The Strategic Logic of the AI Arms Race

Charles Ferguson
https://cferguson.com/
https://www.project-syndicate.org/onpoint/ai-is-transforming-military-power-and-geopolitical-competition-by-charles-ferguson-2026-06
June 19, 2026

SAN FRANCISCO—It is now clear that the AI revolution portends a profound reordering of the determinants of global military power. For the United States and Europe, much will depend on their ability to reverse the severe decline of their defense and industrial sectors.

While there are some encouraging trends, AI's potential effects on the collective security of the world's democracies, as well as on global arms control, nonproliferation, and counterterrorism efforts, are increasingly worrisome. The fact that AI's rise coincides with a transition to an increasingly disordered world—one likely to differ radically from an era defined by US hegemony and nuclear weapons—further amplifies these risks**.**

The advent of nuclear weapons gave rise to an elaborate, US-led global regime designed to contain the Soviet Union, avoid nuclear war, and manage atomic arsenals. It included the US nuclear umbrella, NATO, and systems governing the design, production, testing, deployment, and use of nuclear weapons. It also encompassed intelligence and early-warning systems, the Nuclear Non-Proliferation Treaty (NPT), multiple US-Soviet arms-control treaties and verification mechanisms, and export controls, most notably the Missile Technology Control Regime. Underpinning all this was a vast infrastructure for strategic policy analysis.

Though imperfect, this regime proved remarkably successful: nuclear proliferation remained limited, no direct wars broke out between nuclear powers, and nuclear weapons have not been detonated (aside from testing) since 1945. That order, however, has been gradually crumbling as the West declines and China rises, leading to an increasingly anarchic distribution of military power across many countries and even private companies. The consequences are already visible in the West's lagging AI-enabled warfighting capabilities, weak export controls, and inadequate AI safety regulations.

Over the past several decades, both America's and Europe's defense establishments have become rigid, slow-moving, and woefully out of touch with reality. The US military and its primary contractors remain heavily dependent on enormously expensive—and extremely vulnerable—human-crewed systems: aircraft, ships, tanks, aircraft carriers, submarines. Even individual missiles cost millions of dollars, and only hundreds or thousands of each type are produced annually.

With high-volume commercial markets increasingly driving both economic and military performance, the US has fallen far behind China and other countries in critical industries such as additive manufacturing, semiconductors, robots, drones, personal computers, mobile

phones, and batteries. As a result, US weapons are often orders of magnitude more expensive than the systems they are designed to attack or defend against. These problems are compounded by the instability and dysfunction of American politics, which increasingly affect defense procurement and US policy toward Iran, China, Russia, and Ukraine.

Similar dynamics are evident in Europe, particularly in internal disputes over Ukraine and defense-industry cooperation. At the same time, the tech sector's growing political power has weakened Western security efforts, most notably by undermining export controls on China**,** eroding the Russia sanctions regime, and impeding major governmental AI research and development initiatives.

Absent dramatic changes in US and European behavior, the global balance of power could shift rapidly in favor of China and other innovative countries as the integration of AI, drones, and robots gives rise to new forms of military power. The wars in Iran and Ukraine offer an early glimpse of this emerging reality.

Missing the Drone Revolution

In 2015, US President Barack Obama's administration negotiated the imperfect but meaningful Joint Comprehensive Plan of Action, which sharply constrained Iran's nuclear program, and secured a subsequent United Nations Security Council resolution limiting Iran's ballistic-missile activities. But the first Trump administration abandoned the JCPOA, restoring sanctions and freeing Iran to resume both nuclear and conventional-weapons development, and Joe Biden's administration exacerbated the problem, allowing Iran to reach the brink of nuclear-weapons capability and develop advanced drones. The Biden administration also failed to modernize the US military in response to the rise of drones and AI, even after the Russia-Ukraine war and the rapid pace of commercial AI development underscored the urgency of doing so.

When the second Trump administration attacked Iran, it therefore relied on a military establishment built around increasingly obsolete weapons and hampered by inadequate coordination with Ukraine, a potentially vital ally. As a result, the US lacked critical capabilities, including large-scale AI-enabled drone and anti-drone forces, and ultimately lost a war it should have won.

Some observers have argued that the Iran war shows that the new era of drone warfare inherently favors smaller, less technologically advanced countries over major powers. This is flatly wrong—and a dangerous illusion. Just as the fragmented internet industry was replaced by a few giant firms, so, too, early drone warfare is rapidly being transformed by enormous drone forces controlled by military AI models relying on massive data centers. All else being equal, AI strongly favors major powers over smaller rivals. Just as importantly, however, the AI era also rewards flexible innovators and penalizes obsolete legacy militaries.

