

2 CBRNE

*Dedicated to Global
First Responders*

DIARY

July 2026



**First Plague
Vaccine (?)**

**New vaccine for
Burkholderia
pseudomallei
(melioidosis)**

**KSA goes
Nuclear**

**Humanoid
surgeons**

Al al Qaeda

Sudan CWAs

Agent Orange

An International CBRNE Institute publication

IOI
International
CBRNE
INSTITUTE



DIRTY R-NEWS



What if the smartphone in your pocket was also a radiation detector?

By Andrew Longman and Ephraim Fischbach

Source: <https://thebulletin.org/2026/06/what-if-the-smartphone-in-your-pocket-was-also-a-radiation-detector/>



Smart phone image Image courtesy David / Pixabay

June 14 – In the Middle East, intelligence services furiously hunt for fissile material in Iran. In Japan, residents still worry about radiation exposure from Fukushima Daiichi. In other places, [stolen or missing radioactive sources](#) have made the news. One solution: the cell phone.

The average, everyday smartphone could be equipped with radiation detectors and an app that would make it into a small, highly mobile, radioactive particle mapper.

A network of such phones would mean that gamma-ray spectroscopy would be truly ubiquitous—in other words, this phone technology would be so widespread that it would essentially be encountered everywhere, all the time—and consequently allow for the successful detection and mapping of radiological sources. It would make for a first line of defense, a sort of 21st-century update of the Distant Early Warning system used in the Cold War to spot an incoming attack. The idea of putting radiation detectors—or to be more precise, gamma ray spectrometers—into cellphones isn't new: It began more than 20 years ago, then we reprised it

after the Fukushima disaster, and several times since. But as the technology has improved dramatically—and the world situation has gotten worse—in the time since then, it is worth revisiting the idea. We wish to restate that in our opinion, the only physics solution to track and interdict radiological sources—be they from [theft](#), rogue states, terrorists, or natural disasters—is gamma-ray spectroscopy and real-time data-grams from all cellular telephones.

What's a gamma-ray spectrometer, you ask? When high-energy light waves (gammas) pass through the device, it captures their frequencies (spectra). A prism is a kind of spectrometer for visible light; the color separation is frequency separation. And each radioactive element—plutonium, for example—has its own collection of signature frequencies. While gamma-spectroscopy used to cost big bucks, today a decent quality device could be stuffed into your smartphone.

While readers may be more familiar with the term “Geiger counter” from what they have seen in the movies, a Geiger counter



(or more precisely, a Geiger-Mueller tube) just counts photons. It is old tech.

What's the difference between a gamma-ray spectrometer and a Geiger counter? A Geiger-Mueller counter is sort of like a bucket that provides one count of all the baseballs it catches. In contrast, a gamma-ray spectrometer measures the speed, and counts up fastballs, slowballs, curveballs, and sometimes measures their spin, of all the baseballs it catches. And a gamma ray spectrometer measures each photon by its frequency, which means that you can discern radioactive material composition at a distance.

In addition, by having a vast sea of devices equipped with a gamma-ray spectrometer, you may detect in two dimensions—not just stumble upon hot spots. Tens of millions of gamma ray spectrometers mean that you are getting into making a map of radioactive materials present, not just having a threshold box near a bomb that goes “beep.” The capability of this proposed system would not be met by having a few Geiger counters; it would be met by having hundreds of millions of gamma ray spectrometers moving around randomly.

In 2012, my co-authors and I wrote an [article for the Bulletin](#) on what would be needed to make everyday cell phones into small, highly mobile, radioactive particle detectors. The article was part of a surge of research interest at that time into using networked detectors to locate hidden sources of radioactivity and dealing with the fallout after a radiological event.

When first proposed in the late 1990s, the concept was in the toddler days of smartphone technology. In 2003 [Los Alamos National Laboratory created “RadNet,”](#) a fledgling cellular-telephone-based radiation detector network. Los Alamos and NIST teamed up later in 2006 for field tests at Los Alamos. But while technology components have been researched, widespread social adoption—which is the only way to really work—hasn't gone viral.

Our own crude first prototype attempt began at about that same time, running on a Palm Treo 650 with a LabView application communicating by RS-232 with a Geiger-Müller circuit. Today, billions of people carry in each of their pockets what is in effect a cluster supercomputer with high-speed internet radio—and you can purchase Bluetooth connectable scintillator spectroscopy.

In 2006, we started our quest to persuade everyone of the usefulness of such a system. We went to the cell phone manufacturers. We went to the Pentagon. We went to venture capitalists. We computed. We went to Radio Shack—back when there still was one. We got some small funding from the Indiana Department of Transportation to do a prototype. And while interest was spawned, we have not been successful so far in persuading the world of what we still think is urgent, necessary, and sufficient: The only physical way to protect the world from nuclear terrorism requires real-time gamma-ray spectroscopy in everyone's cell phone.

If this were standard infrastructure now on every Samsung and Apple smartphone, then the entire world would almost always know where most radioactive material was. Since background radiation would be repeatedly measured, and sources constantly measured, any movement would show up as a statistical excursion from the established baseline. This would not be unique to Iran, but would become a property of the planet. And it would have applied a few years ago when a hot radioactive source was stolen in Japan, for example.

It is certainly technically available to essentially make every cell phone into a gamma ray detector. The science, engineering, and technology abundantly exist. But governments and the public and the physics community and the cell phone manufacturers must all say “Yes” in order to make it a standard and ubiquitous feature.

And making them “standard” and “ubiquitous” are essential; a few thousand of these phones is useless. A few million is a poster session. A few hundred million is useful but inadequate. We have simple needs; just give us your 8.3 billion phones.

A few questions inevitably come up; here are the answers: What about [civil liberties](#)? The data would be anonymized, and frankly, if there were radon in my basement, I'd like to know. Couldn't you do it with satellites? No, they are too far away, the signal too weak. Can't you do it with trucks or drones? No, the sampling is too sparse. Aren't they already doing it? No, it's not in there. Can't you do it with cops and firemen using a special device? No. Physics statistics is an unforgiving thing—the hope of effectiveness begins in the tens of millions of measurement nodes. The signals from contraband radioactive materials are dim. Unless you have a good probability of getting multiple measurements within a few meters, you don't have a signal-to-noise hope.

Opposition to the proposed network detector concept has raised good points—that many classes or types of hidden bomb would be missed even if every cell phone in the world had the best gamma-ray spectroscopy we could cram into it. But counting gamma-rays is a statistical pursuit; billions of measurements of the background, in every corner of the world, make all subsequent measurements there better; measuring nothing is important; something is just differential nothing. Ubiquitous spectroscopy gives a better signal-to-noise ratio, and it puts pressure on the terrorist—any screwup, and he could be seen. Today, the terrorist has a free hand, and there is pressure on the good guys—there are very few chances to succeed with today's infrastructure.

Ubiquitous detection makes being a terrorist hard. He can try to resort to heavy lead shielding. But this can become intractably massive and consequently hard to move. A cloud of millions of detectors, counting for thousands of hours, moving unpredictably, can find leaks.

Terrorists make mistakes, and ubiquitous cell phone-enabled gamma detectors



would be in a position to capitalize on those leaks.

Detection is deterrence. A terrorist who thinks he may get caught is a terrorist on the losing end of a cost/benefit analysis. Cell phone density is directly proportional to target desirability for the terrorist, consequently making deterrence in proportion to vulnerability. The proposed system provides insurance for a few dollars per cell phone. The cost/benefit analysis for implementing ubiquitous gamma spectroscopy by cell phone hugely favors its use.

Thermodynamics has given us the concept of emergent behavior. There are properties that billions of particles possess that thousands do not. As the saying goes, "Quantity has a quality all its own." Nowhere is this more true than in cell-phone radiological detection. The ubiquity and density of many millions of phones would give rise to two-dimensional image tracking, instead of simple zero-dimensional thresholding. And once the real-time geographic background has been recorded, innumerable machine learning and astronomy algorithms could compete for detection supremacy in ways we could not have anticipated two decades ago; this democratizes the stopping of nuclear terrorism and converts the problem into a machine learning competition.

Convincing a society is a slow percolation problem. Years ago, Japan's MEXT was interested, and Chiyoda Technol Corp was curious. Eventually, SoftBank released a phone with a small dosimeter in it and an app for putting push-pins on map locations. While we researchers were grateful for these increments, we failed to persuade our colleagues on central points: large detection aperture, spectroscopy, ubiquity, real-time reporting, automatic 2D geo-mapping, without human intervention, with server-side algorithmic source detection, are all *required* components. But so far, the systems that have been built have had only unconnected pieces of that puzzle.

Systems of thousands of nodes have been done as demonstrators. The Defense Advanced Research Projects Agency, or DARPA (the folks who funded the internet, and did much of the early research into satellite navigation, among other things) showed the effectiveness of this approach at first, with 100 detectors, and then 1,000 detectors around the year 2016. By 2018, there was a roll-out to state, local, and federal authorities, but these likely remain in the thousands of nodes. We applaud every baby step, but we wish to be clear that life begins in the tens of millions. Tens of thousands won't cut it. We will keep insisting after every future radiological crisis that ubiquitous, real-time, internet reporting, gamma-ray spectroscopy, has to be on every smartphone if you want to save the world from nuclear terrorism. As physicists and nuclear engineers, we look out today over the landscape of the world, and we see the Trump administration going full tilt on missile defense. Fine enough; inbound nuclear missiles are bad, and stopping them is good. But a missile gives away that it's coming, and it tells everyone who launched it. By contrast, a pickup truck delivered by a drug cartel, or a package shipped commercially, is less detectable and doesn't

have a return address. And today we have a new delivery method unrealized in the early 2000s—a robot trotting from the wilderness into a city. Or a fixed-wing drone, flying at 10 meters above the ground. So, missile defense, while necessary, addresses a 20th-century problem, while robots and drug mules are 21st-century problems: delivery systems that are undetectable under the radar. But all of those are (at least potentially) detectable by ubiquitous cell phone gamma-ray spectroscopy. The 21st-century problem is urgent. According to reports, Iran had in its possession enough fissile material to construct several nuclear weapons. One of the stated primary aims of the Trump and Netanyahu administrations was to destroy the Iranian capability to build and deploy nuclear weapons. And the problem is bigger than just Iran, the IAEA reports that thousands of radioactive sources have been [stolen or went missing](#)—as well as radioactive [accidents](#) in recent years in France, Russia, and the United States. The incident at Fukushima also shows the need for such a detection and measurement system, as does the unresolved question of a plant leak amid the Ukraine war. At its heart, these are all intrinsically physics problems. Once the network of detection is designed to conform to the physics constraints, then military planners, politicians, private enterprise, and policymakers will have new opportunities to protect people from this existential threat. And the usefulness of such a smartphone-based detection system would not be limited to just dealing with existential risk. A fun corollary comes up, regarding how it could be used to advance physics. If every smartphone were a gamma-ray spectrometer, layers of different physics inputs would have to be separated. There would be the terrestrial background from bananas, granite steps, and antique uranium cookware. There would be the signal of interest, the illicit radioactive material. There would also be the astronomical signal, the cosmic ray, and the solar wind. Separating these map layers would be an essential function in making the network detector perform its function. As such, a whole-world X-ray observatory data set with several billion pixels and an aperture on the order of the planet would result; the signal you didn't want would be ... basic science. In other words, building a system like this would also hand scientists the world's largest possible X-ray telescope. The threat posed by smuggled nuclear materials or clandestine weapons programs isn't going to fix itself. Rogue actors will continue to procure fissile material and attempt to construct civilization-ending bombs for motives that run opposite the love-thy-neighbor religion of most people. The probability that the entire terrorist enterprise will be bombed or lectured out of existence is low. Civil society has only one physics option for counting the statistical particle events that provide partial location and signature information of illicit radioactive sources. Stolen radioactive sources, leaking nuclear plants, terrorist threats, plant-war-threats, Fukishimas—these are huge problems that aren't being



treated as huge, in part because no one believes there is a credible solution. We hope that if societies can grasp that there is a physics plus AI solution, then these problems would no longer be treated as “too-big-to-be-solved.”

As we did at Fukushima Daiichi, and now at the Iran crisis, and as we will do in the future at every new radiological event, we call for gamma-ray spectroscopy, with real-time data reporting, in every smartphone.

Andrew Longman holds a PhD in physics from Purdue and is a consulting instrument scientist. He has provided hardware or software instruments to Honeywell, Blue Wave AI Labs, Tri Alpha Energy, Exelon, Parker, and accelerator data acquisition and controls to Purdue. He has taught physics at Purdue and Ball State. A dedicated Christian, he explores the moral interface of the nuclear scientist with physical instrumentation.

Ephraim Fischbach is an Emeritus Professor of physics at Purdue University. His career spans theoretical and experimental work in fundamental interactions, neutrino physics, and unconventional tests of gravity and nuclear decay. He was a researcher at Niels Bohr Institute and Stony Brook before joining Purdue in 1979, retiring in 2024. He's well known for studying “fifth force” effects, reanalyzing the Eötvös experiment, and exploring whether solar neutrinos influence beta decay. His work also extended into the Casimir effect and neutrino detection.

Rump et al. *Military Medical Research* (2021) 8:3
<https://doi.org/10.1186/s40779-020-00291-3>


PERSPECTIVE
Open Access

Preparing for a “dirty bomb” attack: the optimum mix of medical countermeasure resources



Alexis Rump*, Patrick Ostheim, Stefan Eder, Cornelius Hermann, Michael Abend and Matthias Port


Abstract


Background: In radiological emergencies with radionuclide incorporation, decorporation treatment is particularly effective if started early. Treating all people potentially contaminated (“urgent treatment”) may require large antidote stockpiles. An efficacious way to reduce antidote requirements is by using radioactivity screening equipment. We analyzed the suitability of such equipment for triage purposes and determined the most efficient mix of screening units and antidote daily doses.

Methods: The committed effective doses corresponding to activities within the detection limits of monitoring portals and mobile whole-body counters were used to assess their usefulness as triage tools. To determine the optimal resource mix, we departed from a large-scale scenario (60,000 victims) and, based on purchase prices of antidotes and screening equipment in Germany, we calculated efficiencies of different combinations of medical countermeasure resources by data envelopment analysis. Cost-effectiveness was expressed as the costs per life year saved and compared to risk reduction opportunities in other sectors of society as well as the value of a statistical life.

Results: Monitoring portals are adequate instruments for a sensitive triage after cesium-137 exposure with a high screening throughput. For the detection of americium-241, whole-body counters with a lower daily screening capacity per unit are needed. Assuming that 1% of the potentially contaminated patients actually need decorporation treatment, an efficient resource mix includes 6 monitoring portals and 25 mobile whole-body counters. The optimum mix depends on price discounts and, in particular, the fraction of victims actually needing treatment. The cost-effectiveness of preparedness for a “dirty bomb” attack is less than for common health care, but costs for a life year saved are less than for many risk-reduction interventions in the environmental sector.

Conclusion: To achieve economic efficiency, a high daily screening capacity is of major importance to substantially decrease the required amount of antidote doses. Among the determinants of the number of equipment units needed, the fraction of the potentially contaminated victims that actually needs treatment is the most difficult to assess. Judging cost-effectiveness of the preparedness for “dirty bomb” attacks is an issue of principle that must be dealt with by political leaders.





What If North Korea Set Off an EMP Over the Pacific #shorts

Beyond alarmism: A realistic assessment of Bushehr's plutonium risk

By Sasan Karimi

Source: <https://thebulletin.org/2026/07/beyond-alarmism-a-realistic-assessment-of-bushehrs-plutonium-risk/>



July 03 – A recent *Bulletin* [article](#) presents a highly securitized interpretation of Iran's civilian nuclear program. By framing the spent fuel inventory of the Bushehr nuclear power plant as an imminent proliferation threat, Henry Sokolski attempts to construct a new rationale for further political pressures on Iran at a moment when nuclear



diplomacy remains fragile and unresolved. On June 17, the United States and Iran signed a memorandum of understanding, and the parties have until mid-August to find a new agreement.

Sokolski's article attempts to employ technical terminology and quantitative estimates to convey scientific credibility. But its central claims overlook critical operational, legal, and technological realities of the Bushehr nuclear plant. More important, it conflates the theoretical existence of plutonium in spent fuel with the practical capability to manufacture a nuclear weapon. This alarmist narrative reflects a broader pattern of securitization that has long surrounded Iran's nuclear activities: It transforms ordinary aspects of a safeguarded civilian nuclear fuel cycle into extraordinary security threats requiring exceptional international security and political responses. Three dimensions need to be considered to offer a realistic evaluation of the proliferation risk at Bushehr: the technical characteristics of the plant's spent fuel, the safeguards architecture governing the reactor, and the legal rights guaranteed to non-nuclear-weapon states under the Treaty on the Non-Proliferation of Nuclear Weapons (NPT).

No "weapon-ready" plutonium

Sokolski's most dramatic assertion is that Bushehr's existing spent fuel contains enough plutonium for more than 200 nuclear weapons. While technically true in a purely arithmetic sense—plutonium isotopes exist within irradiated fuel assemblies—this statement is deeply misleading because it ignores the distinction between "reactor-grade" and "weapons-grade" plutonium. The plutonium produced in Bushehr is reactor-grade, which is produced under conditions of high burnup during normal power-reactor operations. This material [differs fundamentally](#) from weapons-grade plutonium in isotopic composition. Technical studies, including research conducted at institutions such as Princeton University, [demonstrate](#) that reactor-grade plutonium typically contains roughly 60 percent plutonium 239 and approximately 23 percent plutonium 240, whereas weapons-grade plutonium contains around 93 percent plutonium-239 and only about 6 percent plutonium-240. This distinction is not merely academic. The elevated concentration of plutonium 240 creates severe engineering complications because the isotope undergoes spontaneous fission, [emitting neutrons at random times](#). These stray neutrons significantly increase the risk of pre-detonation, meaning that a nuclear device could begin fissioning prematurely before optimal compression is achieved. Any attempt to manufacture a reliable weapon from reactor-grade plutonium would therefore require extremely advanced implosion engineering, sophisticated diagnostics, and repeated nuclear testing.

Numerous Western technical assessments, including those by the [US Energy Department](#) and the [National Academy of Sciences](#), acknowledge that while reactor-grade plutonium is not unusable in principle, producing a reliable and militarily credible nuclear weapon from it is exceptionally challenging and typically requires sophisticated design capabilities and extensive testing infrastructure. Iran has neither conducted nuclear tests nor demonstrated mastery of such advanced weapons engineering. Still, Sokolski largely treats the conversion of Bushehr's plutonium into operational nuclear weapons as if it were a straightforward industrial process.

