CONTENTS

1

Robot Wars

The New Prometheus

Are we witnessing the birth of a new Prometheus? In 2005, computer scientist and futurist Ray Kurzweil declared that 'the singularity was near'.[1] He believed that computer superintelligence would appear within three decades and, therefore, that machines were about to transform civilization in ways we could only begin to imagine. In 2024, he remained convinced of the epochal powers of artificial intelligence (AI). The 'singularity' was, for him, even nearer.[2] He predicted that with the help of AI, 'We are going to extend our minds many millions-fold by 2045'.[3]

There is little doubt that in the last two decades AI has developed startling—near miraculous—powers. Kurzweil is not alone in his belief that AI will transport humans to a new era. James Lovelock, the celebrated originator of the Gaia Theory, believes that the Anthropocene, the age of humans, is over: we are at the dawn of the Novacene, an age in which AI will control and manage humans' lives. AI could soon function a million times more quickly than the human brain. Eventually, according to Lovelock, the Novacene will regulate the 'chemical and physical conditions to keep the Earth habitable for cyborgs'.[4] In his recent bestseller, *The Coming Wave*, Mustafa Suleyman, an AI pioneer and one of the founders of DeepMind, professed a similar view of AI: 'And now we stand at the brink of another such moment as we face the rise of a coming wave of technology that includes both advanced AI and biotechnology. Never before have we witnessed technologies with such transformative potential, promising to reshape our world in ways that are both awe-inspiring and daunting'.[5]

Suleyman is well-placed to know whether we are, indeed, on the edge of an AI revolution. The London-born son of a Syrian taxi driver and an English

nurse, he was abandoned at sixteen. He then gained a place at Oxford University to study philosophy but dropped out before beginning to work on AI, a field in which he became a major figure. Indeed, he was integral to an event which is widely regarded as one of the seminal moments in the advance of AI in the last decade: he helped to develop AlphaGo, the AI program which defeated the world Go champion in 2016. AlphaGo was a remarkable achievement. Even after IBM's Deep Blue defeated Garry Kasparov, the world chess champion, in 1997, experts claimed that 'it may be a hundred years before a computer game beats humans at Go—maybe longer',[6] because Go is a far more complex game than chess. Go is played on a board of 324 squares. Players place black and white stones on any of the interconnections between the squares, with a view to surrounding each other's pieces. The player who surrounds more of the other player's pieces wins. Go, therefore, has a vastly greater number of potential moves than chess does. After three pairs of moves in a game of chess, there are about 121 million possible configurations; after three moves in a game of Go, there are on the order of 200 quadrillion (2×10^{17}) possible configurations.[7] Experts' scepticism was not wholly misplaced.

Yet the experts were wrong. In 2010, Demis Hassabis, Mustafa Suleyman, and Shane Legg set up a small tech company called DeepMind. They were interested in exploring the possibility of developing artificial general intelligence; as Suleyman put it, 'We wanted to build truly general learning agents that could exceed human performance at most cognitive tasks'.[8] It was a hugely ambitious undertaking. In 2013, DeepMind developed an algorithm, Deep Q-Network, which could play Atari computer games. The success of Deep Q-Network attracted the attention of major companies in Silicon Valley, and in 2014 Google bought DeepMind. In 2015, the DeepMind team began to work on AlphaGo, training the program by having it watch 150,000 games of Go.[9] Initially, the computer failed badly in its attempts to play the game. Gradually it learnt on massive datasets, teaching itself a system purely on the basis of trial and error. Because it was a computer program, it could run through games almost infinitely. Eventually, on 15 March 2016, AlphaGo beat world champion Lee Sedol in a five-game series, 4–1. Famously, in move 37 of game 3, AlphaGo made an extraordinary play, positioning a stone on its own on the edge of the board. Amazed observers declared, 'It's not a human move'.[10] Sedol was plainly disturbed by the unexpected move, and AlphaGo went on to win the game. What was most striking was not just that a computer program had learnt to play Go but that it seemed to have developed a creativity that exceeded the imagination of even the best human player.

AlphaGo was designed purely to play a game. The program was thus, in a certain sense, trivial. Yet the evidence for an AI revolution is becoming incontrovertible. AI has already made major contributions to scientific and medical research, fields in which it has affected the lives of perhaps millions and will

soon affect the lives of billions. Since 1972, biologists have been working on the problem of protein structure: 'Determining the crumpled shapes of proteins based on their sequences of constituent amino acids has been a persistent problem for decades in biology. Some of these amino acids are attracted to others, some are repelled by water, and the chains form intricate shapes that are hard to accurately determine'.[11] In 2020, DeepMind announced that AlphaFold, a program created to identify new chemicals, had developed a method of mapping the structure of folded proteins. By the middle of 2021, the program had mapped 98.5 per cent of the proteins in the human body. AlphaFold had solved the problem of protein folding in just eighteen months.[12]

AI programs now routinely read scans for cancer more accurately than doctors do.[13] AI may have even more radical medical uses. One of the pressing needs of modern medicine is to develop new drugs to overcome antibiotic-resistant bacteria. The Broad Institute at MIT and Harvard, led by Dr Felix Wong, has used AI to make major progress in this area.[14] As Jeremy Hsu has described, Wong's team 'tested the effects of more than 39,000 compounds on *Staphylococcus aureus* and three types of human cells from the liver, skeletal muscle and lungs. The results became the training data for AI models to learn about the patterns in each compound's chemical atoms and bonds. That allowed the AIs to predict both the antibacterial activity of such compounds and their potential toxicity to human cells. The trained AI models then analysed 12 million compounds through computer simulations to find 3646 compounds with ideal drug-like properties'.[15] Wong explained the significance of AI to the research: 'Our [AI] models tell us not only which compounds have selective antibiotic activity, but also why'. Because AI can process huge quantities of data, plotting thousands of variables, it is able to identify patterns which are quite undetectable to human researchers. Humans simply cannot hold all those variables in their minds. Consequently, AI programs have detected new molecular qualities, identifying relations between a molecule's structure and its antibiotic capacity that humans had neither perceived nor defined. Algorithms have been employed with increasing success in many other fields, including oceanography, in which they have been trained to distinguish between submarines, mines, and sea-life better than humans can.[16]

AI has also become integral to business and commerce. It has been essential to the competitive advantage of the largest companies.[17] The rise of Amazon is substantially due to data and AI. Amazon's algorithms have automated buying and selling, and, as a result, they have been able to predict consumers' tastes on the basis of the things that other customers, with similar data profiles, have bought.[18] Walmart has also successfully harnessed the predictive power of AI. Walmart's algorithms noted that when a hurricane was announced, consumers in the American South stocked up on Pop-Tarts (likely because Pop-Tarts are easy to prepare and are high in calories). As Linda Dillman, Walmart's chief

information officer, observed, 'Strawberry Pop-Tarts increase in sales, like seven times their normal sales rate, ahead of a hurricane'.[19] Consequently, Walmart shipped more of that product to its stores in the affected states when a hurricane warning was issued.