What the Iran fiasco actually demonstrates, therefore, is that the US military is increasingly crippled by its own obsolescence. In this regard, the US bears a striking resemblance to Russia,

whose inability to defeat Ukraine similarly reflects deep institutional dysfunction in the face of the AI-drone revolution.

The shift was already apparent by 2024, when it had become clear in Ukraine that drones would dominate future wars and that Western dependence on a global drone industry dominated by China and other potential adversaries posed major strategic risks. The war also showed that effective drone warfare requires both advanced AI systems and real-time feedback between combat units and highly responsive AI and drone industries.

Faced with this profound disruption, the US and Europe could (and should) have launched an urgent effort to build a Western drone industry closely integrated with AI capabilities and frontline combat units. In preparation for potential conflicts with Iran and China, American airbases and carrier battle groups could have been transformed into state-of-the-art platforms capable of deploying tens of thousands of drones, supported by dedicated data centers, real-time R&D, and surge-production capacity.

To this end, drone factories in Taiwan, Japan, and South Korea could have provided critical capacity to hedge against the risk of a Chinese attack on Taiwan. Similar measures could have strengthened NATO and US forces in the Middle East, while military AI programs could have trained on battlefield data from Ukraine, accelerating the development of sophisticated offensive and defensive drone capabilities.

Had such steps been taken, the war with Iran would likely have unfolded very differently. The US could have deployed tens of thousands of coordinated AI-enabled drones per day, blanketing Iran's coastline, weapons depots, drone factories, launch sites, and nuclear facilities. This, in turn, would have enabled the US to locate and destroy many more Iranian military assets, including speedboats, drones, and mobile missile launchers.

Most Iranian attacks, whether conducted by drones or humans, would have been detected and engaged in real time, facing the same near-suicidal odds as Russian troops attacking in Ukraine. Meanwhile, massive "drone walls"—each comprising thousands of coordinated drones— could have protected US forces, Gulf infrastructure, and commercial shipping. Under such conditions, the US could have neutralized or destroyed most of Iran's military assets, safeguarded the Gulf's oil and gas facilities, and possibly kept the Strait of Hormuz open throughout the war.

Lessons from Ukraine

This scenario might seem fanciful, but the Ukraine war has shown that it is entirely feasible. Ukraine, with a population of 40 million, already produces and uses 10,000-20,000 drones per day, and is rapidly expanding production. While Russia now mass-produces drones as well, it remains hostage to legacy institutions riddled with incompetence and corruption. Ukraine, by contrast, has a remarkably vibrant startup ecosystem, world-class engineers, and the world's most sophisticated drone forces. Product cycles are now measured in weeks, as drone and

software producers—directly linked to frontline combat units—continuously adapt to advances in AI and changing battlefield conditions.

During the war's early, primarily conventional phase in 2022-23, most casualties were caused by artillery and bombs, fewer than one-third of casualties were fatal, and Russia's casualties exceeded Ukraine's casualties by two or three to one. Today, drones account for 70-80% of all Russian casualties, more than half of which are fatal. I have been told privately by an investor in Ukraine's defense industry that Russian casualties now vastly outnumber Ukraine's.

The scale of Ukraine's drone warfare is staggering. In 2025, Ukraine conducted more than 800,000 successful drone strikes—one-third of which were directed against Russian personnel and the remainder against weapons, vehicles, and buildings. The pace of Ukraine's drone attacks has continued to grow sharply; in May 2026 alone, Ukraine hit 180,000 Russian military targets, an increase of 12% over the previous month. Russia is now losing 35,000 soldiers per month and recruits only about 29,000. Since re-enlistment is now a hard sell, to put it mildly, Russia's military is shrinking.

With the exception of Russia's ballistic-missile attacks on cities, Ukraine is defending itself effectively, using drone walls that coordinate thousands of increasingly AI-controlled reconnaissance, detection, and kamikaze drones, as well as electronic warfare. The front line has become a 20-mile-wide drone-dominated "grey zone," where human life expectancy is often measured in minutes.

As medium-range Ukrainian drones cripple Russian military logistics, longer-range drones—capable of striking targets up to 1,000 miles away—are steadily degrading Russia's oil industry, military, defense sector, and other strategic assets, forcing rationing and export bans on gasoline and jet fuel. Ukrainian targeting is heavily reliant on AI: targets are identified almost in real time using satellite imagery and video footage from drones, hacked CCTV cameras, and mobile phones. As a result, despite its far larger population, economy, and military, Russia can no longer advance and may now be losing territory.