The omission becomes even more significant when considering the implications of high burnup.

After 15 years of reactor operations, the isotopic degradation of Bushehr's spent fuel further increases the proportion of

undesirable isotopes, including [plutonium 240](#). This intensifies the technical barriers associated with implosion design.

In modern compact nuclear warheads, especially those intended for ballistic missile delivery, highly advanced "two-point" implosion systems are often discussed because they reduce weight and size. But such systems are extraordinarily sensitive to asymmetrical compression and stray neutron emissions. Reactor-grade plutonium, precisely because of its high content in plutonium 240, makes these challenges [dramatically more difficult](#).

Historically, the United States's "Fat Man" device used against Nagasaki relied on a complex 32-point implosion system to achieve symmetric compression of the plutonium core. A two-point system represents a far more advanced level of explosive lens engineering requiring extremely precise geometries and timing mechanisms. The idea that a state could simply extract reactor-grade plutonium and rapidly convert it into usable nuclear bombs overlooks the substantial scientific and engineering challenges involved, as demonstrated by decades of experience accumulated by established nuclear powers. No credible evidence has ever demonstrated that Iran possesses validated designs for such advanced implosion systems or has successfully tested them under operational conditions. The mere presence of reactor-grade plutonium cannot reasonably be equated with an immediate weapons capability.

No industrial reprocessing capability

Another major flaw in Sokolski's article is its treatment of the plutonium separation process itself. Plutonium embedded in spent nuclear fuel is not directly usable. Extracting it requires large-scale industrial reprocessing infrastructure involving highly complex chemical operations, radiation shielding systems, remote handling technology, and specialized waste management capabilities. The article briefly references small-scale Iranian laboratory experiments conducted decades ago, but this is fundamentally different from possessing a functioning industrial reprocessing plant.

While a minimal deterrent might not require a full-scale industrial reprocessing complex, even small-scale plutonium separation remains technically demanding, hazardous, and difficult to conceal. The challenge is not



merely producing plutonium but producing it in a form and quantity suitable for reliable weapons. Without advanced engineering capabilities and testing infrastructure, translating limited reprocessing capacity into a credible nuclear deterrent remains highly uncertain. Iran has neither declared nor been shown to possess any industrial-scale plutonium reprocessing



facility. Constructing such infrastructure would require enormous financial investment, highly specialized technology, extensive personnel training, and considerable time. In addition, the creation of a covert industrial reprocessing capability would be extraordinarily difficult to conceal. Such facilities produce distinct environmental signatures and would almost certainly be detected through [International Atomic Energy Agency \(IAEA\) inspections](#), satellite surveillance, or national intelligence monitoring long before becoming operational. The Sokolski article also draws an oversimplified comparison between plutonium metallurgy and uranium metallurgy. In uranium weapons development, uranium hexafluoride can be converted into metallic uranium through comparatively well-established processes. Plutonium metallurgy, however, is substantially more complicated. Following chemical separation, plutonium must be stabilized in suitable crystalline phases, particularly the [delta phase](#), before precision casting can occur. These processes involve exceptional technical sensitivity, high contamination risks, and elevated failure rates even within advanced nuclear programs. The technical obstacles associated with reactor-grade plutonium, therefore, extend far beyond the simplistic assumption that “10 kilograms equal one bomb.”

Unreasonable demands

The Sokolski article's legal and safeguards analysis is equally problematic because it downplays the extensive monitoring framework already in place for Bushehr.

The entire fuel cycle of the Bushehr reactor has remained [under comprehensive safeguards](#) administered by the IAEA. Fuel loading, unloading, storage, and inventory management are subject to regular verification procedures. According to publicly available safeguards reporting, IAEA inspectors conducted verification activities related to Bushehr's spent fuel as recently as [this month](#), while Iranian authorities stated that the associated Physical Inventory Verification process was completed in late 2025.

Claims suggesting lengthy periods without inspector presence at Bushehr are, therefore, misleading. The reactor operates within a routine safeguards schedule consistent with international practice for civilian nuclear power reactors.

The Sokolski article similarly neglects the longstanding fuel return arrangement between Iran and Russia. Under existing agreements, Russia committed to [taking back spent fuel](#) from Bushehr for reprocessing. Delays

in implementation cannot reasonably be interpreted as evidence of weapons intent.

Particularly striking is the article's suggestion that responsibility for spent fuel management should somehow be transferred to states such as Saudi Arabia or other Persian Gulf countries. While Sokolski does not explicitly propose that these countries take physical possession of the spent fuel, formulations like “assume responsibility” and “including” are vague enough to suggest that covering costs is only one component of a broader set of obligations. Responsibility for spent fuel management typically entails involvement—direct or indirect—in logistical coordination, regulatory oversight, liability frameworks, and security arrangements. These are not trivial extensions of a purely financial role.

Even if interpreted narrowly as cost-sharing, the proposal still raises questions about feasibility and precedent. Transporting and handling highly radioactive spent fuel requires extensive infrastructure, operational experience, regulatory systems, and security arrangements. Assigning these responsibilities—financial or otherwise—to third-party regional actors that are not directly involved in reactor operations or fuel ownership and have limited relevant



experience would create significant safety and security concerns and have little precedent within established international nuclear practice. This proposal, therefore, appears more political than grounded in established technical or institutional practice. The article also advocates forms of “real-time” monitoring that exceed normal safeguards

the Treaty to develop research, production and use of nuclear energy for peaceful purposes without discrimination.”

Attempting to delegitimize the expansion of Iran’s civilian nuclear infrastructure, therefore, risks undermining one of the treaty’s foundational bargains: non-proliferation obligations in exchange for access to peaceful nuclear technology. The



standards. In practice, the IAEA does not require continuous five-minute live-feed surveillance for commercial power reactors, including in Western countries. Even the [Additional Protocol](#) does not mandate such extraordinary levels of intrusive monitoring. Iran’s suspension of voluntary implementation of the Additional Protocol occurred in the context of the unilateral US withdrawal from the Joint Comprehensive Plan of Action (JCPOA) in 2018 and the subsequent failure of European parties to fully implement their commitments. Regardless of political disagreements surrounding that decision, demanding safeguards measures from Iran beyond internationally established norms creates a discriminatory standard inconsistent with the principle of equal treatment under the NPT.

A broader political objective

The policy recommendations advanced in Sokolski’s article ultimately reveal a broader strategic objective extending far beyond nonproliferation concerns.

Calls to halt construction of additional Bushehr reactors effectively amount to demands for permanent limitations on Iran’s civilian nuclear development. But the NPT’s [Article IV](#) explicitly recognizes the “inalienable right of all the Parties to

article’s arguments regarding uranium metal and plutonium chemistry are similarly expansive. Neither the NPT nor existing safeguards agreements prohibit peaceful research involving these materials when conducted under safeguards. Even under the JCPOA framework, Iran did not permanently renounce all such activities. Equally questionable is the suggestion that Bushehr reactor fuel could somehow be diverted, reprocessed, and transformed into nuclear weapons within 90 days. Such a timeline is operationally unrealistic. The transportation of spent fuel, construction or activation of reprocessing capabilities, chemical separation, metallurgical conversion, weapons fabrication, and integration into deliverable systems would require far longer—likely many months at minimum, if not years—while remaining highly visible to international monitoring systems throughout the process.

Any diversion attempt would almost certainly trigger immediate detection through safeguards inspections, satellite imagery, environmental sampling, or intelligence surveillance.

Focus on real issues

Sokolski’s article presents a technically selective and politically charged



interpretation of Bushehr's spent fuel inventory. By emphasizing theoretical plutonium quantities while ignoring the immense practical barriers associated with reactor-grade plutonium, industrial reprocessing, safeguards monitoring, and weapons engineering, the article constructs an exaggerated threat narrative disconnected from operational realities. The broader implication is not merely analytical alarmism but the continuation of a longstanding effort to securitize Iran's peaceful nuclear activities beyond the legal framework established by the NPT and previous negotiated agreements. Framing Bushehr as a latent "200-bomb" threat risks creating political justification for additional coercive measures, including expanded sanctions, indefinite technological restrictions, or even military escalation under the pretext of preventing an "imminent" diversion scenario.

Introducing the Bushehr plutonium issue into ongoing diplomatic negotiations would likely complicate and prolong ongoing talks without contributing meaningfully to realistic nonproliferation objectives. At a moment when achieving a credible and technically grounded diplomatic framework remains urgent, negotiations should focus on verifiable and practical arrangements rather than hypothetical scenarios lacking operational plausibility.

A balanced assessment of Iran's nuclear program must distinguish between theoretical possibilities and realistic capabilities. Failure to preserve that distinction risks transforming technical discourse into political securitization—precisely the dynamic that has repeatedly weakened constructive nuclear diplomacy unilaterally over the past two decades.

Sasan Karimi is an assistant professor of global studies at the University of Tehran. Karimi is also the international politics program director of the Nuclear Watch think tank and vice president of the PAIAB Institute. Karimi holds a B.Sc. in physics from Shahid Beheshti University, continued his M.Sc. at Sharif University of Technology, and has a Ph.D. in political philosophy from the University of Tehran.



CIA's SHOCK AI ARMAGEDDON WARNING: 'Digital Nukes' Could Replace Traditional Warheads | WATCH



Introduction: Doomsday and the movies

By Dan Drollette Jr

Source: <https://thebulletin.org/premium/2026-07/introduction-doomsday-and-the-movies/>

July 08 – Since the dawn of the atomic era, the threat of nuclear annihilation has played a role in the popular imagination. A remarkable number of horror movies—in America, Europe, and Asia— have used radiation as a plot device, while others have focused on the sheer scale of devastation that would occur from an exchange of nuclear weapons.

As time went on and new existential threats arose—in the form of climate change, pandemics, genetic manipulation, artificial intelligence, and the debasement of democracy, to name just a few—films tried to keep up, by incorporating these themes into their worlds.

For this issue of our magazine, the staff of the *Bulletin of the Atomic Scientists* thought it would be enlightening (and fun) to learn more about what the movies have to say about the collective fears of our society when it comes to Doomsday.

And, in the case of science fiction films made a mere couple of decades ago, to find out if Hollywood's predictions of a dystopia in the 2020s turned out to have a kernel of truth.

The essays found here delve into whether these films have or have not been effective in bringing ethical concerns to light, raising public awareness, and inspiring public activism—and even take a stab at the quality of these movies as art. The articles go into some of the best (and mention some of the worst) depictions of existential risk.

The result is by no means exhaustive, and no doubt some readers will notice that one of their favorite Doomsday movies is missing. This was in part purposeful, as one goal of this special issue was to highlight films that were little-noticed when they first appeared, only to become big later.

As an editor, one thing that I found fascinating in preparing this special issue was how often certain titles came up among nearly all of our potential authors. Experts in the realm of nuclear risk and international relations, for example, kept mentioning the name of what to me was an obscure, 1980s British apocalyptic film—which, it turns out, may very well be the first feature-length film to depict the effects of a nuclear winter. In "[The brutal, powerful legacy of *Threads*](#)," Dan Drezner, academic dean of international politics at the Fletcher School of Law and Diplomacy at Tufts University, examines how this film goes far beyond depicting the initial

devastation of a nuclear weapons exchange to delve into the long-term degradation of social order, language, and human


development across generations. He also found that *Threads* contained a powerful critique of civil defense and emergency planning, using the subplot of city officials trapped in a bunker to illustrate the absurdity of what sociologists call "fantasy documents"—disaster plans that are completely detached from reality when it comes to worst-case outcomes.

Not all of the authors for this issue are experts in international relations; this magazine also contains the viewpoint of someone whose full-time, paid job is as an expert on movies. In "[A film about nuclear war 'too horrifying for the medium of broadcasting.'](#)" Rebecca Fons—a film curator and director of programming at the historic Iowa Theater, who is now Head of Cinema at London's

Barbican Centre (arguably the largest performing arts complex in Europe)—describes the curious, twisted tale of a 47-minute movie about nuclear war that those who commissioned it found was too realistic, and too painful, to watch. Called "The War Game," the BBC quickly yanked it from distribution—only for it to be rediscovered, and revered, years later.

And not all Doomsday-related tales are downers. In "[How the movies \(and the rest of the media\) made me into an anti-nuclear weapons campaigner](#)," legendary activist Helen Caldicott describes how books and films such as 1959's "On the Beach" inspired her. Nuclear weapons and nuclear testing—and, just as important, media coverage of them—may have kept following her around like a radioactive shadow. But rather than making her want to run and hide, their presence made her determined to respond, by using the power of that same media.

Now in her late eighties, Caldicott also describes how an effort to muzzle press coverage of her and the anti-nuclear weapons movement backfired: In 1982, a young Canadian filmmaker named Terre Nash recorded a speech Caldicott was giving in a city which had long been home to a large and important US Strategic Air Command Base. Titled *If You Love This Planet*; the film was labeled as so-called "foreign propaganda" by the



Reagan administration's Department of Justice in an attempt at censorship—which only served to create a media firestorm of interest. The film went on to win an Oscar for Best Short Documentary.

In her acceptance speech at the Academy Awards, Nash thanked the US Justice Department “for their tremendous effort in promoting *If You Love This Planet*.”

Speaking from a less glamorous pulpit, David Kirby highlights in “[Gattaca and the quiet Doomsday of genetic determinism](#)” how films can sometimes do a better job than academic treatises in bringing alive concerns about new technologies and how they will be applied. Kirby—a former geneticist

himself—says that this 1997 movie was eerily accurate in many respects, in its imagination of a society where a person's worth, opportunities, and identity are determined exclusively by their genetic profile; meanwhile, willpower, imagination, and labor count for little. Genetic probability is mistakenly reduced to absolute certainty, and a person's future is determined within seconds of being born. The result, Kirby finds, is a different, slow-burning form of existential risk: “A society primed to interpret genetic information as destiny may gradually narrow its imagination of what is possible.” Or, as the tagline for *Gattaca*'s ad campaign put it: “There is no gene for the human spirit.”

Dan Drollette Jr is the executive editor of the *Bulletin of the Atomic Scientists*. He is a science writer/editor and foreign correspondent who has filed stories from every continent except Antarctica. His stories have appeared in *Scientific American*, *International Wildlife*, *MIT's Technology Review*, *Natural History*, *Cosmos*, *Science*, *New Scientist*, and the *BBC Online*, among others. He was a TEDx speaker to Frankfurt am Main, Germany, and held a Fulbright Postgraduate Traveling Fellowship to Australia—where he lived for a total of four years. For three years, he edited CERN's on-line weekly magazine, in Geneva, Switzerland, where his office was 100 yards from the injection point of the Large Hadron Collider. Drollette is the author of “Gold Rush in the Jungle: The Race to Discover and Defend the Rarest Animals of Vietnam's “Lost World,” published in April 2013, by Crown. He holds a BJ (Bachelor of Journalism) from the University of Missouri, and a master's in science writing from New York University's Science, Health and Environmental Reporting Program.

The brutal, powerful legacy of *Threads*

By Daniel W. Drezner

Source: <https://thebulletin.org/premium/2026-07/the-brutal-powerful-legacy-of-threads/>

July 08 – In the annals of apocalyptic cinema, the nuclear war subgenre might have the widest possible variance of quality. There are universally acknowledged classics, like Stanley Kubrick's [Dr. Strangelove](#). There are important but [flawed films](#), like Sidney Lumet's [Fail Safe](#). There are movies that focus more on the post-apocalyptic phase that include some unintentionally goofy offerings, like Jack Smight's [Damnation Alley](#) or the Hughes brothers' [The Book of Eli](#). After a temporary lull in this genre, Kathryn Bigelow's recent [A House of Dynamite](#) offered a pungent reminder that the risk of nuclear war remains nonzero.

Among devotees of this particular subgenre, however, there is one movie that is venerated above all others: Mick Jackson's 1984 UK film *Threads*, currently available to watch on Tubi and to rent on Amazon Prime. *Threads* traces the slow march towards global thermonuclear war stemming from a crisis in the Middle East—and the devastating effects such a conflict would inflict on Great Britain. The reason aficionados venerate it is simple but harsh: Of all the films in the genre, *Threads* has, by far, the grimmest ending. *Threads* is the very epitome of an apocalyptic narrative in which the living would envy the dead. The movie is so unrelenting that a recent documentary about the making of the film concluded that some of its cast and crew “suffered with the trauma of being involved” (Aslett, 2025).

Is that a good thing? For apocalyptic cinephiles, the answer would be an unqualified “yes”—and that was certainly Jackson's intent.

For discerning readers of the *Bulletin of the Atomic Scientists*, however, the answer is slightly more muddled. *Threads* is a difficult watch—and therefore somewhat less effective in attracting viewers who are not fans of apocalyptic cinema. The evidence that *Threads* influenced public sentiment or public policy about nuclear weapons at the time of its release is limited at best. This is in stark comparison to two films that came out a year earlier: Nicholas Meyer's [The Day After](#) and John Badham's [WarGames](#). Unfortunately, however, *Threads*' scenario about how a nuclear war might break out has acquired new relevance in the 21st century. Furthermore, as a statement about the futility of nuclear war planning, the film is mordantly effective.

—
Like its American doppelgänger [The Day After](#) (Stover 2018), the plot of *Threads* is roughly divided between pre-war scenes of quotidian life mixed with rising panic, a terrifying sequence in which the missiles launch and mushroom clouds engulf the horizon, and then post-war scenes of despair and destruction. Another parallel is that both films eschew paying attention to high-level policy makers, focusing instead on a single community's reaction to nuclear brinkmanship followed by nuclear



war. In [The Day After](#), it is Lawrence, Kansas; in *Threads*, it is Sheffield, England.

The pre-war plot of *Threads* mostly follows Ruth and Jimmy, portrayed by Karen Meagher and Reece Dinsdale, respectively. They are two Sheffield teens who let a makeout session in Jimmy's car go a bit too far. Ruth winds up pregnant. Deciding to keep the baby, Ruth and Jimmy decide to get married and rent out a flat. This leads to some considerable British awkwardness. The dinner between Ruth and Jimmy's parents following the revelation of Ruth's pregnancy is as stilted and discomfiting as one would expect from middle-aged English parents meeting each other under such circumstances. The class differences between the two families—Ruth's family is middle-class while Jimmy's is more working class—only heighten the awkwardness.