A similar transformation has been evident in financial markets. Originally, the stock market consisted of human traders physically communicating with each other on the trading floor and exchanging notes documenting the deals they had done in real time, face to face. The process involved famously chaotic scenes in which jacketed traders gesticulated and shouted at one another. In the 1980s, stock markets became predominately computerised; exchanges began to be communicated over email. Electronic trading displaced physical trading. In 2006, entrepreneurs recognised the potential of high-volume electronic trading. High-volume trading involved thousands of small transactions which exploited small price shifts in the market. Traders sold and bought rapidly as prices rose and fell; speed was everything. Since trading had already been digitised—it was electronic—high-volume traders realised that algorithms might be developed to compute and execute trades; software could process financial data far more quickly and accurately than humans could. As a result, from 2006, high-volume traders automated their activities, employing algorithms to sell and buy shares 'at the speed of light'.[20] The stock exchange has been revolutionised by AI; buyers and sellers are connected digitally. Sales are now processed automatically.

AI's progress shows no sign of slowing. On the contrary, it has been exponential and self-reinforcing. On 30 November 2022, OpenAI released ChatGPT, a large generative language model capable of trawling the entire internet for data and producing useful responses to prompts. Many people see generative AI as the next breakthrough. Indeed, some hope that generative AI will help to alleviate poverty in the developing world, by increasing education and economic productivity: 'AI stands to transform lives in the emerging world, too. As it spreads, the technology could raise productivity and shrink gaps in human capital faster than many before it'.[21] In the global south, a lack of teachers, educators, doctors, engineers, and managers is a major obstacle to development; 'AI could ease this shortfall, not by replacing existing workers, but by helping them become more productive'.[22] Locals could draw on AI to help bridge the gaps in expertise. AI might act as a proxy teacher or doctor for locals, accessing the internet for help. For instance, Tonee Ndungu, an entrepreneur in Kenya, has developed two apps to help children learn through engaging with a chatbot.[23] In the West, too, many leaders have advocated AI as a way of transforming the economy and improving the well-being and livelihoods of citizens. For instance, in 2023 the UK prime minister Rishi Sunak, in response to worries about economic stagnation after Brexit, declared that AI represented 'one of the greatest opportunities for the UK': 'Combined with the

computational power of quantum we could be on the precipice of discovering cures for diseases like cancer and dementia or ways to grow crops that could feed the entire world'.[24]

It is difficult to predict how AI will evolve even in the next five years. As a result of advances in AI, humanity is now on the edge of a Fourth Industrial Revolution. The powers of computing, data, and digital communications—all enhanced, enabled, and integrated by AI—are converging to transform every aspect of human existence, just as coal, electricity, and nuclear power successively transformed society in previous eras. A new Prometheus is appearing.

Frankenstein's Monster

Or perhaps, rather than a Prometheus, we are waking a monster. Experts in the field of computing have highlighted the dangers AI may pose. For instance, in an interview with *The New York Times* in May 2023, Geoffrey Hinton, one of the pioneers of neural networks and a seminal figure in the development of machine learning (ML), confessed his fears that AI was approaching a tipping point, when the interconnection of existing systems might trigger the rise of a new class of intelligence. Computers, he said, by sharing their data automatically with each other, could soon become vastly more intelligent than humans: 'Whenever one [model] learns anything, all the others know it. People can't do that. If I learn a whole lot of stuff about quantum mechanics and I want you to know all that stuff about quantum mechanics, it's a long, painful process of getting you to understand it'. Indeed, Hinton described AI as not just troubling but an 'existential threat'. In the near future, he feared, AI would be harnessed neither to play games, nor to make medical advances, but for war. He urged, 'What we want is some way of making sure that even if they're smarter than us, they're going to do things that are beneficial for us', adding, 'but we need to try and do that in a world where there [are] bad actors who want to build robot soldiers that kill people. And it seems very hard to me'.[25] Hinton warned that AI may be used for military purposes. Indeed, AI may automate war, as killer robots, directed by non-human machine intelligence, take over. By February 2024, Hinton's fears had increased. He said: 'If I were advising governments, I would say that there's a 10 per cent chance these things will wipe out humanity in twenty years. I think that would be a reasonable number'.[26]

Hinton might be considered alarmist. Nevertheless, it is striking that many policymakers and scholars who are experts on strategy and security have also worried about the development of AI and its implications for war. They too fear the automation of war. The former US secretary of state Henry Kissinger has been prominent here. Kissinger, controversial though he is, was probably the most important Western diplomat, strategist, and strategic thinker

of the late twentieth century. He was well-placed to make a judgement about the strategic implications of AI for global security. In one of his final acts as a public figure, Kissinger expressed his worries about AI. In 2021, just two years before he died, he published a book with Eric Schmidt and David Huttenlocher about AI and 'our human future'. Kissinger's choice of co-authors was well considered. Schmidt has had a long and illustrious career in Silicon Valley and is profoundly aware of the potential of AI. He served as CEO of Google from 2001 to 2011 and subsequently chaired the Defense Innovation Board and the National Security Commission on Artificial Intelligence (NSCAI) in Washington. Huttenlocher is the dean of computing at MIT, having previously served as the dean of Cornell Tech; both are new centres dedicated to the analysis of digital technology and AI. The book is, then, a statement from a pre-eminent strategist and two tech and AI specialists. It is an important contribution.

In *The Age of AI and our Human Future*, Kissinger, Schmidt, and Huttenlocher avoid the jeremiads of which Hinton is perhaps guilty. They explore the question of how AI will change strategy and the prosecution of war as a political enterprise. In the near future, might AI replace human strategic judgement? Might computers decide when, where, and how to fight wars? The authors' assessment of the political implications of AI for strategic affairs is sobering: 'Only very rarely have we encountered a technology that challenged our prevailing modes of explaining and ordering the world. But AI promises to transform all human experiences'.[27] It is predicted that 'the introduction of non-human logic to military systems and processes will transform strategy'. If the armed forces accrue an advantage from using AI, then they will surely use it—even if only to defend themselves from those states which do use AI. Yet the perils of AI are clear. The authors are particularly concerned with cyberwarfare, lethal autonomous systems, and nuclear weapons. AI operates at a speed which no human can achieve, offering very real benefits to states and their armed forces. The authors warn that if AI gains control of weapons, including nuclear weapons, the norms of deterrence which have operated since 1945 would collapse. The rationale and motivations of AI might be different to those of humans: 'In contrast to the field of nuclear weapons, no widely shared proscription and no clear concept of deterrence (or of degrees of escalation) attend such uses of AI-assisted weapons [. . .] Such reliance will introduce unknown or poorly understood risks'.[28]

Kissinger, Schmidt, and Huttenlocher are therefore disturbed by the prospect of the automation of war, a scenario in which AIs, not humans, develop strategy and prosecute war. In the context of interstate competition, military automation carries many risks. The authors declare that, in order to keep humans in control of lethal weapons, 'We will need to overcome, or at least moderate, the drive to automaticity before catastrophe ensues'. They conclude, 'Defence will have to be automated without conceding the essential elements

of human control'.[29] States need to employ AI, but they must also mitigate and control it.

