

The Guardian

The Coming Wave by Mustafa Suleyman review – a tech tsunami

The co-founder of DeepMind issues a terrifying warning about AI and synthetic biology – but how seriously should we take it?

Scott Shapiro

Fri 8 Sep 2023 04.00 EDT Last modified on Wed 13 Sep 2023 09.38 EDT

On 22 February 1946, George Kennan, an American diplomat stationed in Moscow, dictated a 5,000-word cable to Washington. In this famous telegram, Kennan warned that the Soviet Union's commitment to communism meant that it was inherently expansionist, and urged the US government to resist any attempts by the Soviets to increase their influence. This strategy quickly became known as "containment" – and defined American foreign policy for the next 40 years.

The Coming Wave is Suleyman's book-length warning about technological expansionism: in close to 300 pages, he sets out to persuade readers that artificial intelligence (AI) and synthetic biology (SB) threaten our very existence and we only have a narrow window within which to contain them before it's too late. Unlike communism during the cold war, however, AI and SB are not being forced on us. We willingly adopt them because they not only promise unprecedented wealth, but solutions to our most intractable problems – climate change, cancer, possibly even mortality. Suleyman sees the appeal, of course, claiming that these technologies will "usher in a new dawn for humanity".

An entrepreneur and AI researcher who co-founded [DeepMind](#) in 2010, before it was acquired by Google in 2014, Suleyman is at his most compelling when illustrating the promises and perils of this new world. In breezy and sometimes breathless prose, he describes how human beings have finally managed to exert power over intelligence and life itself.

Take the AI revolution. Language models such as [ChatGPT](#) are just the beginning. Soon, Suleyman predicts, AI will discover miracle drugs, diagnose rare diseases, run warehouses, optimise traffic, and design sustainable cities. We will be able to tell a computer program to "make a \$1 million on Amazon in a few months" and it will carry out our instructions.

The problem is that the same technologies that allow us to cure a disease could be used to cause one – which brings us to the truly terrifying parts of the book. Suleyman notes that the price of genetic sequencing has plummeted, while the ability to edit DNA with [technologies such as Crispr](#) has vastly improved. Soon, anyone will be able to set up a genetics lab in their garage. The temptation to manipulate the human genome, he predicts, will be immense.

The problem, however, is that the same technologies that allow us how to cure a disease could be used to cause one

Human mutants, however, are not the only horrors awaiting us. Suleyman envisions AI and SB joining forces to enable malicious actors to concoct novel pathogens. With a 4% transmissibility rate (lower than chickenpox) and 50% case fatality rate (about the same as Ebola), an AI-designed and SB-engineered virus could “cause more than a billion deaths in a matter of months”.

Despite these risks, Suleyman doubts any nation will make the effort to contain these technologies. States are too dependent on their economic benefits. This is the basic dilemma: we cannot afford not to build the very technology that might cause our extinction. Sound familiar?

The Coming Wave is not about the existential threat posed by superintelligent AIs. Suleyman thinks that merely *smart* AIs will wreak havoc precisely because they will vastly increase human agency in a very short period. Whether via AI-generated cyber-attacks, homebrewed pathogens, the loss of jobs due to technological change, or misinformation aggravating political instability, our institutions are not ready for this tsunami of tech.

He repeatedly tells us that the “wave is coming”, “the coming wave is coming”, even “the coming wave really is coming”. I suppose living through the past 15 years of AI research, and becoming a multimillionaire in the process, would turn anyone into a believer. But if the past is anything to go by, AI is also known for its winters, when initial promise stalled and funding dried up for long periods. Suleyman disregards the real possibility that this will happen again, thereby giving us more time to adapt to and even stem the tide of social change.

But even if progress continues its frenetic pace, it is unlikely that societies will tolerate the ethical abuses Suleyman fears most. When a Chinese scientist revealed in 2018 that he had [edited the genes of twin girls](#), he was sentenced to three years in prison, universally condemned, and there have been no similar reports since. The EU is set to prohibit certain forms of AI – such as facial recognition in public spaces – in its forthcoming [AI Act](#). Normal legal and cultural pushback will probably slow the proliferation of the most disruptive and disturbing practices.

Despite claiming that the containment problem is the “defining challenge of our era”, Suleyman does not support a tech moratorium (he did just start a new AI company). Instead he sets out a series of proposals at the end of the book. They are unfortunately not reassuring.

For example, Suleyman suggests that AI companies spend 20% of R&D funds on safety research, but does not say why companies would divert capital away from rushing their new products to market. He advocates banning AI in political ads, but doing so would violate the first amendment to the US constitution. He proposes an international anti-proliferation treaty, but does not give us any indication of how it might be enforced. At one point, Suleyman hints that the US may need to coerce other countries to comply. “Some measure of anti-proliferation is necessary. And, yes, let’s not shy away from the facts; that means real censorship, possibly beyond national borders.” I don’t know exactly what he means here, but I don’t like the way it sounds.

Suleyman pushes these costly proposals despite conceding that his catastrophic scenarios are tail risks. Yes, the probability of doomsday is low, but the consequences would be so catastrophic that we must treat the possibility as a clear and present danger. One very large elephant in the room is climate change. Unlike the AI apocalypse that may happen in the future, a climate emergency is happening right now. This July was [the hottest on record](#). Containing carbon, not AI, is the defining challenge of our era. Yet here, Suleyman is strikingly and conveniently optimistic. He believes that AI will solve the climate emergency. That is a happy thought – but if AI will solve the climate problem, why can't it solve the containment problem too?

If the book's predictions about AI are accurate, we can safely ignore its proposals. Wait a few years and we can just ask ChatGPT-5, -6, or -7 how to handle the coming wave.

Scott Shapiro is professor of law and philosophy at Yale and author of [Fancy Bear Goes Phishing](#) (Allen Lane). *The Coming Wave* by Mustafa Suleyman and Michael Bhaskar is published by Bodley Head (£25). To support the Guardian and Observer order your copy at guardianbookshop.com. Delivery charges may apply.