As the nuclear crisis unfolds, all of the Sheffield characters begin preparing for the worst—including Ruth and Jimmy's families. Large demonstrations and counter-demonstrations emerge as the prospect of war nears. Hospitals decide to discharge patients so they can prepare for those wounded in the war. The British government cracks down on dissent, arresting a swath of protest leaders. The government also broadcasts some truly ghastly advice of what to do in case of a nuclear attack. Local leaders start organizing for a war footing (more on this later). Ruth's family prepares their cellar as a bomb shelter, whereas Jimmy's family constructs a flimsier lean-to from mattresses and wood. As the global conflict escalates from conventional to nuclear exchanges, the film announces that the Soviet Union has dropped 210 megatons on the United Kingdom. The nuclear weapons have a devastating effect, killing tens of millions, destroying two-thirds of the housing stock, and unleashing fires and devastation across the country.

The effect on the film's protagonists is equally severe. Jimmy, attempting to find Ruth as the missiles are flying, is never seen again after the first nuclear explosion and is presumed dead. Jimmy's parents barely survive the initial attack. Ruth's family has somewhat better initial luck: Ruth, her parents, and her grandmother survive in their basement for a spell.

British society collapses after the nukes are dropped. Survivors are mostly left to fend for themselves (again, more on this later). Food becomes scarce as a result of looting and nuclear winter-induced crop failures. Declaring martial law, the government attempts to use food as currency, awarding it for hard labor and withholding it as punishment. Summary executions occur to deter looting.

The fate of Jimmy and Ruth's families is equally severe. Jimmy's brother is killed in the initial nuclear strike, his mother dies from radiation burns, his father dies soon after from radiation poisoning, and his sister is arrested and imprisoned. As for Ruth, following the death of her grandmother, she flees her family's shelter. Sojourning to Jimmy's house, she finds only his mother's corpse. When she returns to her family's home, she discovers that looters have murdered her parents.

Still pregnant, she escapes Sheffield and makes her way to the countryside. Ruth gives birth to her daughter Jane alone in an abandoned barn, using her teeth to cut the umbilical cord.

Up until this moment in the film, the structural differences between *Threads* and *The Day After* are relatively minor, mostly reflecting the cultural differences between rural America and industrial England. In both movies, ordinary families are torn apart from the horrors of nuclear war and the ensuing nuclear winter. At this point, however, *Threads* begins its brutal final act, daring to go to a level of despair far beyond *The Day After's* more graceful ending.

Threads flash forwards a decade, notifying viewers that Great Britain's surviving population has plummeted to a medieval level of 4,000,000 to 11,000,000 people. (By comparison, Great Britain has about 80,000,000 people today.) With no fuel or electricity, the survivors have little choice but to cultivate crops by hand. The children born after the war are developmentally delayed for multiple reasons: fetal radiation exposure, a dearth of organized schooling, and parents traumatized by the war. Jane is no exception, possessing only a rudimentary vocabulary. Ruth and Jane till agricultural fields along with everyone else. Ruth, prematurely aged and blinded by cataracts from exposure to nuclear radiation, dies when Jane is 10. Jane is completely unmoved by her mother's passing, simply taking the last of Ruth's possessions and heading out to explore the countryside.

Threads then flash-forwards another three years. Some industry and electricity has returned but living conditions remain appalling and food is still scarce. Jane gloms onto two boys, Gaz and Spike, and as a group they steal food to stay alive. When they are caught in the act, one of the boys is shot dead. Jane and the surviving boy fight over the remaining food, a fight that quickly turns into rape. Nine months later, Jane gives birth in a makeshift hospital. The nurse wraps Jane's mute baby in a bloody sheet and gives it to Jane. She looks at her new child with dumbfounded disgust—at which point the film ends.

The final act of *Threads* makes its message clear: as devastating as the initial nuclear war would be to an advanced developed society, successive generations would face even greater challenges. *Threads* basically predicts that even if nuclear war was not an immediate extinction-level event, devolution over the next few generations would likely bring an end to human civilization.

While *Threads* retains a cult following to this day, it would be difficult to claim that it was as influential as *The Day After*. The latter film, which starred Oscar-winner Jason Robards and aired on ABC following weeks of controversy and public debate, garnered more than 100 million viewers, earning the highest ratings ever recorded for a made-for-television movie in American history (McKairnes 2023).





Despite criticism from the Reagan administration, President Reagan watched the film, and it clearly had an impact. Viewing it at Camp David, Reagan wrote in his diary, "It is powerfully done... It's very effective & left me greatly depressed... My own reaction was one of our having to do all we can to have a deterrent & to see there is never a nuclear war" (Ronald Reagan Presidential Library and Museum 1983). Four years later, when Reagan signed the Intermediate Nuclear Forces treaty, he reportedly telegrammed director Nicholas Meyer to say, "Don't think your movie didn't have any part of this, because it did" (*The State Journal-Register* 2015). To be sure, *Threads* generated waves in the United Kingdom, but its global effect was far more muted. It was not reviewed widely in the United States—and, because it came out a year after *The Day After*, occupied familiar ground for American audiences. Although director Mick Jackson has claimed in interviews that he knows from first-hand sources that Reagan also watched *Threads*, there is no contemporaneous evidence to buttress that claim (*The Scotsman* 2009).

While *The Day After* was hailed for its brutal depiction of life after a nuclear war, its portrayal was relatively gentle compared to *Threads*. Most of its major characters are portrayed sympathetically—and most survive the initial nuclear exchange. The hospital where the protagonist worked continued to function, albeit at a bare-bones level. The final scene, in which Jason Robards' character returns to his home to see squatters occupying a pile of rubble, treats all of the characters humanely. In interviews, Meyer has noted that his film presented "nuclear war on a good day as opposed to what the reality is likely to be" (Alter 2022). In contrast, there are few characters to root for in *Threads*—Jimmy cheats on Ruth after she's pregnant, for example. Post-nuclear authoritarianism is also more evident. The unrelenting

downward spiral of *Threads* might be more accurate—but that makes watching it all the more painful.

In 1982, Mick Jackson produced and directed a short documentary for a science program called "A Guide to Nuclear Armageddon." This provided the grist for a later, two-hour drama, that was to be called "Threads" a few years later. Both films focused on the subject of nuclear war from the viewpoint of everyday people. BBC screen capture.

That said, *Threads*' cult following is deserved—though not merely for the bleakness of the final section. In retrospect, the film's depiction of the run-up to nuclear war holds greater contemporary relevancy. And *Threads*' depiction of emergency planning in case of nuclear war might be its

greatest and most powerful legacy—a darkly humorous tribute to the absurdity of planning for the end of the world.

The Day After's origin scenario for the start of World War Three was centered on the then-controversial placement of Pershing II missiles in West Germany and the Soviet Union's bellicose response to it. While extremely relevant to the early 1980s, that particular scenario has obviously been overtaken by events. *Threads*, on the other hand, posits a crisis escalation that feels all too contemporary. The trigger for the start of nuclear war is US-backed regime change in Iran. The Soviet Union invades northern Iran in response. The crisis escalates when the nuclear submarine *USS Los Angeles* sinks in the Persian Gulf. American troops are then deployed to western Iran's oilfields, and the United States sets an ultimatum for a joint withdrawal. US and Soviet forces engage in a limited nuclear exchange in Mashhad, in northeastern Iran. US and Soviet warships subsequently exchange fire, and the aircraft carrier *USS Kitty Hawk* is sunk in the Persian Gulf, where it is stationed along with elements of the Israeli navy. The United States then enforces a blockade on Cuba, and the USSR orders the departure of NATO forces from Berlin. The global launch of ICBMs ensues.

The parallels to geopolitics in 2026 are mildly discomfiting, to say the least. As of the spring of this year—at the time of writing this article—the United States is bombing Iran in a haphazard effort to foment regime change. The Trump administration is also enforcing a blockade of Cuba following military action in Venezuela to cut off its energy sources. In response to the joint US-Israeli bombing, the theocratic regime in Iran is launching missiles and drones across the region from Cyprus to Saudi Arabia. Israel is expanding the war by bombing Beirut repeatedly in an effort to diminish Hezbollah. Iran's functional closing of the



Strait of Hormuz threatens to have ripple effects across the global economy.

The obvious difference, of course, is the absence of the Soviet Union as a superpower rival to the United States. To date, US bellicosity in the Middle East has mildly exacerbated tensions with the People's Republic of China (Chia 2026)—but the risk of a military crisis escalation remains muted. Nonetheless, the end of the rules-based, post-World War II liberal international order has heightened the danger of a regional conflict unexpectedly escalating into a more global conflict. Less than three weeks into the Iran war, President Trump publicly referenced the possibility of nuclear weapons (Trump BlueSky 2026). It is noteworthy that all of the actions described in the previous paragraph took place after what Canadian Prime Minister Mark Carney described in Davos in January of this year as “a rupture in the world order, the end of a pleasant fiction and the beginning of a harsh reality, where geopolitics, where the large, main power, geopolitics, is submitted to no limits, no constraints” (Carney/World Economic Forum 2026). In other words, anything could happen and the parallels between the *Threads* scenario and the present day are more than a little unsettling.

The more darkly amusing legacy of *Threads* is that it is perhaps the best fictional rendering of what sociologist Lee Clarke labeled “fantasy documents”—disaster planning for worst-case outcomes that alter the underlying system, in which the plans seem detached from reality. In fantasy documents, unrealistic assumptions are often made about catastrophic outcomes in order to promulgate standard operating procedures for such contingencies. As Clarke noted in his book *Mission Improbable: Using Fantasy Documents to Tame Disaster*, preparing for nuclear war—particularly in the area of nuclear civil defense—often leads to absurd assumptions and plans. Clarke explained, “To mount an effective case that it makes sense to talk about defending against thermonuclear bombs and restoring to normalcy a society devastated by nuclear destruction demands that the singularity of nuclear war be dispensed with.... [thereby offering] the potential to transform the rarefied world of nuclear war planning into something more mundane and hence more rational” (Clarke 1999).

Any large organization needs to engage in planning in order to convert radical uncertainty into more psychologically manageable forms of risk—a point that Clarke noted in his book. As Dwight Eisenhower famously said, “plans are worthless, but planning is everything” (Dwight D. Eisenhower Presidential Library Quotes 1957). Clarke, however, is hardly the only social scientist to note the absurdity of how governments plan for contingencies surrounding a nuclear war (Ellsberg 2017). Sociologist Lynn Eden concluded in her 2004 book *Whole World on Fire: Organizations, Knowledge, and Nuclear Weapons Devastation* that for decades, US

nuclear planners radically underestimated the damage that would occur from a thermonuclear war from uncontrollable nuclear fire (Eden 2004). As Eden explained in a 2021 article for the *Journal of Peace and Nuclear Disarmament*: “As they work, planners strip out the human meaning of the consequences of the hypothetical actions they are planning. They do this because the complex and detailed tasks that command their attention have already removed vivid references to human society” (Eden 2021).

All of these tropes are on vivid display in a subplot of *Threads* that follows Sheffield city comptroller Clive Sutton, played by Harry Beety. Sutton is tasked by the British central government with forming an emergency preparedness response committee in case a nuclear war does occur. Sutton obliges, performing his duties like any diligent bureaucrat, to the point of ignoring his wife's entreaties to stay with her as the crisis escalates. Clive and his team head to their emergency bunker as the crisis worsens, securing stores of food to distribute if war breaks out. They ride out the attack in their bunker, trying to reach other government units while getting buried under a mountain of rubble.

Despite the committee's preparations, they offer no immediate public response. At first this is to avoid radioactive fallout. However, as the committee's plight becomes more desperate with no conceivable means of exit from their bunker, their callousness towards the other survivors increases. At one point, an exasperated Clive responds to a request to distribute food with, “What's the point of wasting food on people who are going to die anyway?!” Another official suggests providing workers only 1,000 calories a day and everyone else only 500 a day. At the end of their tether, the committee bickers with one another as the situation continues to deteriorate. They accomplish little during their time in the bunker. A month after the attack, Clive's committee is found dead in their bunker, having achieved nothing of consequence. The entire subplot is a vivid demonstration of what happens when fantasy planning meets brute reality.

As the memory of Cold War nuclear brinkmanship recedes, it is all too easy to dismiss *The Day After* and *Threads* as historical artifacts with little relevance to the 21st century. That would be a mistake. Leaders like Trump and Russia's Vladimir Putin are now brandishing nuclear threats in response to military setbacks (Mills 2024). Such threats might sound unrealistic and detached from reality, but as this journal notes, it is now 85 seconds to midnight (*Bulletin of the Atomic Scientists* 2026). In a world in which there has been no actual historical parallels to envision a thermonuclear war, leaders may have no choice but to rely on fictional narratives to guide their thinking (Jones and Paris 2018). If that is the case, then the apocalypse cinephiles are right to advocate for *Threads*. Leaders need to know just how bleak the aftermath of a real nuclear war would be.



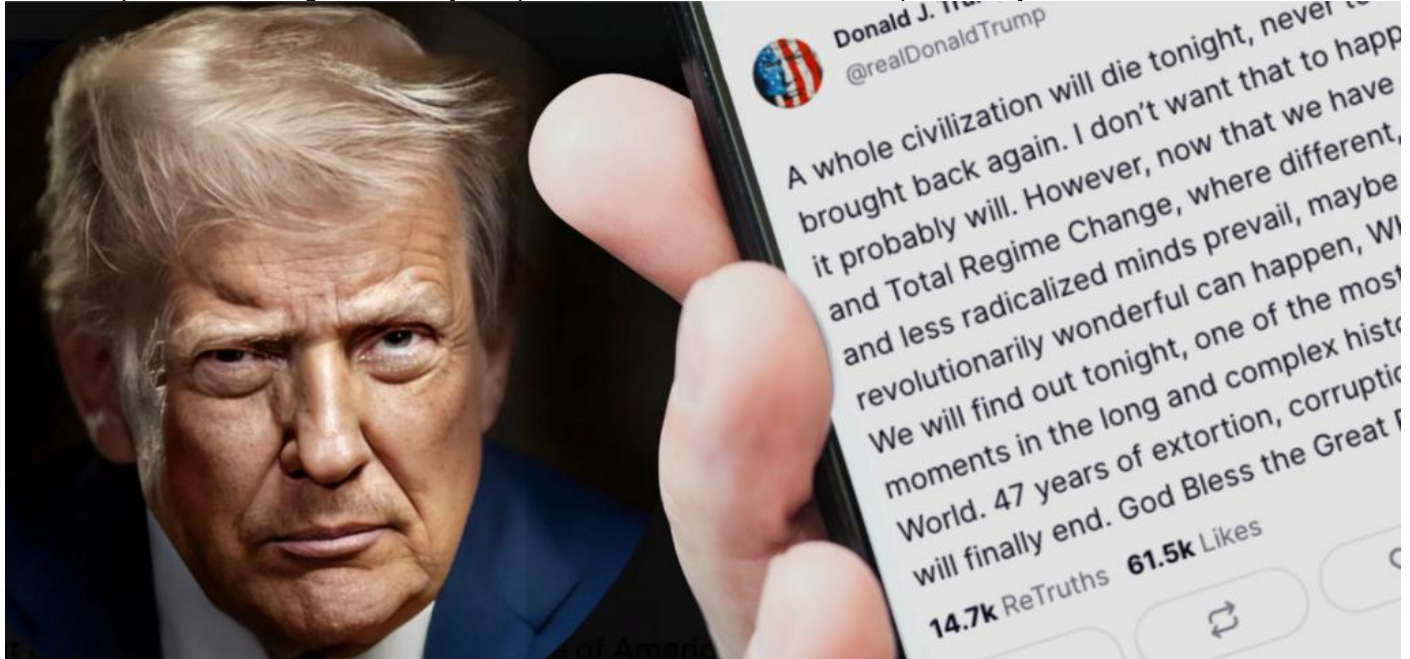
► References are available at the source's URL.

Daniel W. Drezner is academic dean and distinguished professor of international politics at the Fletcher School at Tufts University, where he is co-director of Fletcher's Russia and Eurasia Program. His research focuses on economic statecraft, US foreign policy, global political economy, and power in world politics. Drezner is a nonresident senior fellow at the Chicago Council on Global Affairs, and an academic partner at State Street. He's authored seven books, including *Theories of International Politics and Zombies* and *All Politics is Global*.

Why Trump's nuclear blackmail of Iran failed spectacularly

By Jack Kennedy

Source: <https://thebulletin.org/2026/07/why-trumps-nuclear-blackmail-of-iran-failed-spectacularly/>



On April 7, President Trump posted a message on his Truth Social platform asking Tehran to agree to a ceasefire or face severe consequences, in what many considered as being an implicit nuclear threat. (Illustration by François Diaz-Maurin; original photo by KLYONA / depositphotos.com)

July 07 – On the morning of Tuesday, April 7—38 days into the US-Israeli war on Iran—President Donald Trump posted on his social media platform, Truth Social, that “[a whole civilization will die tonight, never to be brought back again](#).” He added that the promised destruction could be averted—“maybe something revolutionarily wonderful can happen”—if, he implied, the Iranian government made a deal with him. The post capped off a [series of escalating threats](#) the US president had made towards Iran. He had set a [deadline](#)—first April 7, then moved to April 8—for Tehran to agree to a ceasefire and threatened increasingly severe consequences as it approached. Among them, Trump threatened to destroy civilian infrastructure on what he called “[Power Plant Day and Bridge Day](#).”

The statements were condemned around the world, from the UN Secretary General to the Pope. Deliberate targeting of civilian infrastructure is a war crime; eradicating an entire nation or culture would be genocide, and [threatening to commit genocide is itself a crime](#). But many people read the

threat with an additional layer: Trump was threatening to use nuclear weapons.

The logic is simple enough: Destruction on the scale the president promised would be essentially impossible to achieve using conventional munitions, making a nuclear strike the obvious implication of Trump's threat.

Interpreted that way, Trump would enter a small club of leaders who, finding themselves locked in apparently unwinnable wars, have reached for the threat of nuclear escalation to try to secure a favorable outcome on their own terms. [Richard Nixon tried it in 1969](#), putting US nuclear forces on alert across the globe to scare the Soviet Union and North Vietnam into coming to the table to end the Vietnam War, then about five years in. Most familiar to everyone today is Russian President Vladimir Putin's [repeated nuclear threats](#) during his full-scale invasion of Ukraine, which reached [its peak in October 2022](#) when the Russian army appeared to be wavering. There's every reason to believe



the US president was doing exactly this, invoking the nuclear specter to try to scare Tehran into surrender after conventional bombing failed to do so. About 11 hours after Trump's Tuesday morning post, the parties did indeed [agree to a ceasefire](#). But that ceasefire saw Tehran give up nothing, and when a more lasting peace was agreed two months later, it was largely [on Iran's terms](#). In other words, Trump's threat failed. It was probably always going to.