Kissinger, Schmidt, and Huttenlocher are perhaps the most prominent strategic commentators to address the question of AI. Yet they are certainly not alone among strategic experts in their disquiet that AI is about to automate strategy. Many other scholars believe that AI will inform and even automate strategy. In the last decade, the British military scholar Kenneth Payne has been an eloquent voice in these debates. In a series of books and articles exploring the automation of strategy and war, Payne concurs with Kissinger, Schmidt, and Huttenlocher. In the past, human commanders have been inhibited by a range of emotional commitments. They have been scared to put themselves or their troops at risk; fear has emasculated decision-making. In many cases, imagined fears have been crippling. Pity and morality have also constrained decision-making. Commanders have often sought to preserve human life. On other occasions, hatred and vengeance have compelled extreme actions which have no military logic. For Payne, AI potentially cures these strategic inhibitions: 'AI is primarily a decision-making technology. Its effect is on the nature of warfare, insofar as it alters the long-standing human psychology of the decisions made in combat'.[30] That is, AI will not be hampered by the foibles of human psychology. AI will calculate the strategic situation entirely on the basis of a logical analysis of the data, in order to make rational decisions. It will make strategic decisions quickly and accurately to execute those decisions instantaneously. Consequently, 'AI alters the nature of war by introducing non-human decision-making'.[31] War will become an automated competition between computers, not a visceral struggle between peoples.

It is a radical claim, but other scholars have concurred. For instance, the political scientist and arms control expert Denise Garcia echoes Payne exactly: 'The development of artificial intelligence and its uses for lethal purposes in war will fundamentally change the nature of warfare'.[32] In her recent monograph, she claims that 'militarised artificial intelligence' represents an existential threat. She believes that AI will determine how, when, and where wars are fought: 'What is at stake is the potential loss of human control to machines that will kill autonomously in response to an algorithm, with no humans involved'.[33] For Garcia, the only solution to this future is regulation and human control.[34]

Roberto Gonzalez, a military anthropologist, also decries the military application of AI. Unlike Garcia, he claims that the pursuit of AI is practically misguided; AI is not nearly as effective as military leaders believe. Yet the armed forces are on a 'quest for the automated battlefield'[35] and are actively committed to using AI to automate their operations.

In his work on AI and military decision-making, James Johnson has also warned that military automation, and the substitution of human commanders by AI, is imminent. Like Gonzalez, he rejects the claim that AI might perform

the functions of human commanders. He insists, 'Machines cannot reliably complement or augment, let alone replace, the role of humans in command decision-making'.[36] War is non-linear, unpredictable, and chaotic; AI models are unsuited to the complexity of military decision-making. Nevertheless, Johnson believes that, because of weaknesses in human psychology, AI will begin to assume command responsibilities. In a crisis situation, where there is an overwhelming amount of data, humans are liable to place a false trust in AI. They will become victims of the 'automation bias'; they will defer to AI because they are uncertain what to do themselves. Consequently, 'as deep human–machine symbiosis alters and shapes the psychological mechanisms that make us who we are, thus as they learn and evolve, AI agents will likely become—either inadvertently or more probably by conscious choice—de facto strategic actors in war'.[37] Johnson continues: 'The logical end of this trajectory is an AI commander—planners, warfighters, and tacticians. The danger is that decision-makers may seek to reconcile the paradox of war by outsourcing our consciences in the use of lethal force to non-human agents who are ill-equipped to fill this ethical-moral void'.[38] The literature is troubling. Many scholars believe that artificial intelligence is about to assume the role traditionally taken by human commanders and political leaders; AI will make the decisions. AI will arrogate the fatal decision of whether to go to war. AI will assume the role of politician and commander-in-chief. AI will automate strategy. War itself will be directed not by humans but by machines.

Kissinger, Schmidt, and Huttenlocher, and all these other commentators, then, profess an imminent revolution in strategic affairs. It is possible to identify a second theme in the literature on AI and war. Many scholars are not so exercised by the thought of AI automating strategy, but they are deeply concerned that AI will automate *warfare*. They are disturbed by the prospect that AI will automate weaponry. Above all, scholars in this camp fear that AI will necessarily lead to a proliferation of lethal autonomous weapons; drone swarms and robots controlled by AI will dominate.[39]

The fear that AI will automate weapons has been apparent for about a decade. In 2013, activists concerned about the military threat posed by autonomous weapons created a group called the Campaign to Stop Killer Robots. This group advocates the regulation of the military application of AI. As part of its campaign, at the International Joint Conference on Artificial Intelligence on 28 July 2015, Elon Musk, Stephen Hawking, Demis Hassabis, and many other leading AI experts published an open letter documenting their concerns about the military application of AI; AI could be used to turn weapons against humans. Unlike nuclear weapons, autonomous weapons will be easy and relatively cheap to develop. 'The key question for humanity today is whether to start a global AI arms race or to prevent it from starting. If any major military power pushes ahead with AI weapon development, a global arms

race is virtually inevitable, and the endpoint of this technological trajectory is obvious: autonomous weapons will become the Kalashnikovs of tomorrow'.[40]

Stuart Russell, who played an important role in the development of AI from the 1980s, has been a leading opponent of the proliferation of AI-enabled autonomous weapons. He has campaigned vociferously for the regulation of AI for military purposes and was one of the signatories of the Campaign to Stop Killer Robots' 2015 letter. He is scared by the prospect of AI-automated weapons. To highlight his fear, he created a short fictional film called *Slaughterbots*.[41] The film, released in November 2017, dramatized the assassination of a senator by a swarm of killer drones which then attacked a university campus. The implication was that once they have been automated, such swarms will kill without constraint.[42]

