PART III

Computer Politics

From:

Harari, Y., 2024. Nexus. Fern, pgs 303-316.

CHAPTER 9

Democracies: Can We Still Hold a Conversation?

Civilisations are born from the marriage of bureaucracy and mythology. The computer-based network is a new type of bureaucracy that is far more powerful and relentless than any human-based bureaucracy we've seen before. This network is also likely to create inter-computer mythologies that will be far more complex and alien than any human-made god. The potential benefits of this network are enormous. The potential downside is the destruction of human civilisation.

To some people, warnings about civilisational collapse sound like over-the-top jeremiads. Every time a powerful new technology has emerged, anxieties arose that it might bring about the apocalypse, but we are still here. As the Industrial Revolution unfolded, Luddite doomsday scenarios did not come to pass, and Blake's 'dark Satanic Mills' ended up producing the most affluent societies in history. Most people today enjoy far better living conditions than their ancestors in the eighteenth century. Intelligent machines will prove even more beneficial than any previous machines, promise AI enthusiasts like Marc Andreessen and Ray Kurzweil.[1] Humans will enjoy

much better health care, education and other services, and AI will even help save the ecosystem from collapse.

Unfortunately, a closer look at history reveals that the Luddites were not entirely wrong and that we actually have very good reasons to fear powerful new technologies. Even if in the end the positives of these technologies outweigh their negatives, getting to that happy ending usually involves a lot of trials and tribulations. Novel technology often leads to historical disasters, not because the technology is inherently bad, but because it takes time for humans to learn how to use it wisely.

The Industrial Revolution is a prime example. When industrial technology began spreading globally in the nineteenth century, it upended traditional economic, social and political structures and opened the way to create entirely new societies, which were potentially more affluent and peaceful. However, learning how to build benign industrial societies was far from straightforward and involved many costly experiments and hundreds of millions of victims.

One costly experiment was modern imperialism. The Industrial Revolution originated in Britain in the late eighteenth century. During the nineteenth century industrial technologies and production methods were adopted in other European countries ranging from Belgium to Russia, as well as in the United States and Japan. Imperialist thinkers, politicians and parties in these industrial heartlands claimed that the only viable industrial society was an empire. The argument was that unlike relatively self-sufficient agrarian societies, the novel industrial societies relied much more on foreign markets and foreign raw materials, and only an empire could satisfy these unprecedented appetites. Imperialists feared that countries that industrialised but failed to conquer any colonies would be shut out from essential raw materials and markets by more ruthless competitors. Some imperialists argued that acquiring colonies was not just essential for the survival of their own state but beneficial for the rest of humanity, too. They claimed empires alone could spread the blessings of the new technologies to the so-called undeveloped world.

Consequently, industrial countries like Britain and Russia that already had empires greatly expanded them, whereas countries like the United States, Japan, Italy and Belgium set out to build them. Equipped with mass-produced rifles and artillery, conveyed by steam power and commanded by telegraph, the armies of industry swept the globe from New Zealand to Korea, and from Somalia to Turkmenistan. Millions of indigenous people saw their traditional way of life trampled under the wheels of these industrial armies. It took more than a century of misery before most people realised that the industrial empires were a terrible idea and that there were better ways to build an industrial society and secure its necessary raw materials and markets.

Stalinism and Nazism were also extremely costly experiments in how to construct industrial societies. Leaders like Stalin and Hitler argued that the Industrial Revolution had unleashed immense powers that only totalitarianism could rein in and exploit to the full. They pointed to the First World War – the first 'total war' in history – as proof that survival in the industrial world demanded totalitarian control of all aspects of politics, society and the economy. On the positive side, they also claimed that the Industrial Revolution was like a furnace that melts all previous social structures with their human imperfections and weaknesses and provides the opportunity to forge perfect societies inhabited by unalloyed superhumans.

On the way to creating the perfect industrial society, Stalinists and Nazis learned how to industrially murder millions of people. Trains, barbed wire and telegraphed orders were linked to create an unprecedented killing machine. Looking back, most people today are horrified by what the Stalinists and Nazis perpetrated, but at the time their audacious visions mesmerised millions. In 1940 it was easy to believe that Stalin and Hitler were the models for harnessing industrial technology, whereas the dithering liberal democracies were on their way to the dustbin of history.