Diagnosing the threat

President Trump did not make an explicit nuclear threat, in that he didn't use the word "nuclear." But many nuclear threats don't. One prominent threat from [Vladimir Putin](#) in the early days of the Ukraine War was similarly euphemistic, promising "consequences will be such as you have never seen in your entire history" to any state opposing Russia.

Later on April 7, after Trump's "civilization" post, Vice President JD Vance said that the United States might employ capabilities they "[so far haven't decided to use](#)" in the war. After some commentators noted this could imply nuclear weapons, an official White House X account [shot back](#) that "literally nothing [Vance] said here 'implies' this, you absolute buffoons." However, it's difficult to imagine what else it could have meant; the US military has used an [exceptionally wide range of platforms and weapons against Iran](#), including its [largest single conventional munition](#).

The White House denial itself means little: Denials often are thrown in alongside nuclear threats. In March 2022, a Kremlin spokesperson [insisted](#) that "no one is thinking about using, about—even about [the] idea of using a nuclear weapon." Threats, both implicit and explicit, continued unabated. Moreover, the denial extended only to Vance; asked if Trump's original post meant he was considering nuclear use, the White House Press Office declined to definitively say he wasn't and [added](#) that "only the President knows where things stand and what he will do."

The core issue is that it can be difficult to define when a nuclear threat has been made because in almost every case—even when the nuclear dimension is explicit—the threatening party purposefully [maintains some ambiguity](#). In theory, this gives the threatener room to maneuver; by not clearly committing to nuclear use under a given condition, it can back down without damaging its reputation and credibility. And since nuclear weapons are so frightening, even an ambiguous threat may have coercive value.

According to White House sources, Trump [believes](#) "seeming unstable could help spur the Iranians to negotiate." This follows in the footsteps of Nixon's infamous "[madman theory](#)," which held that if he appeared unconstrained and capable of unlimited escalation—including [past the nuclear threshold](#)—this could be a powerful coercive tactic.

Ultimately, the most obvious and compelling reason to call it a nuclear threat is that it was widely interpreted as one.

Those who interpreted Trump's post as a nuclear or potentially nuclear threat include former Vice President and Democratic presidential nominee [Kamala Harris and other US Democratic politicians](#); former Trump allies [Anthony Scaramucci](#) and [Tucker Carlson](#); reporters for outlets like [The Atlantic](#), [The New Republic](#), [AFP](#), and [the Associated Press](#); retired US Army General [Shawn Harris](#); [Pervez Hoodbhoy](#) writing in the *Bulletin*, and many more.

Apply the duck test: If it sounds like a nuclear threat, is treated by a great many people as such, and the threatener doesn't definitively deny it is a nuclear threat, then it is probably a nuclear threat. Whether Trump ever seriously considered using nuclear weapons against Iran is immaterial: A bluffing threat is still a threat.

The results

After the April 8 ceasefire was announced, Washington and Tehran both claimed victory. Trump [said](#) it had been agreed because "we have already met and exceeded all Military [sic] objectives," while the Iranian government [portrayed it](#) as a "victory against the United States and Israel." There was nothing to show for Trump's last-minute threats. The ceasefire was mostly just that: a cessation of active combat. This overwhelmingly benefited Iran, which, despite occasionally destructive retaliation across the region, was still being bombed much more than it was bombing. The sides agreed to negotiations—[from very different starting points](#)—which quickly [fell apart](#). Most crucially, Iran maintained its blockade of the Strait of Hormuz. The United States [protested](#) that this was "not the agreement we have" and imposed a [counter-blockade](#) on Iranian ports, but the waterway remained closed. Clearly, the ceasefire was not a matter of Tehran being intimidated into capitulation but primarily a [climbdown by Washington](#). Over the subsequent two months, the truce largely held, despite numerous violations. And on June 17, Washington and Tehran signed a [memorandum of understanding](#) (MOU), solidifying the ceasefire pending a subsequent, final agreement due by mid-August. [The immediate consensus was](#) (with some [exceptions](#)) that the MOU represents a US defeat, and a large one. Trump himself said that the deal had been signed to [prevent a "worldwide depression"](#) caused by the Hormuz blockade.

Iran has conceded nothing from the pre-war status quo, and both parties have lifted their respective naval blockades. In return, Tehran has been promised relief from pre-war sanctions and [hundreds of billions of dollars](#) for reconstruction and economic development. The United States, by contrast, has conceded on all of its war goals. Initially aiming for regime change, Trump now [says of Iranian leaders](#) that they are "nice to deal with ... strong people, smart people ... not radicalized." Having been adamant that Iran should have [no ballistic missiles](#) and [no nuclear program of any kind](#), Trump now says it would be ["unfair"](#) for Tehran



not to have [both](#). Trump's blackmail failed in the proximate sense, with Tehran making no concessions and remaining in control of the Strait. In the broader sense, it failed catastrophically: Ultimately, the United States was the party forced to make concessions to end the war, buckling under the cumulative economic pressure of the Hormuz blockade.

Why nuclear blackmail fails

Fundamentally, nuclear coercion lacks credibility. As international relations scholars Todd Sechser and Matthew Fuhrmann [conclude](#), "it is exceedingly difficult to make coercive nuclear threats believable."

Blackmail features an imbalance of stakes that deterrence does not. The use of nuclear weapons, even against an enemy that cannot retaliate, [would likely come with huge political costs](#) and unpredictable ripple effects. It's one thing to convince your adversaries you'll accept these costs if your country is being invaded, when the alternative is mass death, domination by a foreign power, and so on. It's another thing entirely to insist you'll do so to avoid making political concessions in a war of choice.

However, Trump got closer than most leaders have, historically. His rhetoric was treated by many as somewhat credible and therefore alarming. US politicians who demanded that Trump clarify he was not threatening nuclear use would not have done so if that threat were obviously implausible. "[Will Trump nuke Iran](#)" was, for a time, a question worth considering seriously.

How did he manage to convince so many that his threat was real? Reputation.

By starting the Iran war, Trump has made clear he is not overly bothered by [legality](#), ethical norms, or [domestic](#) or [international](#) disapproval. His other actions in office—from the capture of Venezuelan President Nicolas Maduro to his attacks on US democracy—have also helped build his reputation as a president constrained only by [his "own morality."](#)

This has parallels to Putin's threats to use nuclear weapons in support of his invasion of Ukraine. These threats were treated as potentially plausible precisely because his decision to invade Ukraine in the first place was so [shocking](#) and laid waste to so many assumptions about his limits.

Nixon may have craved a "madman" reputation, but there's [no evidence he ever had it](#). According to longtime Soviet Ambassador [Anatoly Dobrynin](#), while Moscow was initially wary of Nixon given his anti-Soviet rhetoric, they came to see him as "more pragmatic and realistic than other Cold Warriors," using anticommunist zeal as "just a convenient means to climb the political ladder." Nixon failed to establish himself as a "madman" because he never did anything to earn that reputation. Trump and Putin, by contrast, have real track records of extreme behavior.

Still, Putin mostly failed, and Trump failed completely. In the latter case, the likely reason was the imbalance of stakes: For

Washington, this was a war of choice; for Tehran, it was existential.

One notable difference between Trump's nuclear blackmail and that of Putin and Nixon was that the latter two channeled their threats through third parties: Putin's coercion initially focused mostly on NATO states whose material support was (and is) essential to Kyiv's resistance. Equally, Nixon's threat was calculated to be picked up by the Soviet Union, whose support kept North Vietnam afloat and fighting. By contrast, Trump has only threatened Iran directly, not least because the country [has essentially been fighting alone](#). This [third-party approach](#) has a strategic logic: NATO and the Soviet Union had much less at stake than their respective allies and much less to lose if they backed down.

Nixon's attempt still failed, largely because the signaling was too opaque and Moscow wasn't sure what was being threatened or why. On the other hand, Putin's nuclear threats in Ukraine achieved limited success: Western aid to Kyiv was [initially constrained](#) for fear of provoking escalation, only ramping up slowly over time.

In both cases, though, the primary target of those threats was particularly unmoved. Kyiv has consistently been less concerned about escalation than its allies, [requesting a NATO-imposed no-fly zone](#) early in the war even at the risk of direct combat between Russian and Western forces. In October 2022, when Moscow appeared to be threatening nuclear use against Ukraine itself, Kyiv did not slow or halt its ongoing offensive.

While we don't know if North Vietnam even comprehended Nixon's 1969 nuclear threat, a similar, more public attempt to threaten Hanoi in December 1972 also failed. After negotiations to end the war broke down in their final stages, Nixon [issued North Vietnam with an ultimatum](#) to immediately resume talks, "or else." Hanoi ignored him.

In both cases, these states were fighting for stakes they considered existential: Ukraine for its territorial integrity and independence, North Vietnam to unite its country and drive out what it viewed as an ongoing colonial occupation by the United States. Both states had also already endured enormous destruction, [stiffening their will to resist](#) and making them harder to shock with threats of further attack. It was the same in Iran. For the leaders of the Islamic Republic, losing this war would have meant losing everything. The original US-Israeli war goal was [regime change](#). After they failed to secure a quick victory, Tehran could potentially have made peace based on less maximalist concessions. But capitulation to its two archenemies would have dealt a [devastating, likely fatal blow](#) to the regime's legitimacy at home. The war has also [radicalized the regime and its supporters](#) and strengthened hardline factions, such as younger officers of the Revolutionary Guard.

For this new, more determined Iranian government, balancing the consequences of surrender against the probability that



Trump was bluffing, holding firm was the obvious choice. Ultimately, Trump failed because he tried to shove someone whose back was against the wall.

Consequences of Trump's failure

On April 23, two weeks after his threat, President Trump seemed to [disavow his previous behavior](#): “No, I wouldn’t use it,” adding that “a nuclear weapon should never be allowed to be used by anybody.”

Despite Trump’s failure to coerce Tehran and his ultimate concession in the war, his attempt will have consequences—all of them negative. Following on from Putin, it will serve to further normalize overt nuclear saber-rattling. Even aside from the nuclear dimension, explicit threats to target civilians and civilian infrastructure [are harmful to the international rule of law](#), which is already under strain.

Right from the start, the Iran war risked [encouraging nuclear proliferation](#) rather than halting it. Tehran’s attempt at nuclear hedging clearly failed. Other regimes in Iraq, Libya, and Syria that considered going nuclear but didn’t commit tell the same story. Meanwhile, nuclear-armed North Korea remains secure.

By raising the possibility that non-nuclear states are vulnerable not only to conventional attack by powerful rivals but possibly nuclear coercion or attack too, Trump has reinforced this lesson. Countries that have long considered going nuclear, [like South Korea](#), will take note. Here, too, Trump is continuing the work that [Putin started in Ukraine](#), and the world will be more dangerous for it.

Worse still, maybe, there is the risk of consequences to Trump himself.

Should the famously thin-skinned president feel humiliated by the loss of this war, Trump may become more belligerent and unpredictable as he attempts to save his credibility—or pride. On June 28, with the ceasefire under renewed strain, he [threatened](#) to “militarily complete the job,” adding, “the Islamic Republic of Iran will no longer exist.” Trump’s behavior is hard to predict: Most of his threats are empty, but occasionally he will follow through with shocking violence, as the Iran war has demonstrated. But the more the US president feels embarrassed and no longer taken seriously, the more likely that next time he threatens catastrophe—whether in Iran or [elsewhere](#)—he won’t be bluffing

Jack Kennedy is the *Bulletin*’s nuclear risk editorial fellow. He is a doctoral candidate and research associate at the Centre for International Security at the Hertie School of Governance in Berlin. His research focuses include strategic stability under conditions of multipolarity, extended deterrence, and coercive diplomacy. He previously worked at the European Peacebuilding Liaison Office, and as a journalist. He holds a master’s degree from Johns Hopkins School of Advanced International Studies.

Is London shifting from nuclear deterrence to war-fighting?

By S.M. Amadae and Tom Stevenson

Source: <https://thebulletin.org/2026/07/is-london-shifting-from-nuclear-deterrence-to-war-fighting/>



An F-35A fighter jet releases an inert B61-12 bomb during a dual-capable aircraft (DCA) test flight in the skies above Edwards Air Force Base, California, on June 27, 2019. In June 2025, the United Kingdom announced plans to acquire 12 F-35As, and there are indications that in July 2025 the US Air Force may have moved B61-12 nuclear bombs to Royal Air Force Lakenheath air base, which has recently undergone a significant upgrade to reactivate a nuclear mission. (Los Alamos National Laboratory)



July 10 – For most of the past three decades, one might almost have said the United Kingdom had a glowing record on nuclear non-proliferation: Britain was the only major nuclear power to operate a minimum deterrent based on a single delivery system; it supported international [non-proliferation initiatives](#); and the UK government had spent 30 years reducing its nuclear stockpile. No more. Instead, London now builds momentum supporting a [new wave of nuclear proliferation](#) under the auspices of NATO's nuclear sharing. In the last five years, Britain has quietly moved to increase the maximum size of its nuclear stockpile and revise its nuclear doctrine, and it is now acquiring an additional nuclear capability. In concert with the United States, the UK government is expanding the deployment of non-strategic (tactical) nuclear weapons in a way that could destabilize the strategic balance in Europe.

In doing so, London is reversing decades of gradual progress on disarmament without serious analysis of the strategic implications or public consultation, and at a time when there is less support for nuclear diplomacy than at any point in recent history.

New systems and more weapons

In June 2025, the United Kingdom [announced plans](#) to acquire Lockheed Martin F-35As that are the [first stealth fighter jets](#) certified to carry variable-yield thermonuclear B61-12 gravity bombs. The 12 F-35As would be ordered in place of F-35B fighter jets, which are designed for use on aircraft carriers but are not nuclear-capable. At the Hague summit last year, the United Kingdom then announced it would be joining NATO's air-launched nuclear mission and putting those planes to work carrying US nuclear bombs. Though it has still not been officially confirmed, US B61-12s were [almost certainly moved to RAF Lakenheath](#) in July 2025.¹¹ Media briefings at the time suggested the UK government was also in talks with Washington on acquiring non-strategic B61-12 bombs of its own, but no such agreement has publicly emerged yet.

On the surface, the United Kingdom's decision to acquire tactical nuclear capability appeared to be part of a natural evolution in European strategic planning prompted by Russia's invasion of Ukraine in 2022. But London's acquisition of F-35As and the deployment of US B61-12s are just the latest examples of a longer trend in the United Kingdom's nuclear posture. In 2022, RAF Lakenheath had [already undergone infrastructure modernization](#) necessary for the return of US nuclear weapons.

Before that, in 2021, the UK government decided to raise the maximum cap on the number of warheads in its arsenal by more than 40 percent. In 2010, the UK government stated that it would continue to reduce its warhead stockpile—at the time numbering 225 warheads—and cap its future arsenal at a maximum of [180 warheads](#). But in 2021, it suddenly raised the cap to 260 warheads. Despite being a reversal of decades of

gradual reductions in the United Kingdom's projected nuclear stockpile, that decision was [slipped onto page 76](#) of the government's Integrated Review of Security, Defence, Development, and Foreign Policy.

Policy implications

The United Kingdom has also made a quiet change to its nuclear doctrine, [introducing a "right to review"](#) the use of nuclear weapons against non-nuclear states under certain circumstances. Together with the decision to raise the warhead cap, this constituted a dramatic strategic change for which there was no discernible democratic mandate and no public discussion.

At the time, even the best-informed commentators on UK nuclear affairs were unaware of the motivation. One argument, [put forward by nuclear historian Lawrence Freedman](#), was that a larger stockpile of 260 warheads might allow the United Kingdom to have two fully armed nuclear submarines (each potentially carrying a maximum of 128 warheads on 16 missiles) on patrol at once. That the increase in stockpile size was motivated by plans for the United Kingdom to have the option of acquiring a non-strategic nuclear capability only emerged later.

The decision to station B61-12s on UK soil amounts to a de facto settlement of the long-running debate over adopting a doctrine of [no-first-use](#) or [sole-purpose](#), and over the future of [tactical nuclear weapons in NATO](#). The deployment of high-precision, variable-yield counter-force weapons signals a willingness to carry out retaliatory or pre-emptive nuclear strikes on an adversary.

Due to their enhanced guidance systems, B61-12 bombs, coupled with the dual-capable stealth F-35A, also have the potential to alter the geopolitical balance in Europe. The UK B61-12 adoption paves the way for further proliferation of these weapons throughout Europe, cements the normalization of nuclear weapons, and challenges the guaranteed second-strike survivability of adversaries' nuclear command and control and strategic weapons systems.

Since the end of the Cold War, the United States and, by extension, NATO have become, in the words of the now Under Secretary of Defense for Policy Elbridge Colby, ["accustomed to escalation dominance,"](#) the strategic objective of achieving battlefield superiority on all rungs of escalation. This aim relies on maintaining asymmetrical advantage, which is not a sound doctrine against a nuclear peer competitor: It is [destabilizing](#) and stokes further arms races.

The combination of F-35A fighter jets and B61-12 bombs gives NATO a stealth-enabled, precision counterforce tool designed to deter by denial by holding Russian nuclear and nuclear command and control assets [at credible risk](#). In theory, they provide limited escalation-control options, such as pre-emptive or tit-for-tat strikes at tailored yield. These technological developments make



Russian nuclear forces and command systems both more vulnerable in crisis and less certain for guaranteeing second-strike capability.

From arms control hero to nuclear war fighter

Until recently, UK efforts on exemplary disarmament were comparatively good. The United Kingdom was the only major nuclear state that had limited itself to a single deterrence system (the Trident submarine-launched ballistic missile). The UK government had made incremental reductions in its nuclear stockpile since 1980. With US B61s having returned to the United Kingdom, that record has now been overturned. Given the lack of official transparency, it is worth revisiting the strategic rationale for increasing the deployment of B61-12 across NATO member states and the role non-strategic nuclear weapons are expected to play. During the Cold War, the Soviet Union overmatched the conventional military forces of Western Europe, particularly on land. To offset this conventional weakness, the United States deployed thousands of lower-yield tactical nuclear weapons throughout Western Europe and Turkey. The Soviet Union also maintained non-strategic nuclear options (though, unlike the United States, it formally adopted a 'no-first-use' nuclear doctrine).