Following *Slaughterbots*, Russell dedicated one of his BBC Reith Lectures in 2020 to the military application of AI. He discussed the issue of automated weapons and killer drone swarms almost exclusively.[43] The lecture reached a climax when Russell described a scenario in which a lethal quadcopter the size of a jar could be armed with an explosive projectile device: 'A regular shipping container could hold a million lethal weapons [. . .] The inevitable endpoint is that autonomous weapons become cheap, selective weapons of mass destruction'.[44] He continued: 'Anti-personnel autonomous weapons could wipe out all the males in a city between 12 and 60 or all the visibly Jewish citizens in Israel. Unlike nuclear weapons, they leave no radioactive crater'. As evidence, he cited the Turkish use of an autonomous Kargu-2 drone in Libya in March 2020. Russell concluded that unless governments acted to regulate the military application of AI, 'there are eight billion people wondering why you cannot give them protection against being hunted down and killed by robots'.[45]

Eric Schmidt has, apart from in his work with Kissinger, articulated similar concerns about lethal autonomous weapons. He takes an entirely different political and ethical stance to Hinton and Russell, believing that the US must invest in AI in order to retain its supremacy—and to protect democracy and freedom. Yet he, too, sees the cataclysmic military potential of AI:

> Eventually, autonomous weaponized drones—not just unmanned aerial vehicles but also ground-based ones—will replace soldiers and manned artillery altogether. Imagine an autonomous submarine that could quickly move supplies into contested waters or an autonomous truck that could find the optimal route to carry small missile launchers across rough terrain. Swarms of drones, networked and coordinated by AI, could overwhelm tank and infantry formations in the field.[46]

Warfare will be automated. Drones and robots, controlled by algorithms, will dominate the battles of the future.

The Campaign to Stop Killer Robots, Stuart Russell, and Eric Schmidt might exaggerate the potential of lethal autonomous weapons. It is striking, then, that many prominent security-studies scholars, while eschewing the language of slaughter-bots, have often concurred with their view. They too claim that AI will automate weapons, making war easier and more likely. For instance, in an important article, Jürgen Altmann and Frank Sauer observe: 'Today's unmanned systems have already increased the risk that military force will be used in scenarios where manned systems would previously have presented decision-makers with bigger, caution-inducing hurdles'.[47] The anthropologist Lucy Suchman claims that states will use AI and autonomous weapons to prosecute dehumanised targets anywhere in the world at will: 'These [AI-enabled targeting systems] become ever more dangerous in the contemporary moment, as the figure of the 'imminent threat' is expanded into a horizon of anywhere and of endless war'.[48] Consequently, these and other scholars call for the regulation of autonomous weapons.[49]

In their recent monograph on AI, Ben Buchanan and Andrew Imbrie describe AI as the 'new fire'. For them, AI represents a potentially revolutionary development for the armed forces, and they draw a striking historical parallel:

> Humanity has also wielded fire's destructive forces. The Byzantine Empire used it to great military success, first during the siege of Constantinople in 672 AD, and then in the centuries that followed. In battle, Byzantine troops shot a specially formulated compound at their enemies, one that would burn even when it came into contact with water. Once the compound hit the target, the power of fire would kick in, torching enemy equipment and causing soldiers to flee. Since then, the flames of war have become deadlier. Could there ever be another force so productive and perilous, one so essentially defined by the exponential growth of its core components? Welcome to the age of artificial intelligence.[50]

For Buchanan and Imbrie, AI is the equivalent of ancient Greek fire or the gunpowder weapons of late medieval Europe. AI-automated weapons will magnify the destructive power of weapons. Buchanan and Imbrie have suggested that with the help of AI, 'missiles would fly to an area of concentrated enemy forces and hover. Each missile would release smaller munitions, and each of these would select and attack an enemy target'.[51]

In the last two decades, David Hambling has established himself as a leading expert on drones and remotely controlled systems. Like Buchanan and Imbrie, he has claimed that military automation is approaching. Autonomous drone swarms, in particular, will be revolutionary:[52] 'A swarm of ten thousand small drones could level a town [. . .] A small perching drone could deliver multiple incendiaries the size of bats [. . .] Acting together drones might bring down a bridge or skyscraper, but they could do more than that'.[53]

Kenneth Payne's work on the strategic implications of AI was already discussed; Payne sees great potential for AI in weapons development too. In his view, AI will facilitate the rise of automated weapons; 'warbots', as he felicitously calls them in his most recent book. In his analysis of warbots, he discusses the now-famous AlphaDogfight experiment at the US Air Force Research Laboratory in 2016 and 2020, which tested AI in a virtual simulation of aerial combat. AI programs, which had used massive amounts of data to teach themselves the best aerial manoeuvres, flew fighter jets in simulated combat against a human pilot. Heron Systems' Falco program proved successful in the trials in 2020. Displaying 'superhuman precision in its flying and fighting', Falco beat the human pilot 5–0. There were several reasons for Falco's victory; one of these was that 'the AI agent could pull manoeuvres that a human pilot simply could not physically withstand'.[54] Another was that Falco calculated that the most effective way to attack an opponent was frontally: 'The AI agent showed a strong favour for what pilots called forward-quarter gunshots, when the two aircraft are racing toward each other head-to-head'.[55] Such an approach is extremely difficult and dangerous; human pilots tended to avoid it. Indeed, one pilot described it as 'a gunshot that is almost impossible'. Many pilots flinch when a plane flies at them. By contrast, Falco, experiencing no emotional response, fired its weapons coolly, no matter how likely the chances of a head-on mid-air collision. These simulated dogfights seemed to demonstrate that AI could automate aerial combat. AI could be quicker, more skilled, and more lethal than even the best human pilots.

On the basis of the AlphaDogfights, many other commentators assume that soon AI will automate combat. In his recent best-seller, Paul Scharre, for instance, fears that military forces are developing autonomous weapons systems which will be able to identify and engage targets independently of human control: 'Militaries around the globe are racing to deploy robots at sea, on the ground, and in the air—more than ninety countries have drones patrolling their skies. These robots are increasingly autonomous and many are armed. They operate under human control for now, but what happens when a Predator drone has as much autonomy as a Google car?'[56] James Baker claims that because AI has the ability 'to outperform humans in pattern recognition and anomaly detection', it will soon direct weapons independently of human control.[57] John Antal confirms the point. He has claimed that the Second Nagorno-Karabakh War was 'the first war won primarily by robotic systems'. The future, for him, is clear: 'When these [autonomous] systems are connected into a network and form a multi-domain strike capability that leverages the synchronization in time, space and effect with artificial intelligence (AI), the ability for anyone or anything to hide in the battlespace will become much harder, if not impossible'.[58] Similarly, Seth Frantzman claims that once drones are AI-enabled, war will start 'to look a lot more like a computer game'.[59]

In the 1990s, John Arquilla, with David Ronfeldt, made an important intervention into discussions about the evolution of war, claiming that 'cyber war was coming'. Arquilla was impressed by the Revolution in Military Affairs in the 1990s, when the US harnessed the potential of new surveillance systems, digital communications, and precision munitions. These new systems would soon allow US forces to coordinate seamlessly with each other, converging on decisive locations on the battlefield. Arquilla has been similarly impressed by AI and the potential of military automation. For him, AI will accentuate the trends of the Revolution in Military Affairs. Robots and drones will replace humans: 'Future battles between advanced forces will be incredibly fast-paced, replete with weapons empowered by artificial intelligence and coordinated to strike in networked "swarms"'.[60] The ethicist and philosopher Ronald Arkin has developed an unusual position in these debates. He, too, claims that lethal autonomous weapons will proliferate to become extremely important. However, he welcomes the development. He claims that because their judgement is motivated not by fear or hatred but by reason, AI will make decisions more ethically than human combatants would. AI will not kill unnecessarily.[61] Nevertheless, he still believes that autonomous drone swarms will dominate the battlefield of the near future.