The very existence of competing recipes for building industrial societies led to costly clashes. The two world wars and the Cold War

can be seen as a debate about the proper way to go about it, in which all sides learned from one another, while experimenting with novel industrial methods to wage war. In the course of this debate, tens of millions died and humankind came perilously close to annihilating itself.

On top of all these other catastrophes, the Industrial Revolution also undermined the global ecological balance, causing a wave of extinctions. In the early twenty-first century up to fifty-eight thousand species are believed to go extinct every year, and total vertebrate populations declined by 60 per cent between 1970 and 2014.[2] The survival of human civilisation too is under threat. Because we still seem unable to build an industrial society that is also ecologically sustainable, the vaunted prosperity of the present human generation comes at a terrible cost to other sentient beings and to future human generations. Maybe we'll eventually find a way – perhaps with the help of AI – to create ecologically sustainable industrial societies, but until that day the jury on Blake's satanic mills is still out.

If we ignore for a moment the ongoing damage to the ecosystem, we can nevertheless try to comfort ourselves with the thought that eventually humans did learn how to build more benevolent industrial societies. Imperial conquests, world wars, genocides and totalitarian regimes were woeful experiments that taught humans how *not* to do it. By the end of the twentieth century, some might argue, humanity got it more or less right.

Yet even so the message to the twenty-first century is bleak. If it took humanity so many terrible lessons to learn how to manage steam power and telegraphs, what would it cost to learn to manage bioengineering and AI? Do we need to go through another cycle of global empires, totalitarian regimes and world wars in order to figure out how to use them benevolently? The technologies of the twenty-first century are far more powerful – and potentially far more destructive – than those of the twentieth century. We therefore have less room for error. In the twentieth century, we can say that humanity got a C-minus in the lesson on using industrial technology. Just enough to pass. In the twenty-first century, the bar is set much higher. We must do better this time.

THE DEMOCRATIC WAY

By the end of the twentieth century, it had become clear that imperialism, totalitarianism and militarism were not the ideal way to build industrial societies. Despite all its flaws, liberal democracy offered a better way. The great advantage of liberal democracy is that it possesses strong self-correcting mechanisms, which limit the excesses of fanaticism and preserve the ability to recognise our errors and try different courses of action. Given our inability to predict how the new computer network will develop, our best chance to avoid catastrophe in the present century is to maintain democratic self-correcting mechanisms that can identify and correct mistakes as we go along.

But can liberal democracy itself survive in the twenty-first century? This question is not concerned with the fate of democracy in specific countries, where it might be threatened by unique developments and local movements. Rather, it is about the compatibility of democracy with the structure of twenty-first-century information networks. In chapter 5 we saw that democracy depends on information technology and that for most of human history large-scale democracy was simply impossible. Might the new information technologies of the twenty-first century again make democracy impractical?

One potential threat is that the relentlessness of the new computer network might annihilate our privacy and punish or reward us not only for everything we do and say but even for everything we think and feel. Can democracy survive under such conditions? If the government – or some corporation – knows more about me than I know about myself, and if it can micromanage everything I do and think, that would give it totalitarian control over society. Even

if elections are still held regularly, they would be an authoritarian ritual rather than a real check on the government's power. For the government could use its vast surveillance powers and its intimate knowledge of every citizen to manipulate public opinion on an unprecedented scale.

It is a mistake, however, to imagine that just because computers could enable the creation of a total surveillance regime, such a regime is inevitable. Technology is rarely deterministic. In the 1970s, democratic countries like Denmark and Canada could have emulated the Romanian dictatorship and deployed an army of secret agents and informers to spy on their citizens in the service of 'maintaining the social order'. They chose not to, and it turned out to be the right choice. Not only were people much happier in Denmark and Canada, but these countries also performed much better by almost every conceivable social and economic yardstick. In the twenty-first century, too, the fact that it is possible to monitor everybody all the time doesn't force anyone to actually do it and doesn't mean it makes social or economic sense.

Democracies can choose to use the new powers of surveillance in a limited way, in order to provide citizens with better health care and security without destroying their privacy and autonomy. New technology doesn't have to be a morality tale in which every golden apple contains the seeds of doom. Sometimes people think of new technology as a binary all-or-nothing choice. If we want better health care, we must sacrifice our privacy. But it doesn't have to work like that. We can and should get better health care and still retain some privacy.