Yet Ukraine is prevailing only because of Russia's incompetence and inferior AI technology in the face of Ukrainian ingenuity. The same is true of Iran's stalemate/strategic victory against the US. Had the US modernized its military, it could have leveraged the vastly greater sophistication and scale available to it as the world's largest economy and most advanced AI ecosystem. But doing so would have required a state-of-the-art military capable of producing and deploying tens or even hundreds of thousands of AI-enabled drones per day—potentially millions over the course of the war—backed by real-time, AI-enabled targeting.

Neither the US nor Europe can currently achieve anything remotely close to this. In December 2025, only three months before the Iran war began, the US Defense Department published a "drone dominance" plan aimed at increasing military drone production to 300,000 units by 2027—roughly one-twentieth of Ukrainian production.

The gap becomes even more striking when compared with China. Chinese manufacturers are expected to produce 8-15 million drones this year (Ukraine expects to produce ten million this year, according to a recent statement by President Volodymyr Zelensky). China is also developing military drone swarms and leads the world in swarm-control software. Multiple Chinese firms compete by producing enormous displays, some recently exceeding 30,000 coordinated drones.

The US, by contrast, remains woefully unprepared. Its very first carrier-borne drones were cleared for low-rate initial production only in May. The US military's drone-control systems are equally backward, incapable of managing large drone walls, swarms, or phased attacks. Israel, once a leader in military technology, has fallen behind as well.

But the most important lesson from Iran and Ukraine lies beyond drones themselves. What they reveal is a deeper shift: AI is becoming central to warfare and military power.

Automated Hard Power

Drones and AI were made for each other, and their capabilities are expanding rapidly as they co-evolve. Before long, cars, trucks, ships, planes, submarines, torpedoes, humanoid robots, and tanks will all become AI-controlled drones. Even mines, missiles, radars, rifles, bombs, and artillery shells will have AI built into them; some already do.

This transformation reflects four factors. The first is complexity: no human being, organization, or traditional computer system can manage the deployment of hundreds of thousands of weapons per day, particularly amid the fog of war. The second is effectiveness, as AI systems have already surpassed both humans and traditional software in navigating terrain, identifying targets, and evading defenses.

A third factor is response time. Missile, laser, and drone warfare require reaction times far beyond human capacities. An attack by 10,000 AI-enabled drones or hypersonic missiles can be countered only by AI-enabled defenses.

Lastly, falling semiconductor costs and rising processing power are making it possible to embed powerful AI capabilities in every drone at low cost—sharply improving drone effectiveness while reducing dependence on communications vulnerable to jamming.

Underlying all of these trends is a simple fact: weapons are now far more effective when they do not have to include human beings. Human-crewed systems must be large, relatively slow, safe to operate, and capable of returning to base. Remove the human, and weapons can be smaller, faster, cheaper, and disposable—and you can make a lot more of them.

Earlier generations of drones still required a human controller, who could be far away but was nevertheless essential, making large, highly choreographed drone formations difficult to

deploy. AI is progressively eliminating that constraint. As Russia has discovered, neither armored weapons nor human soldiers can compete against swarms of AI-enabled drones.

President Vladimir Putin may be willing to sacrifice hundreds of thousands of Russian soldiers every year, but many other governments will not. Consequently, AI-enabled systems will increasingly replace human beings on the battlefield. As AI continues to advance, so will AI-enabled weaponry, including systems capable of planning and managing battles and, eventually, even entire wars.

AI Warfare's Winners and Losers

Leaving aside nuclear weapons, military power in the AI-drone era will depend on four determinants: drone-manufacturing capacity, AI capabilities and infrastructure, access to training data, and—perhaps most important—the political and organizational ability to deploy these assets effectively, even in the face of resistance from entrenched industry, military, and government interests. Superpower status in the AI-first era will require massive investments, comparable to those devoted to the US and Soviet nuclear-weapons infrastructures during the Cold War.

Given these conditions, China appears to be the most obvious beneficiary of the AI-drone revolution. It has enormous resources and fewer legacy constraints, dominates global drone manufacturing, and its AI industry is approaching parity with that of the US.

While much of that progress reflects the long-standing failure of US export controls—and their subversion by American and European firms—China now has a strong, self-sustaining AI and technology ecosystem. Yet it remains severely deficient in military AI training data, largely because China has not fought a war since 1979. Although it probably acquires data from Russia, Iran, and cyber operations, those sources are vastly inferior to the data possessed by Ukraine, the US, and even Israel.