A consensus is developing across the study of security, armed conflict, and war. In a field which is typically riven with debate and disagreement, it is surprising that so many scholars and commentators have eventually converged on essentially the same position regarding the military application of AI. Despite the wide divergence in their political and critical viewpoints, Henry Kissinger, Ken Payne, David Hambling, Roberto Gonzalez, Denise Garcia, Jürgen Altmann, Frank Sauer, and many others believe that AI is about to automate war—or significant parts of its prosecution. AI is about to displace humans to make strategic decisions as to when and how to go to war. AI will increasingly direct weapons, killing people independently of human control. It is a troubling vision of the future.

AI Scepticism

The concerns of Kissinger, Hinton, Russell, Payne, and others are not baseless. On the contrary, these authors have good reasons to argue the way they do. It is absolutely true that, today, states are actively seeking to harness the power of AI for military advantage. China, for instance, has announced its intention to become the world leader in AI by 2030. Its 'New General AI Plan' proclaimed that 'AI is a strategic technology that will lead the future'.[62] China is determined to have the world's premier AI-enabled military within a decade. Similarly, the Russian president Vladimir Putin declared, 'Whoever becomes the leader in this sphere [artificial intelligence] will become ruler of the world'.[63] Although Putin has suffered a terrible setback in Ukraine, there

is little doubt that he and his successors will attempt to enhance the Russian armed forces with AI as quickly as they can.

In response to the challenge posed by China and Russia, in 2014 the US committed to a 'Third Offset Strategy'. The US has invested heavily in AI, autonomy, and robotics to sustain its advantage in defence and will continue to do so. Some have declared that the US is in an 'AI arms race'.[64] Indeed, Alex Karp, the CEO of Palantir Technologies, a leading tech defence company, went further: 'The power of advanced algorithmic warfare systems is now so great that it equates to having tactical nuclear weapons against an adversary with only conventional ones'.[65] In September 2018, the Defense Advanced Research Projects Agency (DARPA) announced a $2 billion campaign to develop the next wave of AI.[66] The US Department of Defense issued its AI strategy in 2019, accompanied by a major increase in AI funding;[67] in 2024, the Department of Defense budget for AI was $1.8 billion.[68] The US has established the Defense Innovation Unit and the Joint Artificial Intelligence Center to germinate, accelerate, and enhance its armed forces' AI capability. Smaller states are equally committed to the military development of AI. The UK and Israel, for instance, are developing their AI capabilities. AI has become an existential security question which no serious military power can ignore any longer. It is becoming as central to defence policy as aircraft carriers, tanks, or atomic bombs were in the twentieth century.

Today, of all the automated and robotic systems being developed for military usage, the drone swarm has attracted by far the most attention and seems to show the most potential. The trajectory of the drone, or the uncrewed aerial system (UAV), over the last two decades is remarkable. The drone first began to be commonly used by the US in 1999, as a surveillance system; by 2024, it was a ubiquitous weapon, used by almost every combatant on a daily basis. There have already been notable developments in autonomous swarming. In October 2016, the US Department of Defense demonstrated a swarm of 103 Perdix micro-drones capable of 'advanced swarm behaviours such as collective decision-making, adaptive formation flying and self-healing'.[69] The Chinese have also made significant advances in swarm intelligence. In 2017, a formation of a thousand UAVs flew at Guangzhou Airshow, and China Electronic Technology Group flew a swarm of 119 fixed-wing drones.[70] The US Army has procured and tested the TSM-800 drone swarm, manufactured by Booz Allen. At Fort Irwin, California, in recent trials in 2023, operators successfully flew a preprogrammed swarm of ninety-seven TSM-800 drones to attack a designated target; one human controller oversaw the attack remotely (with the capability of aborting the mission), but the swarm was essentially autonomous. The swarm was divided into five subgroups of twenty drones, programmed to attack on different vectors, so that it was more difficult to defend against them.[71] The US Navy has tested superswarms which look and fly like flocks of birds in order to deceive enemy radar. The possibility of automated drone swarms controlled entirely by AI is evident.

No one should doubt the military utility of AI, then. Yet it is easy to be entranced by AI. AI has, after all, made extraordinary advances in a very short time. No matter how tempting it is to enthuse about AI, though, scepticism is in order. Current predictions about AI are more fragile than they appear. Hinton, Russell, Kissinger, Schmidt, Payne, and other experts project a vision of the future based on their understanding of AI today. Yet they take a relatively narrow view of AI, examining only a few exceptional cases; there is little discussion in their work of the difficulties and shortcomings of AI. In addition, they propose the most extreme future scenarios, on the presumption that further major strides are inevitable and obvious. There are serious epistemological dangers to prognostications of this type, especially in a field as empirically complex as war. Many scholars have been too quick to draw causal conclusions about AI and the inevitable automation of war. They predict an AI military revolution on the basis of thin, narrowly selected evidence which supports only the case for automation while ignoring the limitations of AI and the difficulties of applying it to strategy, to war, and to warfare. Indeed, there is a tendency towards circularity in contemporary work. Because these scholars presume the future of AI, they read the evidence about the performance of AI in the present in only one categoric way, which, they claim, leads to that inevitable, already assumed future. It is a pure case of teleology.