But with the collapse of the Soviet Union, the strategic dynamic in Europe reversed.

Now Russia occupies the weaker position in the conventional military balance. The United States [greatly reduced](#) the number of non-strategic nuclear weapons deployed in Europe and eliminated its land- and sea-launched tactical nuclear weapons. Critically, however, it [refused to change](#) any aspect of NATO's nuclear policy, which included the possibility of first use. Russia has since remedied some of its conventional arms deficit, and the current US government's ambivalence about NATO does introduce some uncertainties. But Russian strategic documents consistently express concern about the conventional balance, something nuclear policy expert [Kristin Ven Bruusgaard describes](#) as "a sustained perception of conventional inferiority." As the Centre for Strategic and International Studies [noted in June 2024](#), "NATO's cumulative capabilities far exceed Russia's—even excluding the United States."

American scholars of nuclear warfighting and counterforce strategy, Keir Lieber and Daryl Press, point out that now the military balance in conventional weapons favors the United States and NATO, the tide has turned. They [observe that](#) weaker states, like Russia in Europe, face large incentives to use nuclear weapons in a conventional war against a superior opponent, particularly if they fear battleground losses pose an existential threat to national sovereignty or regime stability.

Ukraine's successful drone strikes against both [Russia's nuclear command and control](#) early warning system and its [strategic bomber fleet](#) have been dramatic demonstrations of Russia's limited capability to provide territorial security against

homeland attacks. And the United States has launched two pre-emptive military operations crafted to initiate regime change in Venezuela and Iran. Russian leaders could perceive themselves as facing a greater existential threat now in a crisis with NATO.

US and allied defense planners should be especially focused on preventing escalation with Russia that might lead to a scenario in which nuclear use would be incentivized and seen as potentially beneficial. But instead, the United Kingdom and the United States are now expanding the deployment of tactical nuclear weapons that are designed for first use and honed for precision and pre-emptive attacks against hardened military targets.

No rationale, real danger

The strategic need for the United Kingdom to operate air-launched tactical nuclear weapons has not been clearly presented either by US officials or the UK government. In addition, these weapons arguably increase the United Kingdom's reliance on US weapons platforms at a time of transatlantic political divergence.

The United Kingdom cancelled its domestic nuclear gravity bomb program [nearly 30 years ago](#) in response to the proliferation of ground-based anti-air systems. US nuclear bombs haven't been deployed in Britain since 2008. The B61-12 program dates to the [2010 Nuclear Policy Review](#), which set out a full-scope life extension of the B61 gravity bombs deployed in Europe in the Cold War. A narrative can be presented that the [F-35A](#) and [B61-12s](#) are well-suited to make NATO's nuclear sharing credible again, given the crisis of the Ukraine war, but the underlying planning dates to the early 2000s.

More seriously, the idea that any nuclear exchange with Russia could be [limited to the sub-strategic level is dubious](#). There can be no guarantee that after initiating a nuclear war with tactical nuclear weapons, or even using them to retaliate against a nuclear strike, it would be possible to stop further moves up the escalatory ladder. Even the deployment of dual-capable F-35A fighter jets that also carry out conventional military missions raises the likelihood of accidental escalation into nuclear war: Adversaries have no way of differentiating whether in-flight F-35As are carrying [conventional or nuclear warheads](#), and they may erroneously perceive an incoming nuclear strike.

The UK government and the United States should carefully revisit the assumptions that led to this point. There is no justifiable reason for the United Kingdom and NATO to expand non-strategic capabilities while also maintaining a no-first-use doctrine unless they are planning a pre-emptive nuclear strike deep into Russian territory, or the territory of another nuclear weapons state, or against a non-nuclear weapons state. NATO member states, like Finland and Poland, are promoting their [participation in nuclear](#)



[sharing](#) and planning. Both understandably perceive an existential threat from Russia, and they are now also considering [hosting US nuclear weapons](#) and [joining France's forward deterrence initiative](#). But with the United Kingdom having already made that choice, it has the dubious distinction

of leading the way in normalizing more first-use tactical nuclear capabilities in Europe without comprehensive analysis of the strategic implications. This nuclear brinkmanship, directly pertinent to the Ukraine war, offers the false promise of security for some at the cost of dooming security for all.

Note

[1] The piece does not argue that the United Kingdom already has its own F-35A fighter jets (which of course have not yet been delivered) “mated” to B61 bombs. But the movement of B61s to RAF Lakenheath is very definitely linked to the decision to purchase F-35As. The UK government formally announced as much last year at the Hague summit. The United States has already deployed B61s in the United Kingdom alongside its own air capability. But the meaning of the United Kingdom “joining NATO’s air-launched nuclear mission” is precisely that UK F-35As, once they arrive and the pilots are certified, would become what NATO calls “certified Allied aircraft” carrying US B61-12 bombs. There is no assumption here. More important, all these developments are very significant changes to the UK nuclear posture.

[Sonja Amadae](#) is the director of the Centre for the Study of Existential Risk at the University of Cambridge and a tenured political scientist at the University of Helsinki.

[Tom Stevenson](#) is a contributing editor at the London Review of Books.

EDITOR’S COMMENT: My humble opinion is that the UK should take care of its land and sea armed forces and then rely on a few nuclear bombs. The Royal Navy has 9 nuclear-powered submarines. This fleet is split between 4 ballistic submarines carrying the nuclear deterrent and 5 fleet submarines. As of April 2025, the combined surface fleet of the Royal Navy and the Royal Fleet Auxiliary (RFA) consisted of 70 vessels. This includes: 2 Queen Elizabeth-class aircraft carriers; 6 Type 45 air-defense destroyers; 7 Type 23 frigates (with only a few being fully operational at any one time due to maintenance and decommissioning); and 13 RFA support vessels. The UK is actively building new Type 26 and Type 31 frigates to replace the aging Type 23 fleet over the next decade. The official figure for the number of Challenger 2 tanks held by the British Army is 288. This number increased significantly in 2025 when 69 vehicles were brought out of storage. The number of “battle-ready” or fully operational tanks is far lower. In 2023, it was reported that only about 157 out of a fleet of roughly 200 were operational. More recent warnings from former senior military officers suggest the number of tanks ready for combat could be as low as 25. Numbers do not lie!

Chinese nuclear weapons, 2026

By Hans M. Kristensen, Matt Korda, Eliana Johns, and Mackenzie Knight-Boyle

Source: <https://thebulletin.org/premium/2026-07/chinese-nuclear-weapons-2026/>

A Julang-2 (JL-2) submarine-launched ballistic missile (SLBM) of the People's Liberation Army Navy breaks above the water during an undated test. On July 6, China conducted a new test launch of what is believed to be a JL-2 SLBM fired from one of its nuclear-powered ballistic missile submarines. The missile fell back into the Pacific Ocean, more than 7,200 kilometers away from the launch point, thereby testing the JL-2's maximum range for the first time. This range is enough to reach Alaska, Guam, Hawaii, Russia, and India from waters near China, but not the continental United States. The Pentagon estimates that China has replaced all of its deployed JL-2 SLBMs with longer-range JL-3s. (Photo by China Military via Wikimedia; modified by François Diaz-Maurin)



July 08 – Within the past decade, China has significantly expanded its ongoing nuclear modernization program by fielding more types and greater numbers of nuclear weapons than ever before. Since our previous edition on China in March 2025, China has continued to develop its three new missile silo fields for solid-fuel intercontinental



ballistic missiles (ICBMs), continued the construction of new silos for its liquid-fuel DF-5 ICBMs, has been developing new variants of ICBMs and advanced strategic delivery systems, and has likely produced excess warheads for these systems once they are deployed. China has also further expanded its dual-capable DF-26 intermediate-range ballistic missile force, which appears to have completely replaced the medium-range DF-21 in the nuclear role. At sea, China has been refitting its Type 094 nuclear-powered ballistic missile submarines (SSBNs) with the longer-range JL-3 submarine-launched ballistic missile, and is building additional Type 094 SSBNs as well as a new class known as Type 096. In addition, China has recently reassigned an operational nuclear mission to some of its H-6 bombers with an air-launched ballistic missile that might have nuclear capability. To what extent this expansion will affect China's nuclear policy and strategy for using nuclear weapons remains to be seen.

To arm this force, we estimate that China has produced a stockpile of approximately 620 nuclear warheads for delivery by land-based ballistic missiles, sea-based ballistic missiles, and bombers (see Table 1).

Table 1. Chinese nuclear forces, 2026.

Type	NATO designation	Number of launchers ^a	Year deployed	Range (kilometers)	Warheads x yield ^b (kilotons)	Warheads
Land-based ballistic missiles						
<i>Intermediate-range</i>						
DF-26	CSS-18	300 ^c	2016	4,000	1 × 425	100 ^d
<i>Subtotal:</i>		300				100
<i>Intercontinental range</i>						
DF-5A	CSS-4 Mod 2	6	1981	12,000	1 × 4.5 MT	6
DF-5B	CSS-4 Mod 3	12	2015	13,000	Up to 5 × 425	60
DF-5C	(CSS-4 Mod 4)	5	(2025)	13,000	1 × "multi-MT"	5
DF-31A/AG	CSS-10 Mod 2 ^e	104	2007/2018	11,200	1 × 425	104
DF-31B	(CSS-10 Mod 3)	320	2023	12,000	1 × 425	150 ^f
DF-41	CSS-20 (mobile)	28	2020	12,000	Up to 3 × 425	84
DF-61	CSS-? (mobile)	?	(2025)	12,000	?	?
<i>Subtotal:</i>		475				409
<i>Total land-based</i>		775				509
Submarine-launched ballistic missiles						
JL-2	CSS-N-14	0 ^g	2016	7,000+	1 × 700	0
JL-3	CSS-N-20	6/72	2022 ^h	9,000+	1 × 425	72
Aircraftⁱ						
H-6N	B-6	20	2020	3,100+	2 × ALBM	20
H-20	?	.	(2030)	?	(ALCM?)	.
Subtotal		867				601
Warheads produced for incoming delivery systems						19 ^j
Total						620^k

Table 1. Chinese nuclear forces, 2026.

The Pentagon reported in 2025 that China's nuclear stockpile "remained in the low 600s through 2024, reflecting a slower rate of production when compared to previous years" (US Department of Defense 2025, 22). But Chinese warheads are not "operational" like the US and Russian nuclear warheads deployed on operational missiles and at bomber bases; nearly all Chinese warheads are thought to be stored separately from the launchers. It is possible, but uncertain, that a small number of land-based missiles during training exercises are brought to full alert status with warheads. Moreover, some US officials have told us that SSBNs on intermittent deterrent patrols are carrying nuclear warheads, but there is no confirmation of this. In any case, President Xi Jinping's purge of military officials—including leaders of the People's Liberation Army Rocket Force—makes it seem unlikely that nuclear warheads would be handed over to the military under normal circumstances.

Overall, we cannot replicate the Pentagon's warhead estimate based on the reported and observable force structure unless we assign warheads to a significant number of China's new silos.

The Pentagon also estimates that China's arsenal will surpass 1,000 warheads by 2030, many of which will probably be "deployed at higher readiness levels" (US Department of Defense 2024, IX). The Pentagon's 2025 report to Congress, like in 2024 and 2023, notably did not include the projection made in previous Department of Defense reports that



China might field a stockpile of about 1,500 nuclear warheads by 2035 (US Department of Defense 2022b, 94, 98).

These projections depend on many uncertain factors, including:

- How many missile silos China will ultimately build;
- How many silos China will load with missiles;
- How many warheads each missile will carry;
- How many DF-26 intermediate-range ballistic missiles will be deployed and how many of them will have a nuclear mission;
- How many ballistic missile submarines China will field and how many warheads each missile will carry;
- How many bombers China will operate and how many weapons each will carry; and
- Assumptions about the future production of fissile materials and the number of warheads China will be able to produce.

The Pentagon warhead projection appears to simply apply the same growth rate of new warheads added to the stockpile annually between 2019 and 2023 to the subsequent years until 2030. Using that same growth rate until 2035 produces the 1,500 warheads the Pentagon previously projected. We assess that this projected growth trajectory is potentially feasible but depends significantly upon answers to the above questions.

Research methodology and confidence

The analyses and estimates made in the Nuclear Notebook are derived from a combination of open sources: (1) state-originating data (e.g. government statements, declassified documents, budgetary information, military parades, and treaty disclosure data); (2) non-state-originating data (e.g. media reports, think tank analyses, and industry publications); and (3) commercial satellite imagery. Because each of these sources provides different and limited information that is subject to varying degrees of uncertainty, we crosscheck each data point by using multiple sources and supplementing them with private conversations with officials whenever possible.

Analyzing and estimating China's nuclear forces is a challenging endeavor, particularly given the relative lack of state-originating data and the tight control of messaging surrounding the country's nuclear arsenal and doctrine. Like most other nuclear-armed states, China has never publicly disclosed the size of its nuclear arsenal or much of the infrastructure that supports it. This degree of relative opacity makes China's nuclear arsenal difficult to quantify, particularly given that it is likely the fastest-growing arsenal in the world. China may become more transparent about its nuclear forces over the coming decade if it deepens its participation in arms control consultations—the first of which took place in November 2023—although building a culture of nuclear transparency from scratch will take time (Gordon 2023). In September 2024, China surprisingly notified the United States, France, and other countries before it test-launched a DF-31B intercontinental ballistic missile from Hainan Island, although it did not notify its nearby neighbors, Japan or the Philippines (US Department of Defense 2025, 23).

Despite blind spots, it is possible to develop a much more comprehensive picture of the Chinese nuclear arsenal today than just a few decades ago by examining videos of China's People's Liberation Army (PLA) military parades, translations of strategic documents, and commercial satellite imagery. The relative degree of structure and standardization within the various PLA services also allows researchers to better understand the structure and mission of missile brigades and individual units.

Additionally, other countries—particularly the United States—regularly produce public assessments or statements about China's nuclear forces. Such statements, however, are rarely coupled with evidence, as doing so could reveal sensitive sources and methods. As a result, they must be verified as they can be institutionally biased and reflect a mind-set of worst-case thinking rather than the most likely scenario.

Analysis produced by think tanks and non-governmental experts can also be highly useful in informing estimates: The transparency surrounding China's missile forces, in particular, has been greatly enhanced in recent years by the unique work of Mark Stokes, Decker Eveleth (Eveleth 2023), Ben Reuter, and the US Air Force's China Aerospace Studies Institute.

It is important to view external analysis with a critical eye, as there is a high risk of citation and confirmation bias, in which governmental or non-governmental reports build on each other's estimates—sometimes without the reader knowing that this is occurring. This practice can inadvertently create a cyclical echo chamber effect, which may not necessarily match the reality on the ground.

In the absence of reliable or official data, commercial satellite imagery has become a particularly important resource for monitoring new developments and analyzing China's nuclear forces. Satellite imagery makes it possible to identify air, missile, and navy bases, as well as potential underground storage facilities. The PLA's standardization has also enabled researchers to better understand developments at China's military bases, as layouts and construction dynamics now increasingly follow certain patterns, designs, and dimensions. Considering all these factors, we maintain a relatively higher degree of confidence in our Chinese nuclear force estimates than in those of some other nuclear-armed countries where official and unofficial information is even more scarce (Pakistan, India, Israel, and North Korea). However, our estimates about Chinese nuclear forces come with relatively more uncertainty than those of countries with greater nuclear transparency (the United States, the United Kingdom, France, and Russia).



Fissile materials production

How much and how fast China's stockpile can grow will depend upon its inventories of plutonium, highly enriched uranium (HEU), and tritium. The International Panel on Fissile Materials assessed in 2023 that China had a stockpile of approximately 14 tonnes (metric tons) of HEU and approximately 2.9 tonnes of separated plutonium in or available for nuclear weapons (Kütt, Mian, and Podvig 2023, 328–329). Those inventories were sufficient to support a doubling of the stockpile over the past five years and potentially an increase to approximately 1,000 warheads by the turn of the decade. However, producing more than 1,000 additional warheads by 2035 would require additional fissile material production. The Pentagon confirmed this assessment in 2024 by saying China “probably will need to begin producing new plutonium this decade to meet the needs of its expanding nuclear stockpile” (US Department of Defense 2024, 107). The Pentagon also assesses that China is expanding and diversifying its capability to produce tritium (108), and China reportedly began in 2023 to operate two large new centrifuge enrichment plants—one at Emeishan and another at Lanzhou (Zhang 2023a, 2023b).

Chinese production of weapon-grade plutonium reportedly ceased in the mid-1980s (Zhang 2018), and the Pentagon stated in 2024 that China “has not produced large quantities of plutonium for its weapons program since the early 1990s” (US Department of Defense 2024, 107). However, Beijing is combining its civilian technology and industrial sector with its defense industrial base to leverage dual-use infrastructure (US Department of Defense 2023, 28). It is believed that China likely intends to acquire significant stocks of plutonium by using its civilian reactors, including two commercial-sized CFR-600 sodium-cooled fast-breeder reactors that are currently under construction at Xiapu in Fujian province (Jones 2021; von Hippel 2021; Zhang 2021a). Rosatom—Russia's state-controlled nuclear energy company—completed the final delivery of fuel to supply the first fuel loading in December 2022 (Rosatom 2022), and steam possibly seen emanating from a cooling tower on satellite imagery in October 2023 suggests the first CFR-600 reactor may have begun operation (Kobayashi 2023). In December 2023, the International Panel on Fissile Materials reported that the first reactor had begun operating at low-power mode in mid-2023 but was likely still undergoing testing as of early 2025 (Zhang 2023a; US Department of Defense 2025, 24). As of the beginning of 2026, it was unclear if the reactor had been connected to the grid and begun generating electricity. The second was scheduled to come online by 2026.