Entire books are dedicated to outlining how democracies can survive and flourish in the digital age.[3] It would be impossible, in a few pages, to do justice to the complexity of the suggested solutions or to comprehensively discuss their merits and drawbacks. It might even be counterproductive. When people are overwhelmed by a deluge of unfamiliar technical details, they might react with despair or apathy. In an introductory survey of computer politics, things should be

kept as simple as possible. While experts should spend lifelong careers discussing the finer details, it is crucial that the rest of us understand the fundamental principles that democracies can and should follow. The key message is that these principles are neither new nor mysterious. They have been known for centuries, even millennia. Citizens should demand that they be applied to the new realities of the computer age.

The first principle is *benevolence*. When a computer network collects information on me, that information should be used to help me rather than manipulate me. This principle has already been successfully enshrined by numerous traditional bureaucratic systems, such as health care. Take, for example, our relationship with our family physician. Over many years she may accumulate a lot of sensitive information on our medical conditions, family life, sexual habits and unhealthy vices. Perhaps we don't want our boss to know that we got pregnant, we don't want our colleagues to know we have cancer, we don't want our spouse to know we are having an affair and we don't want the police to know we take recreational drugs, but we trust our physician with all this information so that she can take good care of our health. If she sells this information to a third party, it is not just unethical; it is illegal.

Much the same is true of the information that our lawyer, our accountant or our therapist accumulates.[4] Having access to our personal life comes with a fiduciary duty to act in our best interests. Why not extend this obvious and ancient principle to computers and algorithms, starting with the powerful algorithms of Google, Baidu and TikTok? At present, we have a serious problem with the business model of these data hoarders. While we pay our physicians and lawyers for their services, we usually don't pay Google and TikTok. They make their money by exploiting our personal information. That's a problematic business model, one that we would hardly tolerate in other contexts. For example, we don't expect to get free shoes from Nike in exchange for giving Nike all our private information and allowing Nike to do what it wants with it. Why should we agree to get

free email services, social connections and entertainment from the tech giants in exchange for giving them control of our most sensitive data?

If the tech giants cannot square their fiduciary duty with their current business model, legislators could require them to switch to a more traditional business model, of getting users to pay for services in money rather than in information. Alternatively, citizens might view some digital services as so fundamental that they should be free for everybody. But we have a historical model for that too: health care and education. Citizens could decide that it is the government's responsibility to provide basic digital services for free and finance them out of our taxes, just as many governments provide free basic health care and education services.

The second principle that would protect democracy against the rise of totalitarian surveillance regimes is *decentralisation*. A democratic society should never allow all its information to be concentrated in one place, no matter whether that hub is the government or a private corporation. It may be extremely helpful to create a national medical database that collects information on citizens in order to provide them with better health-care services, prevent epidemics and develop new medicines. But it would be a very dangerous idea to merge this database with the databases of the police, the banks or the insurance companies. Doing so might make the work of doctors, bankers, insurers and police officers more efficient, but such hyper-efficiency can easily pave the way for totalitarianism. For the survival of democracy, some inefficiency is a feature, not a bug. To protect the privacy and liberty of individuals, it's best if neither the police nor the boss knows everything about us.

Multiple databases and information channels are also essential for maintaining strong self-correcting mechanisms. These mechanisms require several different institutions that balance each other: government, courts, media, academia, private businesses, NGOs. Each of these is fallible and corruptible, and so should be checked by the others. To keep an eye on each other, these institutions must have

independent access to information. If all newspapers get their information from the government, they cannot expose government corruption. If academia relies for research and publication on the database of a single business behemoth, could scholars still criticise the operations of that corporation? A single archive makes censorship easy.

A third democratic principle is *mutuality*. If democracies increase surveillance of individuals, they must simultaneously increase surveillance of governments and corporations too. It's not necessarily bad if tax collectors or welfare agencies gather more information about us. It can help make taxation and welfare systems not just more efficient but fairer as well. What's bad is if all the information flows one way: from the bottom up. The Russian FSB collects enormous amounts of information on Russian citizens, while citizens themselves know close to nothing about the inner workings of the FSB and the Putin regime more generally. Amazon and TikTok know an awful lot about my preferences, purchases and personality, while I know almost nothing about their business model, their tax policies and their political affiliations. How do they make their money? Do they pay all the tax that they should? Do they take orders from any political overlords? Do they perhaps have politicians in their pocket?