Recently, some scholars have begun to question some of the presumptions which have become so established in the debates around security. Rather than advocating a single AI future, they have highlighted the limitations of AI and the difficulty of applying it to military operations. For instance, in an important recent article, American security-studies scholars Avi Goldfarb and Jon Lindsay have punctured the hyperbole around AI, saying that 'AI, from this perspective, is not a simple substitute for human decision-making'.[72] They fully recognise that AI is capable of better, faster, cheaper statistical prediction than humans are. AI has consequently proved highly successful in the commercial world, allowing companies to predict customer demand and market trends with striking accuracy. There is no doubt AI will be useful to the armed forces. Nevertheless, Goldfarb and Lindsay stress the distinctiveness of military operations: 'the conditions that have made AI successful in the commercial world—quality data and clear judgement—may not be present or present to the same degree for all military tasks'.[73] In the commercial sector, markets are relatively stable; demand is predictable. The data on which companies make their decisions is generally clean, reliable, and adequate. Rival companies are serious competitors, but their actions, too, are broadly predictable, operating from within regulatory parameters. Not so in war: 'In contrast with assumptions about rapid robot wars and decisive shifts in military advantage, we expect AI-enabled conflict to be characterized by environmental uncertainty, organizational friction, and political controversy'. The authors conclude,

'War, by contrast, occurs in a more anarchic environment'.[74] In war, data will, therefore, be incomplete, messy, and inaccurate. Moreover, the enemy will actively seek to corrupt and poison data: 'The importance of data and judgment creates incentives for strategic competitors to improve, protect, and interfere with information systems and command institutions'.[75] Moreover, the military decision-making process cannot be reduced to statistical prediction; it is not reducible to an algorithm. A command decision is a complex process. Commanders do not just order a weapon to fire at a threat. They have to define a mission, in which all their forces and all their weapons are organised and oriented to a single goal. Commanders, therefore, must consider many different factors before they make a decision. They have to understand the situation; they must comprehend what they have been directed to do by political leaders. Balancing that direction with a variety of military, civil, and political stakeholders, they must work out what is possible, not only militarily but politically. No matter how impressively it processes data, AI does not possess the judgement that underlies decision-making.[76] 'AI will alleviate some of the data processing burden', Goldfarb and Lindsay allow, but, in war, human intelligence will remain critical. Indeed, AI, data, and machine learning will make 'human beings even more vital'.[77]

Goldfarb and Lindsay are not alone in their scepticism about AI. In a closely related article, Benjamin Jensen, Christopher Whyte, and Scott Cuomo also take a sceptical view of AI. They fully recognise the potential of AI for military affairs, as AI can perform and indeed has already performed a variety of useful military functions. They acknowledge that 'deep learning has the potential to create combat-advising software agents that anticipate both the natural and human environment, offering predictions about enemy actions'.[78] AI could prove very useful in military logistics; it could anticipate supply needs, thereby revolutionizing military readiness. It could simulate defence scenarios to improve reactiveness and decision-making. Alternatively, 'AI advances have the potential to perform a wide range of intelligence tasks faster and with higher accuracy than human analysts'.[79] There are many military practices to which AI might be usefully applied. However, the authors also highlight the operational limitations of AI. War is a complex, bewildering phenomenon: 'As a nonlinear system, every battle and campaign is contingent and subject to emergent properties'.[80] On contemporary battlefields, civilians, friendly and enemy forces are often intermingled and indistinguishable from one another in blasted, ruined urban areas. War is an agonistic enterprise: 'The enemy gets a vote, producing a complexity unique to war. Every change to military capabilities—the hardware—and their battlefield employment through new concepts and organizations—the software—is subject to a corresponding reaction'.[81] The smallest bias or gap in the dataset would generate egregious targeting errors. AI would be extremely susceptible to errors of

targeting in the confusion of a battlefield:[82] 'Consider what would happen if military intelligence professionals entered into a similarly flawed image recognition system hundreds of pictures of adversary fighters assessed to be located in an urban area filled with hundreds of thousands of non-combatants'.[83] The risks are obvious. An AI agent might easily target civilians or friendly fighters; more likely, it might simply stop functioning at all. AI is powerful, but it is also limited. It is very unlikely, according to Jensen, Whyte, and Cuomo, that combat could be completely automated. As a result, the authors dismiss the utopian vision promoted by so many other commentators and scholars. They do not see AI taking over: 'AI does not yet promise to change states' abilities to prevail in major conflict'.[84]

In an indignant recent article, Cameron Hunter and Bleddyn Bowen have made a similar argument and rebutted the claim that AI could ever supersede human commanders. Because AI has been successful under closed conditions, they explain, AI proponents describe war as a similarly prescribed system: 'Decision-making in war under this implied vision is within a closed, rule-based system [. . .] Conceiving of war as a kind of game or closed system allows AI optimists to envisage a future in which AI will be able to make or advise on command decisions'.[85] Hunter and Bowen vehemently disagree with that view; war is an open, complex—indeed, chaotic—environment. Strategy, command, and military decision-making, therefore, require more than mere calculation: 'Command in strategy and tactics requires abductive logic—an ability to think and make decisions based on the constant presence of unknowns and unknowable things that may never appear in a historical dataset or past experience'.[86] Strategists need a subtle awareness of other actors and the range of factors at play as a state moves to war: 'AI currently cannot make judgements, but rather makes probabilistic inferences. Nor can it make useful decisions in the absence of comprehensive data in a closed system'. It is, therefore, difficult to see how second-generation AI could automate military decision-making—much less war more widely. It will be a good deal more difficult for AI to automate war than many scholars presume.

There is much evidence to support the arguments of sceptical scholars like Goldfarb and Lindsay. For instance, there is a common error in much of the literature about the application of AI to military affairs. Many AI advocates extrapolate from military simulations that make use of AI to presume that the same situation would pertain on the battlefield itself. On this account, the evidence from simulations transposes immediately onto the battlefield; what happens in virtual reality will soon inevitably happen in reality.

The heavily referenced AlphaDogfight trials illustrate the problem of this evidential carelessness rather well. AlphaDogfight has been recurrently cited by AI advocates to prove the superiority of AI over human pilots. In the simulations, the AI pilots won. On this basis, it is presumed—by Kenneth Payne and

by Kissinger, Schmidt, and Huttenlocher—that AI will soon fly real planes in combat more successfully than human pilots can. Yet the conditions in those trials were vastly in the favour of AI: 'The AI agent [Falco] was given perfect situational awareness of the simulated environment, including the location of the opposing fighter'.[87] In addition, the human pilot was constrained in a way which Falco was not. In training, human pilots are not permitted to conduct head-on shots; it is too dangerous to practice them in the air.[88] And human pilots do not like taking head-on shots. However, in actual combat, human pilots might well adopt this kind of tactic. Trained on large amounts of pristine data from previous simulations, Falco performed supremely. Yet the real world is vastly more complex than the world of such simulations. Human pilots have to deal with weather—clouds, rain, wind, unusual lighting conditions, and so on—unexpected enemy action, mechanical failures, human errors, air-defence systems, and more. They have to land and take off; they have to coordinate with their colleagues. Their mission changes. It is sometimes difficult to distinguish friend from foe. To an AI pilot, by contrast, even the smallest change to the environment might be confusing.