To extract plutonium from its spent nuclear fuel, China completed construction of its first civilian “demonstration” reprocessing plant at the China National Nuclear Corporation (CNNC) Gansu Nuclear Technology Industrial Park in Jinta, Gansu province, which is expected to be operational in 2025 (Zhang 2024). China has started the construction of a second plant to reprocess spent light-water reactor fuel at the same location, which is expected to be up and running before the end of the decade (Zhang 2021b, 2024). Construction of a third plant in a new, extended area of the park began in 2023 and is expected to be completed in the early 2030s (Zhang 2024). The MOX plant and fuel reprocessing capacity at Jinta and the 50 tonne-per-year capacity at Jiuquan (Plant 404) could support the two CFR-600 reactors, which will begin operation with highly enriched uranium (HEU) rather than mixed oxide (MOX) fuel through a supply agreement with Russia (US Department of Defense 2023, 109; Zhang 2021b). The ambiguity of Chinese nuclear warhead types and uncertainty on the exact amount of fissile material required for each warhead design make it difficult to estimate how many weapons China could produce from its existing HEU and weapons-grade plutonium stockpiles. If both fast-breeder reactors operate as planned (despite other countries having serious difficulties operating fast-breeder reactors), they could potentially produce large amounts of plutonium and, by some estimates, could enable China to acquire over 330 kilograms of weapon-grade plutonium annually for new warhead production (Kobayashi 2023)—which would be consistent with the Pentagon's most recent projections (US Department of Defense 2024, 107). China, however, insists that its CFR-600 reactors are for civilian use only, and some experts have pointed out that fast-breeder reactors are an extremely inefficient way of producing weapons-grade plutonium (Park 2024).

Although China's production and reprocessing of fissile materials could potentially be consistent with its nuclear power efforts and its goal of reaching a closed nuclear fuel cycle, the Pentagon asserts that Beijing intends to use this infrastructure “to produce nuclear warhead materials for its military in the near term” and “reestablish China's ability to produce weapons-grade plutonium” (US Department of Defense 2024, 107; 2025, 24). Satellite imagery also indicates improved and expanded infrastructure at the Pingtong Nuclear Facility—believed to be involved in manufacturing plutonium pits—as well as the Zitong Nuclear Weapons Complex—involved in the fabrication and assembly of weapons components (Babiarz 2025a, 2025b).

The degree of transparency surrounding China's nuclear materials production and its suspected expansion of uranium and tritium production has decreased since China has not reported its separated plutonium stockpile to the International Atomic Energy Agency since 2017 (US Department of Defense 2024, 108).

US estimates and assumptions about Chinese nuclear forces

Evaluation of current US projections about the future size of China's nuclear weapons stockpile must take earlier projections into account, some of which did not come to pass. During the 1980s and 1990s, US government agencies published several projections for the number of Chinese nuclear warheads. A US Defense Intelligence Agency study from 1984 inaccurately estimated that China had 150 to 360 nuclear warheads and projected it



could increase to more than 800 by 1994 (Kristensen 2006). Over a decade later, another Defense Intelligence Agency study published in 1999 projected that China might have over 460 nuclear weapons by 2020 (US Defense Intelligence Agency 1999). While this latter projection ultimately proved to be closer to the warhead estimate the Pentagon published in 2020, it was still more than twice the “low-200s” warhead estimate announced by the Pentagon (US Department of Defense 2020, ix).

Current US projections also come with significant uncertainties. In November 2021, the Pentagon’s annual China Military Power Report (CMPR) to Congress projected that China could have 700 deliverable warheads by 2027, and as many as 1,000 by 2030 (US Department of Defense 2021, 90). The 2022 Pentagon report increased the projection even further, claiming that China’s stockpile of “operational” nuclear warheads had surpassed 400 and might reach about 1,500 warheads by 2035 (US Department of Defense 2022b, 94). The estimate increased to more than 500 warheads in 2023 and to over 600 warheads in 2024 and repeated the projection that China might possess over 1,000 operational warheads by 2030 (US Department of Defense 2024, IX). The 2025 Pentagon report stated that China’s stockpile “remained in the low 600s through 2025, reflecting a slower rate of production when compared to previous years.” Still, it maintained that the PLA is “on track” to have “over 1,000 warheads by 2030” (US Department of Defense 2025, 22). However, the available public information and observable force structure do not allow us to replicate the estimate of more than 600 warheads unless we assign a significant number of nuclear warheads—up to 160—to the new silos. As a result, we assess that China’s stockpile may include approximately 620 warheads, but this estimate must include a significant number of warheads that have been produced for missiles that are still being fielded.

Chinese officials have pushed back against what they see as the “sensationalized” or “exaggerated” claims and “false narratives” made in successive Pentagon CMPRs (Li 2022; Ministry of National Defense of the People’s Republic of China 2023; *China Daily* 2024; Yin 2024). However, the Chinese government has not denied—and has barely acknowledged—the expansion of the mobile ICBM force or the construction of three large new missile silo fields.

The Pentagon’s projected increase has unsurprisingly triggered a wide range of speculations about China’s nuclear intentions. Over the past five years, high-ranking US officials—including the former US STRATCOM commander—have suggested that China has moved away from its longstanding “minimum deterrence” posture, and that it “seeks to match, or in some areas surpass, quantitative and qualitative parity with the United States in terms of nuclear weapons” (Billingslea 2020; Bussiere 2021; US Strategic Command 2022; Cotton 2023).

US officials and individuals in the public debate have repeatedly claimed that China’s expansion will make it a nuclear “peer” or “near peer” in the future. However, there is no evidence that China’s ongoing nuclear expansion will result in parity with the US arsenal. Even the highest projection from 2023 of 1,500 warheads by 2035 amounts to less than half of the current US nuclear stockpile. When reminded about that discrepancy, some US defense officials have sought to downplay the importance of numbers: “We don’t approach it from purely a numbers game,” according to then-deputy commander of the US Strategic Command, Lt. Gen. Thomas Bussiere. “It is what is operationally fielded, ... status of forces, posture of those fielded forces. So, it is not just a stockpile number,” he said (Bussiere 2021).

Nuclear testing

The projection for how much the Chinese nuclear stockpile will increase also depends on the size and design of its warheads. China’s nuclear testing program of the 1990s partially supported the development of the warhead (#535) type currently arming the DF-31-class ICBMs. This warhead design may also have been used to equip the liquid-fueled DF-5B ICBM with multiple independently targeted reentry vehicle (MIRV) technology, replacing the much larger multi-megaton warhead (#506) used on the DF-5A. The large DF-41 and the JL-3 missiles could potentially use the same smaller #535 warhead or the even smaller #5 × 5 warhead (Zhang 2025). The Pentagon reported in 2024 that China was probably seeking a “lower-yield” nuclear warhead for the DF-26 (US Department of Defense 2024, 124); however, it was unclear if that implied the production of a new warhead and how low such a “lower” yield might be. The Pentagon specified further in its 2025 report, stating that China is “probably pursuing nuclear weapons with yields below 10 kilotons,” and that the “highly precise theater weapons” DF-26 and the H-6N’s air-launched ballistic missile potentially “would be well suited for delivering a low-yield nuclear weapon” (US Department of Defense 2025, 24). China is thought to already have a lower-yield warhead, the #5 × 5 warhead (Zhang 2025). For instance, if Beijing wanted a low-yield warhead, it could potentially do so by using existing warheads and “turning off” the uranium secondary so only smaller-yield plutonium primary is used, similarly to what the United States did with its W76-2 warhead more than five years ago.

Developing significantly different warhead designs would probably require additional nuclear test explosions. To avoid such tests, China could potentially make simpler designs that use a previously tested nuclear explosive package, advanced computer simulations, and sub-critical (or very low-yield) underground explosive experiments. In recent years, the United States has claimed that some of China’s actions at Lop Nur “raise concern” about its adherence to the United States’ “zero-yield” standard, but had not explicitly accused China of conducting critical tests that produced a yield (US Department of State 2022, 29; 2023, 18; 2024). In a February 2026 speech at a Conference on Disarmament meeting in Geneva, however, the US undersecretary of state for arms control and international security, Thomas DiNanno,



accused China of conducting a “yield producing nuclear test” on June 22, 2020, and of “preparing for tests with designated yields in the hundreds of tons” (US Mission to International Organizations in Geneva 2026). China’s ambassador to the United Nations called the accusation “completely groundless” (Chinese Mission to the United Nations in Vienna 2026). The Comprehensive Nuclear Test Ban Treaty Organization confirmed in a February 17 statement that the International Monitoring System did indeed detect two small seismic events near China’s Lop Nur test site on the June 2020 date, but that “it is not possible to assess the cause of these events with confidence” due to the detected yield of the events being far lower than the 500-ton threshold consistent with nuclear test explosions (Comprehensive Nuclear Test Ban Treaty Organization 2026). NORSAR, a Norwegian seismic activity monitoring entity, additionally reported that “the data do not rule out that the event could have been a small natural earthquake” (NORSAR 2026). US officials have not provided further evidence of the claim. While the truth about whether China conducted an explosive nuclear test in June 2020 remains uncertain, there have been high levels of activity observed at China’s Lop Nur site in recent years, leading some analysts to conclude that China could be preparing to conduct underground nuclear testing (Babiarz 2023).

Analysis of commercial satellite imagery shows significant construction at the Lop Nur site, including approximately a dozen new concrete buildings near the site’s airfield, as well as at least one new tunnel at the site’s northern testing area (Brumfiel 2021b). The imagery shows what appears to be new drainage areas, drill rigs, roads, spoil piles, and covered entrances to potential underground facilities, as well as new construction at the main administration, support, and storage areas (Babiarz 2023; Brumfiel 2021a; J. Lewis 2023). Many of these activities appeared to still be ongoing as of the time of writing this report. In addition to new activity at the northern tunnel test area, satellite imagery also indicated activity at a possible new eastern test area at Lop Nur (Babiarz 2023). However, although the construction works are significant, they do not prove that China plans to conduct new nuclear explosive tests at the site. Should China conduct low-yield nuclear tests at Lop Nur, it would violate its responsibility under the Comprehensive Nuclear Test Ban Treaty it has signed but not ratified.

Nuclear doctrine and policy

Since its first nuclear test in 1964, China has maintained a consistent narrative about the purpose of its nuclear weapons and, in its words, upheld the 1980 consensus that “a nuclear war cannot be won and must not be fought” (Information Office of the State Council 2025). This narrative was reaffirmed in a white paper released by the Ministry of Foreign Affairs in late 2025 and maintained in China’s national defense policy, which states:

China is always committed to a nuclear policy of no first use of nuclear weapons at any time and under any circumstances, and not using or threatening to use nuclear weapons against non-nuclear-weapon states or nuclear-weapon-free zones unconditionally. ... China does not engage in any nuclear arms race with any other country and keeps its nuclear capabilities at the minimum level required for national security. China pursues a nuclear strategy of self-defense, the goal of which is to maintain national strategic security by deterring other countries from using or threatening to use nuclear weapons against China. (Ministry of National Defense of the People’s Republic of China 2025)

PLA doctrine has described China’s nuclear forces as one of several deterrence capabilities for ensuring “global strategic balance and stability” (Communist Party of Central China Committee 2025). The role of nuclear weapons for “strategic counterbalance” appears to suggest that China seeks to offset a military disadvantage against the United States and could mean it believes nuclear weapons play a role in shaping the geostrategic balance of global power (US Department of Defense 2025, 10; Zhao 2024).

Despite its declaratory policy of no-first-use that emphasizes a “defensive” posture, China has never defined how large a “minimum” capability is or what activities constitute an “arms race.” The stated policies evidently do not prohibit the unprecedented expansion of China’s nuclear arsenal currently underway. The posture apparently seeks to “adapt to the development of the world’s strategic situation,” part of which involves the “organic integration [of] nuclear counterattack capability and conventional strike capability” (China Aerospace Studies Institute 2022, 381–382).

Such capabilities require investing significant resources to ensure the survivability of the nuclear arsenal against a nuclear or conventional first strike, including practicing “nuclear attack survival exercises” to ensure that troops could still launch nuclear counterattacks if China were to be attacked (Global Times 2020). It also involves improving early-warning systems and the stealth capabilities of its nuclear forces to be able to elude enemy detection (Kaufman and Waidelich 2023, 42, 45).

The PLA maintains what it refers to as a “moderate” readiness level for its nuclear forces and keeps most of its warheads at its regional storage facilities and its central hardened storage facility in the Qinling mountain range.^[1] The 2024 Pentagon report reaffirmed this posture, stating that China maintains “a portion of its units on a heightened state of readiness while leaving the other portion in peacetime status with separated launchers, missiles, and warheads.” But the report also described that the People’s Liberation Army Rocket Force (PLARF) brigades conduct “combat readiness duty” and “high alert duty” drills, which “includes assigning a missile battalion to be ready to rapidly launch” (US Department of Defense 2024, 106).

The readiness of the Chinese nuclear missile force was challenged in early 2024 with the disclosure that a US intelligence assessment had found that corruption within the PLA had led to an erosion of confidence in its overall capabilities, particularly when it comes to the Rocket Force (Martin and Jacobs[2] 2024). The 2024 Pentagon



report noted that the subsequent corruption investigation “likely resulted in the PLARF repairing the silos, which would have increased the overall operational readiness of its silo-based force” (US Department of Defense 2024, 159). This could explain the heightened activity seen through satellite imagery at several new Chinese silos between 2023 and 2024.

Increased readiness and alert drills do not necessarily require nuclear warheads to be installed on the missiles or prove that they are installed at all times, but it cannot be ruled out either. However, the dismissals of top defense officials and widespread corruption in the past few years might degrade the Chinese leadership’s willingness to arm missiles with warheads in peacetime.

A nuclear attack against China is unlikely to come out of the blue and is more likely to follow a period of increasing tension and possibly conventional warfare, allowing the warheads to be mated to the missiles in time. Moreover, a credible retaliatory capability doesn’t require missiles to be on alert in peacetime: With launchers inside tunnels dispersed across vast mountainous areas, it is impossible—even for a large nuclear adversary—to prevent China from retaliating with at least some of its missiles intact. In its national report submitted to the 11th NPT Review Conference in April 2026, the Chinese government reiterated a generic description of its alert posture and the stages Chinese nuclear forces would go through in a crisis:

The command of nuclear force in China is highly centralized and the command process is lean and efficient. It ensures the strictest and most accurate execution of orders of the Central Military Commission. Chinese nuclear forces maintain a moderate level of alert under normal circumstances. When the country is under a nuclear threat, the alert status is raised upon the order of the Central Military Commission to prepare for nuclear counterstrikes, thereby deterring the enemy from using nuclear weapons against China. In the event that the country suffers a nuclear attack, a resolute counterstrike will be mounted against the enemy. (People’s Republic of China 2026)

In peacetime, the “moderate level of alert” might involve designated units to be deployed in high combat-ready condition with nuclear warheads installed, or in nearby storage sites under the control of the Central Military Commission that could be released to the unit quickly if necessary. Analysts have described Chinese nuclear forces as following a six-stage alert sequence based upon increasingly actionable intelligence (J. W. Lewis and Xue 2012; Wood, Stone, and Corbett 2024):

1. Standing war preparedness alert: day-to-day readiness condition.
2. Class 3 operational preparations alert: missile bases accelerate preparations for launching missiles and base security is upgraded.
3. Class 2 operational preparations alert: missile bases, including associated air defense and ground crews, shift to maximum readiness.
4. Class 1 operational preparations alert: gives authority to designated base commanders to launch nuclear weapons upon receipt of formal orders from the Central Military Commission.
5. Preparatory order: includes precise timing and instructions for mobile launch units to enter launch sites and for siloed units to conduct necessary pre-launch activities.
6. Formal order: official launch order from the Central Military Commission authorizing the use of nuclear weapons.

China is building several underground facilities at some of its newer sites, including at its three solid-fuel missile silo fields, which could potentially be used for warhead storage. Each PLARF regional base has a dedicated “equipment inspection” unit that is responsible for storage, management, and transportation of the nuclear weapons assigned to missile brigades in that base area. These units will coordinate the “transfer” of warheads with base “logistical support” regiments, which possibly handle the shipment of warheads to missile brigades within the base using both rail and road transport capabilities (Bruzzeze 2026, 23). While it is unknown if they are armed with nuclear warheads, the 2025 Pentagon report stated that the PLA has likely loaded “more than 100” missiles at its three silo fields. It is unclear how many warheads are kept on-site at the base or brigade level and how many would be mated with the launchers in case of a launch order. Warhead mating at the three silo fields likely takes place “on-site by brigade personnel” and at the base-level for other brigades (Bruzzeze 2026, 25).

The Pentagon assesses that the construction of hundreds of new silos for quick-launch solid fuel missiles and development of a space-based, ground-based, and airborne early warning systems indicate China’s intent to move to a launch-on-warning (LOW) posture, referred to as an “early warning counterstrike” (预警反击), giving China time to launch its missiles before they would be destroyed (US Department of Defense 2024, 110; 2025, 15, 23–24). The Pentagon says that China has launched two additional early warning satellites for a likely total of five as of early 2025, and that the PLARF continues to conduct exercises involving “early warning of a nuclear strike and LOW responses” (US Department of Defense 2024, 110; 2025, 23–24).

In addition to the technical means for protecting the missiles against a first strike, the PLARF has also emphasized “survival protection” for its land-based nuclear forces (China Aerospace Studies Institute 2022, 386). This involves training soldiers to perform additional tasks beyond their primary roles, including a “role switch” where a transporter erector launcher (TEL) driver would also know how to launch a missile, or a measurement specialist who knows how to command (Baughman 2022).

During one “survival protection” training exercise in November 2021, a launch battalion was informed they would be “killed” by an enemy missile strike in five minutes. Rather than attempting to evacuate—the standard “survival protection” procedure—the battalion commander ordered his troops to conduct a surprise “launch on the spot” of



their ballistic missile before the enemy missile hit their position (Baughman 2022; Lu and Liu 2021). While the report did not specify whether the battalion had a nuclear or conventional strike role, the exercise suggests that the PLARF is practicing launching missiles in a launch-on-warning scenario.

These data points, however, are not necessarily evidence of a formal shift to a more aggressive nuclear posture (Fravel, Hiim, and Trøan 2023). They could just as likely be intended to allow China to disperse its forces and, if needed, launch rapidly—but not necessarily first or “on warning”—in the context of a crisis, thereby safeguarding its forces against a surprise conventional or nuclear first strike. For decades, China has deployed silo-based DF-5s and road-mobile ICBMs that, in a crisis, would be armed with the intention to launch them before they are destroyed. Although there has been considerable discussion in China about the size and readiness of the nuclear arsenal as well as when the no-first-use policy would apply, there is little evidence to suggest that the Chinese government has deviated from it, which is also reiterated in its 2023 national defense strategy (Ministry of National Defense of the People’s Republic of China 2023b; Santoro and Gromoll 2020). China potentially could maintain its current strategy even with many new silos and improved early-warning systems.