Democracy requires balance. Governments and corporations often develop apps and algorithms as tools for top-down surveillance. But algorithms can just as easily become powerful tools for bottom-up transparency and accountability, exposing bribery and tax evasion. If they know more about us, while we simultaneously know more about them, the balance is kept. This isn't a novel idea. Throughout the nineteenth and twentieth centuries, democracies greatly expanded governmental surveillance of citizens so that, for example, the Italian or Japanese government of the 1990s had surveillance abilities that autocratic Roman emperors or Japanese shoguns could only have dreamed of. Italy and Japan nevertheless remained democratic, because they simultaneously increased governmental transparency and accountability. Mutual surveillance is another important

element of sustaining self-correcting mechanisms. If citizens know more about the activities of politicians and CEOs, it is easier to hold them accountable and to correct their mistakes.

A fourth democratic principle is that surveillance systems must always leave room for both *change and rest*. In human history, oppression can take the form of either denying humans the ability to change or denying them the opportunity to rest. For example, the Hindu caste system was based on myths that said the gods divided humans into rigid castes, and any attempt to change one's status was akin to rebelling against the gods and the proper order of the universe. Racism in modern colonies and countries like Brazil and the United States was based on similar myths, ones that said that God or nature divided humans into rigid racial groups. Ignoring race, or trying to mix races together, was allegedly a sin against divine or natural laws that could result in the collapse of the social order and even the destruction of the human species.

At the opposite extreme of the spectrum, modern totalitarian regimes like Stalin's USSR believed that humans are capable of almost limitless change. Through relentless social control even deep-seated biological characteristics such as egotism and familial attachments could be uprooted, and a new socialist human created.

Surveillance by state agents, priests and neighbours was key for imposing on people both rigid caste systems and totalitarian re-education campaigns. New surveillance technology, especially when coupled with a social credit system, might force people either to conform to a novel caste system or to constantly change their actions, thoughts and personality in accordance with the latest instructions from above.

Democratic societies that employ powerful surveillance technology therefore need to beware of the extremes of both over-rigidity and over-pliability. Consider, for example, a national health-care system that deploys algorithms to monitor my health. At one extreme, the system could take an overly rigid approach and ask its algorithm to predict what illnesses I am likely to suffer from. The algorithm

then goes over my genetic data, my medical file, my social media activities, my diet and my daily schedule and concludes that I have a 91 per cent chance of suffering a heart attack at the age of fifty. If this rigid medical algorithm is used by my insurance company, it may prompt the insurer to raise my premium.[5] If it is used by my bankers, it may cause them to refuse me a loan. If it is used by potential spouses, they may decide not to marry me.

But it is a mistake to think that the rigid algorithm has really discovered the truth about me. The human body is not a fixed block of matter but a complex organic system that is constantly growing, decaying and adapting. Our minds too are in constant flux. Thoughts, emotions and sensations pop up, flare for a while and die down. In our brains, new synapses form within hours.[6] Just reading this paragraph, for example, is changing your brain structure a little, encouraging neurons to make new connections or abandon old links. You are already a little different from what you were when you began reading it. Even at the genetic level things are surprisingly flexible. Though an individual's DNA remains the same throughout life, epigenetic and environmental factors can significantly alter how the same genes express themselves.

So an alternative health-care system may instruct its algorithm not to *predict* my illnesses, but rather to help me avoid them. Such a dynamic algorithm could go over the exact same data as the rigid algorithm, but instead of predicting a heart attack at fifty, the algorithm gives me precise dietary recommendations and suggestions for specific regular exercises. By hacking my DNA, the algorithm doesn't discover my preordained destiny, but rather helps me change my future. Insurance companies, banks, and potential spouses should not write me off so easily.[7]

But before we rush to embrace the dynamic algorithm, we should note that it too has a downside. Human life is a balancing act between endeavouring to improve ourselves and accepting who we are. If the goals of the dynamic algorithm are dictated by an ambitious government or by ruthless corporations, the algorithm is likely to

morph into a tyrant, relentlessly demanding that I exercise more, eat less, change my hobbies and alter numerous other habits, or else it would report me to my employer or downgrade my social credit score. History is full of rigid caste systems that denied humans the ability to change, but it is also full of dictators who tried to mould humans like clay. Finding the middle path between these two extremes is a never-ending task. If we indeed give a national healthcare system vast power over us, we must create self-correcting mechanisms that will prevent its algorithms from becoming either too rigid or too demanding.

THE PACE OF DEMOCRACY