US Air Force commanders know all this full well and recognise that, in reality, a completely autonomous aeroplane is improbable. They were certainly still interested by the results of the AlphaDogfight trials, but they saw the experiment as a way of improving the performance of human pilots by augmenting them with AI, rather than replacing them.[89] The air force recognised that this trial was only a simulation in a virtual world. For the air force, it is important to distinguish between the virtual and the real. Yet, in many discussions of AI, evidence taken from simulations is assumed to apply immediately in the real world.

A recent furore surrounding the US Air Force demonstrates the fallacy with even more force. In May 2023, Colonel Tucker 'Cinco' Hamilton precipitated a Twitter storm when he seemed to claim that, in a recent exercise, a rogue autonomous drone had attacked its own command post. It was reported that the drone had been unable to find an enemy headquarters and, therefore, logically following its algorithms, attacked a friendly one instead. Many people took this incident as proof that military automation was imminent. They presumed the incident was real. Hamilton later admitted that he had 'misspoken'. The episode had not really happened at all; it had occurred within a simulation. Although the US Air Force is certainly experimenting heavily with AI—with autonomous and quasi-autonomous aircraft—the replacement of piloted combat planes with completely autonomous ones is unlikely. As Bill 'Evil' Gray, a test pilot, observed: 'We are trying to figure out how to integrate artificially trained neural networks, trained in a simulation[, . . .] into the real world'.[90] That is not easy. Prophecies about the imminent AI revolution in military affairs are overstated and under-evidenced.

AI at War

Is AI about to automate war? This question is the central theme of this book. In order to address this issue, I focus on recent and contemporary military practice, rather than projecting into the future. Specifically, I answer two subordinate questions: first, in the last two decades, how has AI been employed in military operations? Second, how have the armed forces reorganised themselves in order to exploit AI? I eventually address a third issue: how has AI changed the character of war in the last decade, and, consequently, how might it change the character of war in the next ten years? The method is deliberately historical; it looks to the past and present. It looks at how militaries have actually applied AI to their activities and operations in the recent past and how they are planning to use AI in the near future. Their plans cannot be taken as a reality in themselves, though they may be organisationally relevant for the present practices. I try not to speculate about how AI might be used or how it might change war and warfare ten or more years from now. In focusing on the past—and therefore on actual evidence—I adopt a sceptical, empiricist approach to AI. I consciously follow the philosophy of the eminent Scottish philosopher David Hume here.

A great deal of contemporary scholarship on AI presumes the future. This is a problem, because there is no evidence about the future. So, any prediction, however plausible it might seem, can be only speculation in the proper philosophical sense. Hume highlighted the dangers of prediction over two hundred years ago from his position of 'determined scepticism'. In a famous passage in his *Treatise of Human Nature*, he showed that cause and effect, so often presumed by philosophers and theologians, can never actually be assumed. He considered the example of billiard balls striking each other. Because billiard balls have collided in the same way in the past, observers naturally presume that they will interact in the same way in the future.[91] Although practically—and empirically—it is correct to assume this eventuality, there is no logical necessity that the balls should strike each other as they have before. Philosophically, there is no necessary bridge forward from the present to the future. In any future case, anything might happen; factors of which we were ignorant might suddenly influence events. Cause and effect are not inevitable or obvious. Humans infer necessary cause from seeing events repeat themselves regularly; they presume a certainty to which they are not entitled. Consequently, Hume famously concluded, 'We have no other notion of cause and effect but that of certain objects which have always conjoin'd together and which in all past instances have been found inseparable'.[92] In the future, even the most apparently ineluctable causal links might not operate. The future development of AI and its application to war is far more indeterminate than billiard balls colliding on a flat baize-covered table.

In this book, I try to avoid prediction and prognostication. Instead, I consciously look backwards to what has actually happened. I examine military

developments in the present and the recent past. I look at how the armed forces have sought to adopt AI, to train with it, and to apply it to military operations. Recent wars are plainly a vital part of the evidence base. Since 2001, conflicts have proliferated and intensified in Ukraine, Georgia, the Middle East, Afghanistan, the Sahel, sub-Saharan Africa, and South-East Asia. War has been a constant, and the amount of potentially relevant material is vast. In particular, the Russo-Ukraine War is of prime significance; it continues to generate nearly endless, often surprising, evidence about war in the twenty-first century, defying many predictions. For instance, as General Mark Milley, the chair of the US Joint Chiefs of Staff, claimed, the Russo-Ukraine War has delivered a unique insight into the potential of AI for war: 'We are witnessing the way wars will be fought, and won, for years'.[93] In a speech to the Royal United Services Institute in November 2022, General Sir Jim Hockenhull, the head of the UK's Strategic Command, discussed the Russo-Ukraine War and AI at length. He used the conflict as a way of illustrating the growing importance of AI, data, and open-source intelligence, declaring, 'The conflict in Ukraine can in some ways be viewed as the first digital war'.[94] The war has involved an explosion of data. The wars in Nagorno-Karabakh and Gaza are also immediately relevant for understanding the military application of AI.

The recent operations of the Israel Defense Forces (IDF) and the current war in Gaza are also instructive. In May 2021, the IDF conducted an eleven-day campaign, called Operation Guardian of the Walls, against Hamas in Gaza. It was described as the 'first AI war', as AI was employed extensively to facilitate targeting. Following the attacks of 7 October 2023 (see chapter 10), Israel has been engaged in a major war with Hamas. The campaign has been brutal, with many thousands of civilians killed and hundreds of thousands more displaced. Nevertheless, the IDF have once again drawn on AI to help them target Hamas militants.

The war in Gaza and the Russo-Ukraine War may be a turning point in the history of war. They may mark the moment when AI first began to be indispensable to military operations. These wars may disappoint the AI proponents, though. There is no sign in Ukraine that AI is about to take over, despite both sides' profligate use of drones and loitering munitions. On the basis of the evidence from Ukraine, AI will not automate war—that is more fantasy or science fiction than reality. Nevertheless, the war in Ukraine has categorically demonstrated that AI has indeed become crucial to military operations. Although President Zelensky, General Zaluznyi, General Syrskyi, and their subordinates still make all the decisions for the Ukrainian military, AI has played a crucial role, harvesting intelligence from a great quantity of diverse data. AI algorithms have helped the Ukrainians to plan and helped them to target the enemy. It has enhanced their military capability. The war shows the salience of AI in contemporary warfare. This connection is likely to deepen in the next decade. AI is

becoming more potent every week, and the armed forces will draw ever more heavily on it. Even if we dismiss the apocalyptical claims about military revolution, AI will inevitably continue to reconfigure warfare. Like gunpowder, railways, telegraphy, automobiles, aeroplanes, wireless, and nuclear weapons, AI will inevitably have a major impact on the way wars are fought now and in the future.