America's future position relative to China remains an open question, given its lack of a competitive drone industry, its legacy defense establishment, and its self-interested technology sector. But the US also retains significant advantages: a world-leading AI ecosystem, a strong startup culture, extensive warfighting experience, and potentially valuable allies—especially Ukraine—with access to the world's best military AI training data. Whether the country will be able to overcome its internal dysfunctions, however, remains to be seen.

That leaves two clear losers in the AI-drone era: Western Europe and Russia. Both suffer from deep structural weaknesses. In Europe, these include an obsolete defense establishment, limited drone-production capacity, and a deteriorating position in high-volume, high-tech manufacturing. Apart from the United Kingdom, it also has little recent warfighting experience and occupies a weak—and perhaps worsening—position in AI. Those weaknesses could leave Europe vulnerable to Russia, which has far greater military experience.

That said, Ukraine may prove to be Europe's biggest asset if the European Union and NATO embrace it. Ukraine is rapidly emerging as a global leader, not only militarily but also technologically, thanks to an extraordinarily dynamic startup ecosystem that is increasingly compared to Israel's.

Russia, by contrast, finds itself in the worst position of all major powers, despite its military might. Within months of Putin's invasion of Ukraine, roughly one million people fled the country, including many of its most educated, productive, and technologically sophisticated citizens. With its technology ecosystem in ruins, Russia now lags badly across every layer of the AI technology stack, from semiconductors and data centers to advanced AI models. Putin's war may leave it with natural resources and nuclear weapons—and virtually nothing else.

AI Arms Control in a Disordered World

When it comes to arms control and nonproliferation, the outlook is not encouraging. Unlike nuclear weapons and the facilities required to produce them, the key components of AI weaponry—additive manufacturing, drone designs, explosives, GPUs, and open-source AI models—are already widely distributed, relatively easy to conceal, and unlikely ever to be fully controlled, even if the world's major powers agreed to try. Just one AI marketplace company, Hugging Face, hosts thousands of open-source models and datasets from a wide variety of countries and companies. While superpowers hold decisive advantages, many second-tier powers are therefore poised to acquire significant AI-drone capabilities.

The US debacle in Iran notwithstanding, the two likely AI superpowers—the US and China—will probably retain substantial capacity to enforce some degree of global peacekeeping, provided they are willing to cooperate and assume that role. Otherwise, smaller, AI-driven regional conflicts may proliferate. Unfortunately, AI-enabled crime is already a reality, and AI-enabled terrorism will probably follow, especially given dangerously weak AI safety regulation and industry guardrails.

The prospects for strategic arms control are similarly bleak. Some have argued that the US should trade export controls for an AI safety treaty with China, while others have proposed an NPT-like treaty banning lethal AI drones. Earlier this month, Anthropic went further, expressing concern that its models might be approaching unmanageable levels of self-improvement and advocating a pause in AI development.

Despite the unquestionable virtue of their goals, these proposals are both unwise and unlikely to succeed. The fundamental problem is that they require universal compliance, yet reliable verification is nearly impossible, as software and training data can be concealed, falsified, or manipulated with relative ease. At the same time, all major powers care deeply about the security of their own AI systems while devoting enormous resources to penetrating, disabling, or misleading those of their rivals. Inspections would pose serious security risks and create powerful incentives and opportunities for espionage and the development of offensive cyberweapons.

Looming over all this is the prospect of artificial general intelligence (AGI)—or even superintelligence—and the race to develop it first. Neither the US nor China is likely to accept any limitations in that domain, particularly when Russia, India, North Korea, Israel, and multiple private companies are pursuing similar ambitions.

It is therefore difficult to imagine major powers curtailing AI development and allowing one another to inspect their most advanced AI systems and weapons. And even if they did, there is little reason to believe they would be satisfied with the completeness or accuracy of the resulting disclosures.

To be sure, China and the US may still cooperate in areas of mutual interest—for example, ensuring that AI models don't accidentally trigger a war or facilitate terrorism. Those are certainly worthy goals, but broad strategic cooperation on AI arms control between two rival superpowers is highly unlikely. A truly global regime is even less plausible.

Where, then, does that leave us? As a quote often attributed to the science-fiction writer William Gibson puts it, "the future is already here; it's just not evenly distributed." At the moment, our distributed future also looks far more complex and dangerous than many anticipated.

Charles Ferguson is a technology investor, policy analyst, and director of documentary films, including the Oscar-winning *Inside Job*. He is a limited partner in six AI venture capital funds, a nonexclusive partner in Davidovs Venture Collective, and a Venture Partner with EWOR. His direct investment positions include three technology incumbents (Apple, Microsoft, and Nvidia) and many AI startups, including Perplexity, Etched, CopilotKit, Paradigm, and Browser Use. The opinions expressed here are his own alone.