Notably, both the United States and Russia operate large numbers of solid-fuel silo-based missiles and early-warning systems to be able to detect nuclear attacks and launch their missiles before they are destroyed. The two countries also insist that such a posture is both necessary and stabilizing. It seems reasonable to assume that China would seek a similar posture to safeguard its own retaliatory capability. A Chinese early-warning system could potentially also be intended to enable a future advanced missile defense system.

As with other nuclear-armed states, there is uncertainty and ambiguity about what circumstances could cause the Chinese leadership to order the use of nuclear weapons. Despite its no-first-use policy, Chinese officials have privately stated in the past that China reserves the right to use nuclear weapons if its nuclear forces were attacked with conventional weapons. The Pentagon suggested China could seek to “deter non-nuclear military actions with its nuclear forces” and could consider a nuclear strike in response to a “non-nuclear attack threatening the viability of China’s nuclear forces or C2, or that approximates the strategic effects of a nuclear strike” (US Department of Defense 2025, 25; 2024, 102).

The modernization of the nuclear forces could potentially influence Chinese nuclear strategy and declaratory policy gradually in the future by offering the Chinese leadership more expansive and efficient ways of deploying, responding, and coercing with nuclear or dual-capable forces. The 2022 US Nuclear Posture Review suggested that China’s trajectory of expanding and improving its nuclear arsenal could “provide [China] with new options before and during a crisis or conflict to leverage nuclear weapons for coercive purposes, including military provocations against US Allies and partners in the region” (US Department of Defense 2022a, 4). Conversely, a larger and more capable nuclear arsenal could also make the Chinese leadership less concerned about the survivability of its force and therefore less likely to overreact.

Regardless of what the specific red lines may be, China’s no-first-use policy probably has a high threshold. The significant modernization of non-nuclear forces seems to indicate that the Chinese leadership is interested in keeping it that way. Many experts believe there are very few scenarios in which China would benefit strategically from a first nuclear strike even in the case of conventional conflict with a military power such as the United States (Tellis 2022, 27). However, in a situation in which the stakes would be high—such as in a military clash over Taiwan—both China and the United States appear to reserve the option of using nuclear weapons—including first—if deemed necessary. In the Pentagon’s words: “Beijing probably would consider nuclear first use if a conventional military defeat in Taiwan gravely threatened [the Chinese Communist Party] regime survival” (US Department of Defense 2024, 102).

Land-based ballistic missiles

China continues the long-term modernization of its land-based, nuclear-capable missile force. But the pace and scope of this effort have increased significantly over the past few years with the construction of approximately 350 new missile silos and several new bases for road-mobile missile launchers. Overall, we estimate that the PLARF currently operates approximately 775 launchers for land-based missiles that can deliver nuclear warheads. But not all of those are necessarily assigned nuclear warheads. Of those launchers, 475 can be loaded with missiles that can reach the continental United States. Many of China’s ballistic missile launchers are for short-, medium-, and intermediate-range missiles intended for regional missions, and most of those do not have nuclear strike missions. We estimate that China has roughly 100 nuclear warheads assigned to regional missiles, although this number comes with significant uncertainty.

The PLARF, which is headquartered in Beijing, has recently undergone several management shakeups: In July 2023, the PLARF commander and political commissar, along with several other senior officers, were removed from their positions following an anti-corruption investigation. Notably, the top two PLARF officials were replaced by generals from outside the PLARF itself: the new commander and political commissar come from the People’s Liberation Army Navy (PLAN) and the People’s Liberation Army Air Force (PLAAF), respectively (London, McCarthy, and Chang 2023). In October



2023, China's Minister of National Defense Li Shangfu, who would have been responsible for approving nuclear weapons acquisitions, was also removed from his post (US Department of Defense 2024, XII).

The PLARF controls nine individually numbered bases: six for missile operations distributed across China (Bases 61 through 66), one for overseeing the central nuclear stockpile (Base 67), one for maintaining infrastructure (Base 68), and one that is assumed to be for training and missile tests (Base 69) (Xiu 2022, 2). Each missile operating base controls six to eight missile brigades, with the number of launchers and missiles assigned to each brigade depending on the type of missile (5).

In November 2025, CNN revealed that China was undertaking a significant expansion and overhaul of its missile production industry. According to the report, missile production-related sites and other sites connected to the PLARF expanded by 21 million square feet since early 2020 (Qiblawi et al. 2025).

To accommodate the growing missile force, the total number of Chinese missile brigades has increased, too. This increase is predominantly caused by the growing inventory of conventional missiles, but it is also a product of China's nuclear modernization program. We estimate that the PLARF currently has approximately 45 brigades with ballistic or cruise missile launchers. Approximately 34 of those brigades either operate ballistic missile launchers with nuclear capability or are upgrading to do so soon (see Table 2). This is close to the 50 nuclear missile brigades that Russia operates—known as regiments in the Russian military vocabulary (Kristensen, Korda, and Reynolds 2023).

Table 2. Chinese land-based missile brigades, 2026.^a

Base Number (Provinces)	Unit	Location ^b	Weapon Type ^c	Nuclear role	Notes
PLARF HQ		Beijing (40.0352°N, 116.3197°E)			
Base 61	HQ	Huangshan (29.6956°N, 118.2997°E)			
(Anhui, Fujian, Guangdong, Jiangxi, Zhejiang)	611 Brigade	Qingyang (30.6903°N, 117.9011°E)	DF-26	Yes	Major upgrade underway.
	612 Brigade	Leping (28.9797°N, 117.1205°E)	DF-31AG or DF-27? ^d	Yes	New TEL seen in 2025.
	613 Brigade	Shangrao (28.4745°N, 117.8954°E)	DF-15B (DF-17?) ^e	No	Possibly upgrading to new missile.
	614 Brigade	Yongan (26.0596°N, 117.3151°E)	DF-17 ^f	No	First DF-17 brigade.
	615 Brigade	Meizhou (24.2828°N, 115.9708°E)	DF-11A ^g	No	
	616 Brigade	Ganzhou (25.8992°N, 114.9587°E)	DF-17 ^h	No	New base added since 2020. ⁱ
	617 Brigade	Jinhua (29.1508°N, 119.6153°E)	DF-16A ^j	No	Second DF-16 brigade; integration base?
	618 Brigade	Nanchang (28.5004°N, 115.9214°E)?	DF-26	Yes	Probably upgraded to DF-26. ^k
Base 62 ^l	HQ	Kunming (24.9888°N, 102.8346°E)			Base expansion underway.
(Guangxi, Guangdong, Hainan, Sichuan, Yunnan)	621 Brigade	Yibin (28.7607°N, 104.7914°E)	DF-31AG? ^m	Yes	Upgraded from DF-21A.
	622 Brigade	Yuxi (24.3601°N, 102.4942°E)	DF-31A	Yes	Potentially to upgrade to DF-31AG.
	623 Brigade	Liuzhou (24.3856°N, 109.5726°E)	DF-10A	No	First DF-10A brigade.
	624 Brigade	Danzhou (19.4721°N, 109.4570°E)	DF-21D	No	Possibly upgrading to new missile.
	625 Brigade	Jianshui (23.7354°N, 102.8713°E)	DF-26	Yes	Possibly second DF-26 brigade.
	626 Brigade	Qingyuan (23.6845°N, 113.1768°E)	DF-26 ⁿ	Yes	Possible third DF-26 brigade.
	627 Brigade	Puning (23.4122°N, 116.1816°E)	DF-17 ^o	No	Base expansion underway.
	HQ	Huaihua (27.5747°N, 110.0250°E)			
Base 63	631 Brigade	Jingzhou (26.5783°N, 109.6703°E)	DF-5B (DF-5C?)	Yes	6 silos, adding 6 more plus training. ^p
(Huaihua, Hubei, Hunan)	632 Brigade	Shaoyang (27.2532°N, 111.3859°E)	DF-31AG	Yes	Upgrades underway.
	633 Brigade	Huitong (26.8935°N, 109.7388°E)	DF-5A	Yes	6 silos. ^q
	634 Brigade	Yueyang (29.3321°N, 113.2150°E) ^r	(DF-5C?)	(Yes)	New 12-silo field under construction.
	635 Brigade	Yichun (27.8869°N, 114.3862°E)	DF-17?	No	Previously with DF-10A.
	636 Brigade	Shaoguan (24.7579°N, 113.6797°E)	DF-16A	No	First DF-16A brigade.

Table 2. Chinese land-based missile brigades, 2026.

Intercontinental ballistic missiles

Of China's 475 ICBM launchers, we estimate that over 300 may have been assigned missiles that can deliver over 400 warheads. This allegedly includes over 100 missiles loaded in three new silo fields in northern China, although it remains to be seen how many of the remaining silos will be loaded. These 320 new silos for solid-fuel missiles and the construction of 30 new silos for liquid-fuel missiles in three mountainous areas of central-eastern China constitute the most significant recent development in China's nuclear arsenal (Eveleth 2023; Korda and Kristensen 2021; Lee 2021; J. Lewis and Eveleth 2021; Reuter 2023).




China's new silo fields

Since March 2020, China has been constructing three large missile silo fields in its northern desert areas: The Yumen field, located in Gansu province in the western military district, covers an area of approximately 1,110 square kilometers and includes 120 individual silos; the Hami field, located in Eastern Xinjiang in the western military district, spans an area of approximately 1,030 square kilometers, and includes 110 missile silos; and the Yulin field, located near Hanggin Banner west of Ordos, measures roughly 830 square kilometers and includes 90 missile silos. All three fields include support facilities, communications and (or) surveillance towers, and potential air defense sites.

The silo fields at Yumen and Hami bear distinct similarities: They both have perimeters that span the entire complex; the individual silos are positioned roughly three kilometers apart in an almost perfect triangular grid pattern; they have many support facilities in common; and they both have at least one underground installation. In contrast, the silos at the Yulin field have individual perimeters rather than one larger perimeter around the field; the silos are positioned more randomly, although still three kilometers apart; and we have not seen any nearby underground facilities. Only the Hami field includes a rail transfer facility.

Overall, the silo fields are located deeper inside China than nearly all other known ICBM bases, and beyond the reach of the United States' conventional cruise missiles.



UNDERGROUND CONSTRUCTION AT YUMEN MISSILE SILO FIELD, CHINA

In late 2024, China began construction of an adit into the side of a mountain range located within the boundaries of the Yumen intercontinental missile silo field, very close to the likely brigade headquarters. A second adit was added in spring 2025. It is possible that this area, once complete, could be used for munitions or warhead storage.

Satellite imagery @ 2026 Airbus

FAS FEDERATION OF AMERICAN SCIENTISTS

Figure 1. Underground construction at Yumen missile silo field, China. (Images: Airbus; Annotations: Federation of American Scientists)

For some time between April and May 2024, we observed more than 30 silos at Yumen being covered with a camouflage structure to conceal possible maintenance on the silos. (It is possible that the remaining silos were covered on days that were not visible through commercial satellite imagery.) These structures were not tall enough to allow for potential loading operations. We assess that this work is likely related to US intelligence reports from early 2024 that pervasive corruption inside China's rocket forces and industrial base had resulted in silo "lids that don't function in a way that would allow the missiles to launch effectively" (Martin and Jacobs 2024; US Department of Defense 2025, 29). To conduct maintenance



on the silo lids—or to replace the lids entirely—the silos would have to be left open for extended periods of time, thus necessitating the environmental covers that were visible at dozens of silos using satellite imagery during this period.

The Pentagon's 2025 CMPR states that China "has likely loaded more than 100 solid-propellant ICBM missile silos at its three silo fields with DF-31 class ICBMs" (US Department of Defense 2025, 24). However, a thorough investigation using different types of imagery from multiple commercial satellite imagery companies has not revealed any widespread loading operations across any of the three silo fields. This does not necessarily mean that such operations have not taken place, but rather that they have not been observed through commercially available means.

Even if China's newer silos are being loaded with missiles, it remains unclear whether these missiles would carry warheads during peacetime, or whether the upload of warheads to the missiles would constitute an important signaling element for a heightened nuclear alert. In either case, it would be logical for each silo field to maintain localized nuclear warhead storage facilities. Given the harsh desert conditions of the three silo fields—which can be prone to flooding, extreme heat, dust storms, or other unpredictable weather systems—any on-site nuclear storage facilities would need to be climate-controlled, possibly necessitating underground storage.

The silo fields at both Yumen and Hami include underground installations: one at Yumen and two at Hami.

Since early 2025, tunneling operations have been underway at a mountainous section located between the northern and southern halves of the Yumen silo field, next to an installation that we assess to be the headquarters for operations at the Yumen field. As of March 2026, two adits had been constructed with considerable work still ongoing (see Figure 1).

The underground complexes at Hami are more distributed. The one on the western side of the field, which includes three adits, was completed in late 2022, and the one on the eastern side, which includes five adits located within a wider area, was completed in 2023.

We have not yet located any underground installations at the Yulin silo field, although the area surrounding that particular field is considerably less mountainous than at the other two. It is possible that such an installation, if it exists, is located further away.

China's ICBM force structure

In total, these discoveries suggest that China is constructing 320 new silos for solid-fueled ICBMs across the three fields of Yumen, Hami, and Yulin, excluding the approximately 15 training silos at the Jilantai site. In addition, China is upgrading and expanding the number of silos for the liquid-fueled DF-5 ICBM and increasing the number of silos per brigade (US Department of Defense 2023, 107). This appears to include doubling the number of silos of at least two existing DF-5 brigades from six to 12 and adding two new brigades each with 12 silos. Once completed, based on current observations, this project will increase the number of DF-5 silos from 18 to 48.

Combined, these construction efforts for silo-based ICBMs (in addition to new road-mobile ICBM bases) constitute the largest expansion of the Chinese nuclear arsenal ever. The 350 new Chinese silos exceed the number of silo-based ICBMs operated by Russia and constitute about three-quarters of the entire US ICBM force.

In addition to the construction of new ICBM facilities, there is uncertainty about how many ICBMs China currently operates. The US Department of Defense's (2025) report about China's military and security developments assessed that China had 550 ICBM launchers (silo and mobile) with 400 missiles in its inventory (US Department of Defense 2024, 66), bumping its number of launchers up by 50 from the previous year while keeping its number of missiles steady (US Department of Defense 2025, 71). The difference between 550 launchers and 400 missiles indicates that China may still be producing missiles for the new launchers. Analysis of satellite imagery shows that construction is still ongoing at all three silo fields, suggesting that they may still be some years away from full operational capability.

If each silo in the three new silo fields is loaded with a single-warhead DF-31-class ICBM, the total number of warheads in China's ICBM force could eventually reach approximately 600 warheads, up from just over 400 today. The sheer number of silos will likely have a significant effect on US strike plans against China because the US targeting strategy has been typically focused on holding nuclear and other military targets at risk, although this strategy does not mean that all silos must necessarily be held at risk at the same time.

At this stage, it is unclear how these hundreds of new silos will alter the existing brigade structure for China's missile forces. Presently, each of China's ICBM missile brigades is responsible for six to 12 launchers, and it might be expected that each new missile silo field would be organized as a single brigade. However, some analysts have hypothesized that the three new silo fields could lead to the creation of entirely new PLARF "Bases" (each with several brigades)—an extremely rare event that has not taken place in more than 50 years (Xiu 2022, 255). For now, the Pentagon's 2025 report on China lists the Hami and Yumen missile silo fields as "Missile Brigades" in the Western Theater organized under Base 64, and the Yulin missile silo field as a "Missile Brigade" in the Northern Theater organized under Base 65 (US Department of Defense 2025, 77, 87, 91).

Although China has deployed ICBMs in silos since the early 1980s, building missile silos on this scale is a significant shift in China's nuclear posture. The decision to do so has probably not been caused by a single event or issue but, rather, by a combination of strategic and operational objectives, including protecting the retaliatory



capability against a first strike, overcoming the potential effects of adversarial missile defenses, better balancing the ICBM force between mobile and silo-based missiles, increasing China's nuclear readiness and overall nuclear strike capability to account for improvements in the Russian, Indian, and US nuclear arsenals, elevating China to a world-class military power, as well as national prestige.

Currently two versions of the DF-5 are deployed: the DF-5A (CSS-4 Mod 2) and the MIRVed DF-5B (CSS-4 Mod 3). Since 2020, the Pentagon's annual reports to Congress have noted that the DF-5B can carry up to five MIRVs (US Department of Defense 2020, 56), and that this version is likely being upgraded (US Department of Defense 2024, 103). We estimate that two-thirds of the DF-5s are currently equipped to carry MIRVs. In its 2024 annual report, the Pentagon indicated that a third modification with a "multi-megaton yield" warhead—known as the DF-5C (CSS-4 Mod 4)—is currently being fielded and suggests that at least two brigades will likely field the DF-5C (103). It is possible that the DF-5C will replace the DF-5A, which is already equipped with a multi-megaton yield warhead. There was no update on the DF-5C's status in the 2025 CMPR, indicating that the Pentagon's assessment had not changed from the previous year. However, satellite images show that a handful of the new DF-5 silos may have been completed. If so, it is potentially possible that they have been loaded with the new DF-5C.

In 2006, China debuted its first solid-fuel road-mobile ICBM—the DF-31 (CSS-10 Mod 1)—which had a range of about 7,200 kilometers, meaning that it could not reach the continental United States from its deployment areas in China.^[2] Since then, China has iterated on its original DF-31 design, producing newer versions of the system: The extended-range DF-31A and DF-31AG (both of which are designated by the United States as CSS-10 Mod 2, indicating that they carry the same missile), as well as the DF-31B, which may have both mobile and silo variants. The original DF-31 version has been retired.

The DF-31A (CSS-10 Mod 2) is an extended-range version of the DF-31. With a range of 11,200 kilometers, the DF-31As can reach most of the continental United States from most deployment areas in China. Each DF-31A brigade used to operate only six launchers but they have recently been upgraded to operate 12 (Eveleth 2020). We estimate that China still deploys a total of about 24 DF-31As in two brigades.