To this end, it is necessary to examine how the armed forces are actually harnessing AI for military operations. Most major military powers are already trying to use AI. A global survey of all these powers would be welcome. Yet a complete survey of how every military force is using AI would be impractical. The empirical focus of this book is, therefore, deliberately circumscribed to achieve a level of evidential adequacy. However, while it is impossible for me to be comprehensive, it is useful to employ a comparative method. In their excellent recent work on technology and civil-military fusion, Yoram Evron and Richard Bitzinger use comparison to great effect.[95] They select four case studies—the US, China, Israel, and India—to show how these states have differentially adopted or failed to adopt new military technologies. The comparisons provide a better understanding of the process in each state, as well as the general pattern of change. I have followed Evron and Bitzinger's method here, adopting a comparative approach focusing on examples from the US, the UK, and Israel. Because the armed forces of these states are Western or Westernised powers, it has been easier for me to gain access to them than it would be to gain access to those of Russia or China. There are also good substantive reasons for focusing on these three powers. The US and Israel are pioneers in the application of AI to military operations. They provide excellent evidence about the military application of AI. And despite the small size of the British armed forces, the UK remains a major European power; as such, it is a leading proponent of the military application of AI. France, Germany, and other European countries are certainly looking to employ AI, but the UK usefully stands as an example of how a medium-sized NATO member is adopting this technology. The evidence presented here is certainly not comprehensive.

In the following chapter, I discuss AI as a technology. However, a major part of the analysis examines not AI as a discrete technology but rather the way in which the armed forces have reorganised themselves in order to be able to employ AI. This is a vital and often under-appreciated issue. AI has not simply automated war or the armed forces, nor will it. In order to exploit AI, the armed forces have already begun to reform their organisational structures and practices. Profound institutional reconfigurations are occurring. The organisational transformations are just as important as the technological developments, because without those alterations in human organisation, it is impossible to use AI. The armed forces are, therefore, changing their command hierarchies and the structure of their headquarters; they are altering their doctrine and practices.

Above all, a profound organisational transformation is taking place. A new partnership between the armed forces and commercial tech companies—such as Google, Amazon Web Services, Microsoft, SpaceX, Palantir, and Anduril—is appearing. In this book, I plot the emergence of this new relationship between the armed forces and tech companies.

The armed forces have, of course, long been dependent on the private sector. In the twentieth century, and especially during the Cold War, private arms companies were contracted to produce weapons and platforms for the armed forces. A military-industrial complex developed. Since the 1990s, private military and security companies have been contracted to perform specific services in support of the armed forces; for the most part, they have provided dining facilities, technical support, and close security. Occasionally, they have provided combat forces—traditional mercenaries. Outsourcing has become a major feature of the defence sector.

The relationship which is crystallizing today between the armed forces and tech companies is different. Tech companies are not providing the armed forces with pristine platforms or weapons. Neither are they supplying peripheral support services. They are providing software, data, and expertise. In addition, they do not merely deliver these AI-enabled services and then leave it to the armed forces to apply them—on the contrary, software and data are immediately related to current operations and need constant revision. Consequently, to harness AI, tech companies are being integrated into the armed forces and into military operations themselves. They are actively partnering with active military forces and deploying their employees forward into operational headquarters. There, the civilian data scientists, programmers, and coders are integrating with military personnel. The pursuit of AI is thus precipitating a major organisational restructuring. A hybrid private sector–public sector, civil-military configuration—a military-tech complex—is emerging. The appearance of this strange new complex is of profound significance not just to warfare but also to civil-military relations. The rise of a military-tech complex raises serious political, legal, and ethical questions which are equally as vexing as current debates about military automation. The problem is not that computers are about to take over strategy and war but that private-sector tech companies are increasingly influencing the conduct of war.

It is already possible to see the emergence of a military-tech complex in Ukraine. In order to harness AI, the Ukrainian armed forces have relied not only on traditional allies, such as the US, but also on close partnership with private-sector tech companies; they have needed the support of Google, Microsoft, Starlink, Palantir, and Anduril. They have fought the Russian invasion with the assistance of tech companies which have provided them with the data, the AI, and the software to be able to execute operations effectively. As General Hockenhull himself noted, 'Much of that digital capability is coming

from commercially available services rather than necessarily traditional military capabilities'.[96] 'Commercial networks' have been a 'force multiplier'. For instance, 'The availability of commercial satellites has enabled an extension of reach in the Ukrainian military's situational awareness and their ability to conduct surveillance and reconnaissance. We're seeing artificial intelligence used alongside commercial software applications to increase the speed of action.'[97]

It is vital that we recognise and try to understand this military-tech complex, especially since, as we have seen, so much of the literature has fetishized AI as a technology, ignoring its organisational aspects. It is also necessary that we acknowledge that any account of the military application of AI and the military-tech complex can be only preliminary. The armed forces are only just beginning to employ AI. The military application of AI is a very new development, one that in most cases has transpired in the last five years. The armed forces and tech companies are at the very beginning of a profound transformation. Studying AI may, therefore, have some equivalence to studying the genesis of strategic bombing forces or tank warfare in the 1920s and 1930s. The potential of AI is clear. The outlines might be visible, but we are examining a volatile process, not a stabilised institution. Analysing the process of construction is always far more difficult than understanding the finished edifice. In the case of AI, it may be even more difficult. The military application of AI is diffused across a transnational organisational complex. The network is still crystallizing. Finally, AI and the military's use of it are developing so rapidly that it is nearly impossible to map the landscape with complete confidence. Even the AI pioneers themselves have been staggered by the pace and scale of the changes—as Hinton's and Russell's warnings illustrate.

Consequently, this book is avowedly provisional. In it, I describe the application of AI to recent military campaigns, especially in the Russo-Ukraine War, and I take examples from US, British, and Israeli armed forces as they try to apply AI to their operations. Only at the end of the book do I offer some tentative predictions about the future trajectory of military transformation and therefore the likely character of warfare between AI-enabled, digitised militaries. Yet, for all my efforts to adopt an empirical method and to limit my claims to what is empirically verifiable, I must allow that even if the picture I depict in this book is broadly accurate for now, it may be overtaken by events. No one knows, for instance, how quantum computing will transform AI and therefore military operations too. However, AI is an existential security issue. Scholars are duty-bound to analyse its development and its implications as best they can. Although this book must be only preliminary, offering conditional findings, it seems imperative for me at least to proffer some interpretation about the military implications of AI. It would be a dereliction of duty to do otherwise.

INDEX