks like, our best guess is that it could resemble that of the octopus. Being an octopus must be pretty different from being human. An octopus has 60% of its brain distributed into its eight legs. It essentially has nine brains and, hence, possesses a much more distributed intelligence. It would be interesting if we could ever understand what it is like to be an octopus.

The other metaphor I want to leave you with concerns flight, both artificial and natural. "Natural flight' refers to that of birds and bees—you know, flapping wings, feathers, and all that sort of stuff. "Artificial flight," such as what we do in airplanes, as I will be doing this evening, involves having a big, fixed wing and a large jet engine behind you. These examples represent two completely different solutions to the same problem: bees and airplanes are heavier than their own motion. We have to resolve the same laws of physics, the Navier-Stokes equations of aerodynamics, but apply a completely different engineering solution to each example. The early indications that we have from the limited AI built today are, again, that we need a completely different solution to the same problem. It is wrong yet very easy to try to apply our experience of intelligence, our human intelligence, and project that understanding onto machines and think that their intelligence is going to be very similar to ours, rather than to realize now that it is probably going to be wildly dissimilar, just as artificial flight differs from natural flight.

Finally, I want us to remember that AI requires us to think about the impact it will have, the way it will break things, as well as do things faster, cheaper, and more efficiently. We must still apply many of the questions that we have asked for thousands of years about any other technology. If you would like to know more about these AI-related topics, I have written a couple of books about the ethical questions of machines behaving badly and about the future of 2062, as mentioned in the introduction. My most recent book is about generative AI and how all the AI fakes. Fake AI is going to be increasingly deceiving us. I know that some of you would prefer not to read books in English, so I am pleased to tell you that most of these books are also available in Chinese. A great pleasure to talk to you, and thank you very much!

Declaration of competing interests

The author declares that he has no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

The author is also the editorial board member of the journal, but was not involved in the peer review process. This paper was handled by the Editor-in-Chief.

Editor's note

This perspective paper is adapted from the transcript of the author's keynote speech at the First International Forum on Quantum and Laozi, held in Shanghai, China, on July 7th, 2024. It has undergone the same rigorous peer-review process as regular submissions to ensure its quality and relevance.

References

[1] T. Walsh, 2062: The World that AI Made, La Trobe University Press, Melbourne, 2018.

[2]K.McAleer, "AI is the New Electricity": Insights from Dr. Andrew Ng. <https://scet.berkeley.edu/ai-is-the-new-electricity-insights-from-dr-andrew-ng/>, 2023. (accessed 09.06.24)

[3] PwC, AI to drive GDP gains of $15.7 trillion with productivity and personalization improvements. <https://www.pwc.com/hu/en/pressroom/2017/ai.html>, 2017. (accessed 09.06.24)

[4] Neil Postman, Five Things We Need to Know About Technological Change. <https://web.cs.ucdavis.edu/~rogaway/classes/188/materials/postman.pdf>，1998. (accessed 09.06.24)

[5] Wikipedia, A* search algorithm. <https://en.wikipedia.org/wiki/A*_search_algorithm>, 2023. (accessed 09.06.24)