In his March 2023 testimony before Congress, US STRATCOM Commander Gen. Anthony Cotton surprisingly suggested that the DF-31A ICBM could carry MIRVs. This differs from the US Air Force's National Air and Space Intel Center's (NASIC) 2020 estimate that the DF-31As are equipped with only one warhead per missile, as well as from the Pentagon's 2022 annual China report, which referred to the DF-41 as "China's first road-mobile and silo-based ICBM with MIRV capability," therefore indicating that the DF-31A is not MIRV-capable (Cotton 2023; National Air and Space Intelligence Center 2020; US Department of Defense 2022b, 94). It remains unclear whether the discrepancy can be attributed to updated intelligence, an incorrect statement by the US STRATCOM commander, or divergent assumptions by different branches of the intelligence community. It is also unclear how the DF-31 family could be MIRV-capable unless China has also designed a smaller-diameter MIRV warhead; some claim China has developed a smaller warhead with a lower yield (Zhang 2025). Adding warheads would also reduce the range of the missile due to a heavier payload. For these reasons and in the absence of further information, we continue to attribute one warhead to each DF-31A.

Since 2017, China's road-mobile ICBM modernization effort has focused on supplementing and possibly replacing the DF-31A version with the newer and more maneuverable DF-31AG and increasing the number of associated bases; we estimate that seven brigades operate the DF-31AG. The new DF-31AG eight-axle launcher is thought to carry the same missile as the DF-31A launcher but has improved off-road capabilities. The US Air Force NASIC's 2020 missile report listed the DF-31AG as having an unknown ("UNK") number of warheads per missile in contrast to the DF-31A, which was listed with only one warhead. This suggests that the AG version could potentially have a different payload (National Air and Space Intelligence Center 2020, 29). However, for the same reasons as for the DF-31A, we currently assume that the DF-31AG is also equipped with a single warhead.

China has also developed a DF-31B ICBM (very likely designed as CSS-10 Mod 3 by the United States) that may have both road-mobile and silo basing modes. The 2025 CMPR assessed that China's ICBM launch from Hainan Island in September 2024—a first since 1980—used a road-mobile DF-31B (US Department of Defense 2025, 23). This was the first mention of the DF-31B in the CMPR since the 2022 edition, when it was rumored to be in development (US Department of Defense 2022b, 65). During China's September 2025 military parade, China unveiled a vehicle carrying a canister with the designation "DF-31BJ." It is possible that the vehicle was a missile loader, and the "J" likely indicates a silo basing mode, as the Chinese character "井" or "jing" means "well" and is used by the PLA to describe silos (Eveleth 2025). This would indicate that the DF-31B is the same missile that is being assigned to the silo fields. However, the fact that the CMPR still describes those missiles as "DF-31 class ICBMs" rather than "DF-31B ICBMs" suggests a degree of uncertainty.

The next phase of China's ICBM modernization is the integration of the new DF-41 ICBM (CSS-20) that began development back in the late 1990s. China displayed 16 DF-41s at its 70th National Day parade in October 2019, with the launchers being said to come from two brigades (New China 2019). In April 2021, the commander of US Strategic Command testified to Congress that the DF-41 "became operational [in 2020], and China has stood up at least two brigades" (Richard 2021, 7). A third base appears to have been completed (see Table 2). The number of garages at the bases indicates that there may be approximately 28 DF-41 launchers deployed.



In previous Nuclear Notebooks, we estimated that the DF-41 could carry up to three MIRVs, which the Pentagon's 2023 and 2024 China reports appeared to validate (US Department of Defense 2023, 107; 2024, 104). It is unknown if all DF-41s will be equipped with MIRVs or if some will have only one warhead to maximize range. In addition to road-mobile launchers, the Pentagon said in its 2024 report that China "appears to be considering additional DF-41 launch options, including rail-mobile and silo-basing" (US Department of Defense 2024, 65; 2022b, 65). It remains unclear whether "silo basing" would refer to China's new silo fields at Yumen, Hami, and Yulin, as this would suggest a combination of different missiles across the silo fields.

During the 2025 military parade in Beijing, China displayed a new ICBM launcher designated the DF-61. The launcher looked identical to the DF-41. It is potentially possible that the DF-61 uses the same launchers as the DF-41, but with a slightly modified missile. The four DF-61s at the 2025 parade apparently came from 651 Brigade at Chifeng (Reuter 2026); a satellite image from March 2026 showed what could either be a DF-41 or a DF-61 at the base (see Figure 2). The 2025 DOD China report does not mention the DF-61.



PLARF DF-41 ICBM Base in Chifeng, China

Vantor FAS FEDERATION OF AMERICAN SCIENTISTS

China is continuing the modernization of its expanded nuclear arsenal by integrating new launchers at brigade bases constructed during the past several years. One example is the suspected 651 Brigade base on the outskirts of Chifeng in the southeastern part of Inner Mongolia province. The base includes four twin-garages possibly for road-mobile launchers for the DF-41 intercontinental ballistic missile. A Vantor satellite image from February 6, 2026, appears to show a DF-41 launcher driving near the garages. The eight DF-41 launchers are supported by a large number of support vehicles and base facilities that include two highbay garages, two underground facilities, and administrative and personnel buildings. Rumors that the base is upgrading to the new DF-61 have not been confirmed.

Figure 2. PLARF DF-41 ICBM base in Chifeng. (Image: Vantor; Annotations: Federation of American Scientists)

China is also developing a new dual-capable missile, known as the DF-27 (CSS-X-24), which reportedly has a range between 5,000 and 8,000 kilometers and likely has a hypersonic glide vehicle payload option (US Department of Defense 2024, 65, 109). This range class is somewhat redundant for the nuclear strike mission, as these distances can already be easily covered by China's longer-range ICBMs. It is therefore potentially possible that the system could ultimately be used in a conventional strike role, and the 2025 CMPR lists the DF-27 in an overview of China's conventional missiles and its role as "land attack/anti-ship" (US Department of Defense 2025, 79). Reporting surrounding the DF-27 is highly unclear, however: The Pentagon's 2024 report states that the missile may be deployed. Moreover, a US intelligence assessment of February 2023 notes that "land attack and antiship variants [of the DF-27] likely were fielded in limited numbers in 2022," whereas in May 2023 the South China Morning Post reported that the DF-27 has been in service since 2019, citing a Chinese military source (R. Chan 2023; US Department of Defense 2023, 67). In June 2021, Chinese state media broadcast videos of what was rumored to be a military exercise featuring the DF-27 (Tiandao 2022), which strongly resembles the DF-26 with an attached conical hypersonic glide vehicle (HGV). This would be similar to how the DF-17 resembles a DF-16 with an attached HGV. US intelligence assessed in February



2023 that China conducted a developmental flight test of a “multirole HGV” for the DF-27, which flew for around 12 minutes and traveled approximately 2,100 kilometers (R. Chan 2023).

Imagery of a new missile launcher, thought to be the DF-27, appeared in China in 2026. One video showed what was said to be three DF-27 launchers driving through a city that was later geo-located to Zhouzhou, south of Beijing (Kristensen 2026).

The Pentagon’s 2024 report noted that “China probably is developing advanced nuclear delivery systems such as a strategic HGV and a fractional orbital bombardment (FOB) system” (US Department of Defense 2024, 65). As of January 2025, China had tested each of these systems at least once. In July 2021, China conducted a test of a new FOB system equipped with a hypersonic glide vehicle, an event described as an unprecedented achievement for a nuclear-armed country (Sevastopulo 2021). According to the Pentagon, the system came close to striking its target after flying around the world, and “demonstrated the greatest distance flown (~40,000 kilometers) and longest flight time (~100+ minutes) of any [Chinese] land-attack weapons system to date” (US Department of Defense 2022b, 65). An operational FOB/HGV system would pose challenges for missile tracking and missile defense systems, as it could theoretically orbit around the Earth and release its maneuverable payload unexpectedly with little detection time, although the US missile defense system is not intended to defend against Chinese missiles.

Figure 3. New possible DF-26 IRBM brigades. (Images: Vantor and Google Earth; Annotations: Federation of American Scientists)

Medium- and intermediate-range ballistic missiles

The most significant development in China’s medium- and intermediate-range ballistic missile force is the production and fielding of the dual-capable DF-26. According to Pentagon estimates, the number of reported DF-26 launchers has increased from 18 to 300 since 2018, with 550 missiles in 2025 (US Department of Defense 2025, 71). We estimate that approximately 300 DF-26 launchers are now fielded across eight brigades, with two additional possible brigades under construction (see Table 2). The DF-26 intermediate-range ballistic missile (IRBM) is dual-capable and launched from a six-axle road-mobile launcher. With its approximate 4,000-kilometer range, the DF-26 can target important US bases in Guam and Northeast Asia, as well as large parts of Russia and all of India.



New Possible DF-26 IRBM Brigades

China is undertaking a significant expansion of its DF-26 intermediate-range ballistic missile (IRBM) force. 647 Brigade in Zhangye (bottom) now appears operational, with garage space for up to 24 TELs. 618 Brigade in Nanchang (middle) recently underwent significant renovation, including a new high-bay garage area where a possible DF-26 TEL can be seen in March 2026. Construction is ongoing at 654 Brigade in Huangling (top), possibly for the DF-26. Many DF-26 have been spotted at the base in recent years.

The DF-26 base structure is undergoing significant developments. One of the most interesting is the expansion of the 611 Brigade (30.6903°N, 117.9011°E) near the town of Rongcheng, east of Chizhou in Anhui province. Chinese President Xi Jinping visited the base in October 2024 and was briefed on “the brigade’s newly introduced weaponry and equipment and examined its training in operating the arms” (Chinese State Council 2024). The new weapon referred to was the DF-26, which has been integrated with the base since 2021. A massive expansion—one that began in 2023 and is now almost complete—shows a completely new structure for a DF-26 base or any known PLARF base.



Significant developments for hosting the DF-26 are also underway at other bases (see Figure 3). Construction appears to have been completed at the 647 Brigade (38.8575°N, 100.4357°E) in Zhangye, which now seems operational. The brigade's two new high-bay garages include 12-vehicle bays each, for a total capacity of 24 TELs. In March 2022, an estimated 24 DF-26 launchers conducted a large exercise on a string of launch pads recently constructed along a road in the valley west of Base 64's Equipment Inspection Brigade garrison outside Changwu town (36.852°N, 101.376°E). It is possible, but unconfirmed, that the DF-26s were visiting Base 64's training center in Xining as part of their preparation for eventually integrating with the 647 Brigade garrison 150 miles (240 kilometers) north.

As a dual-capable system, it seems unlikely that all DF-26s are assigned a nuclear mission. Some sources previously suggested that there is not a singular nuclear-capable variant of the DF-26; rather, they describe every DF-26 as having a "fast switch" capability between nuclear and conventional warheads. The capability potentially extends even after the missile has been loaded onto its launch vehicle (Pollack and LaFoy 2020). However, the special hardware, software, and personnel required for executing the nuclear mission probably place more constraints on the nuclear role. The Pentagon's 2024 and 2025 annual reports to Congress listed three DF-26 variants: the CSS-18 Mods 1, 2, and 3. And at its annual Victory Day parade in September 2025, China publicly unveiled a new conventional variant: the DF-26D (Europe 2025). It is unclear whether the DF-26D is a fourth version or if it is one of the three already identified by the Pentagon (US Department of Defense 2024, 2025). If the DF-26 variants follow the naming convention of the missile's predecessor, the DF-21, then the DF-26A would be nuclear, the DF-26C would be conventional, and the DF-26D would be a conventional missile with an anti-ship mission. There is considerable uncertainty, however, as the DF-26D constitutes the first time that China has publicly displayed a DF-26 variant with a letter suffix designation. The Pentagon describes the DF-26 as capable of both precision land-attack and anti-ship strikes (US Department of Defense 2024). Most DF-26s probably serve conventional missions, given that nuclear warheads have been produced for use by only some of the launchers. We estimate that perhaps a total of 100 warheads is assigned to the DF-26 force. One brigade, the 646 Brigade at Korla, is reportedly tasked with both nuclear and conventional strike missions, the first time this type of dual mission had been confirmed within a single brigade (Xiu 2022, 129, 131). The dual-capable role of the DF-26 (certainly if it is possible to change the launcher's warhead in the field) raises some thorny issues about command and control and the potential for misunderstandings in a crisis. Preparations to launch—or the actual launch of—a DF-26 with a conventional warhead against a US base in the region could potentially be misinterpreted as the launch of a nuclear weapon and trigger nuclear retaliation—or even preemption. China is one of several countries (including India, North Korea, Pakistan, and Russia) that mix nuclear and conventional capabilities on medium- and intermediate-range ballistic missiles.

Citing Chinese defense industry publications, official media commentary, and military writings, the US Department of Defense assessed in 2024 that the DF-26 could eventually be used to "field a lower-yield warhead in the near term" (US Department of Defense 2024, 110). In addition, the US STRATCOM Commander testified in March 2023 that China was making an "investment in lower-yield, precision systems with theater ranges" (Cotton 2023, 6), and the US Department of Defense described the DF-26 in 2025 as a "highly precise theater weapon that would be well suited for delivering a low-yield nuclear weapon" (US Department of Defense 2025, 24).

Separately, previous claims that the medium-range hypersonic glide vehicle-equipped DF-17 may be dual-capable have not been substantiated. The Pentagon's 2022 China report had noted that "[w]hile the DF-17 is primarily a conventional platform, it may be equipped with nuclear warheads" (US Department of Defense 2022b, 65). However, this language was not included in subsequent annual reports, which only describe the DF-17 as a conventional weapon (US Department of Defense 2023; 2024, 2025). Consequently, we do not include the DF-17 in our estimate of Chinese nuclear forces.

The former mainstay of China's regional nuclear missile capability, the DF-21A (CSS-5 Mod 2/6) medium-range ballistic missile (MRBM), has not been listed as operational in the last three Pentagon reports of Chinese military developments and may have been replaced by the DF-26 in the nuclear mission.

Submarines and sea-based ballistic missiles

China currently fields a submarine force of six second-generation Jin-class (Type 094) nuclear-powered ballistic missile submarines (SSBNs), which are based at the Yalong naval base near Longpo on Hainan Island. The two newest SSBNs are believed to be improved variants of the original Type 094 design. Some Chinese journals refer to it as the Type 094A, but this has not been confirmed by either the Pentagon or the Chinese government. These SSBNs include a more prominent hump, which initially triggered some speculation as to whether they could carry up to 16 submarine-launched ballistic missiles (SLBMs), instead of the usual 12 (Suciu 2020; Sutton 2016). However, satellite images subsequently confirmed that the new subs are equipped with 12 launch tubes each (Kristensen and Korda 2020). The upgrades were later assessed to be related to sound silencing (Carlson and Wang 2023, 18).

Per the Pentagon's 2024 China Military Power Report, China has equipped its Jin-class SSBNs to carry either the 7,200-kilometer range JL-2 (CSS-N-14) SLBM or the longer-range JL-3 (CSS-N-20) SLBMs, and China has likely begun replacing the JL-2s with JL-3s on a rotational basis as each submarine returns to port for routine maintenance and overhaul (US Department of Defense 2023, 55). The range of the JL-2 is sufficient to target



Alaska, Guam, Hawaii, Russia, and India from waters near China, but not the continental United States—unless the submarine sails deep into the Pacific Ocean to launch its missiles. With the JL-3's longer range of roughly 10,000 kilometers, a submarine will be able to target the northwestern parts of the continental United States from northern Chinese waters, but not from the South China Sea. And it would still not be able to target Washington, DC without sailing past northeast Japan (National Air and Space Intelligence Center 2020, 33). Unlike the JL-2, the JL-3 allegedly can deliver “multiple” warheads per missile (33). It is unknown if the JL-3 is indeed equipped with multiple warheads; US intelligence has not explicitly stated that it is. The People's Liberation Army Navy reportedly conducted its first test of the JL-3 in November 2018 (Gertz 2018) and appears to have conducted at least two—possibly three—additional tests since then (M. Chan 2020; Guo and Liu 2019).


Although the Jin-class is more advanced than China's first experimental SSBN—the single and now inoperable Xia (Type 092)—it is a noisy design compared with current US and Russian missile submarines. It is estimated that the initial Type 094s was two orders of magnitude louder than the stealthiest Russian or American SSBNs (Coates 2016). However, the enhanced Type 094A SSBN is probably less noisy (Lin and Singer 2017). This may explain why China has decided to continue construction of additional Type 094A SSBNs, rather than transitioning entirely to the next-generation Type 096. Another reason may be due to production delays: The new Type 096 was scheduled to begin construction in the early 2020s, but after apparent delays, the Pentagon's 2024 CMPR states

that China will likely begin construction of the Type 096 “soon” (US Department of Defense 2024, 53).

Figure 4. SSBN activity at Longpo Naval Base, China. (Images: Planet Labs PBC and Vantor; Annotations: Federation of American Scientists)

In 2022, China completed a new construction hall at Huludao shipyard, where the PLA's submarines are built, which could be for the production of the Type 096 SSBN (Sutton 2020). Satellite images show wider hull sections at Huludao, which would make sense for the new SSBN that is expected to be larger and heavier than the Type 094 (Sutton 2021). However, we could not confirm whether the new facilities at Huludao correspond to the larger Type 096 SSBN or a new unknown attack submarine.

As with all new designs, the Type 096 is expected to be quieter than its predecessor. Some experts even believe it could be as quiet as Russia's new Borei-class SSBNs (Carlson and Wang 2023, 30). Although that would be a significant



SSBN Activity at Longpo Naval Base

Satellite imagery shows Jin-Class ballistic missile submarines (SSBNs) being escorted by tugboats to and from the tunnel at Longpo Naval Base, where weapons handling likely takes place. Submarines undergo light maintenance and weapons handling before embarking on, and after returning from, near-continuous deterrence patrols or transits to other locations. Construction continues at the base, including for what appears to tunnels for a new underground submarine facility.

Satellite Imagery © 2026 Vantor and Planet Labs PBC

FAS
FEDERATION OF
AMERICAN SCIENTISTS

Vantor Planet

technological leap for China, The Financial Times reported in September 2024 that, according to US naval researchers, Russia was providing technological assistance to China to achieve a quieter propulsion system for its Type 096 submarines (Foy et al. 2024). Some anonymous defense sources have speculated the Type 096 will carry 24 missiles (M. Chan 2020), but no public official sources have confirmed this. Current and projected missile inventories seem to indicate that the new SSBN will more likely carry 12 to 16 missiles. The Pentagon's 2024 report stated that the Type 096 SSBNs “will reportedly be



armed with a follow-on longer range SLBM,” and that these SLBMs can be MIRVed (US Department of Defense 2024, 53, 88). Given that China’s SSBNs are assumed to have a service life of approximately 30 to 40 years, the US Department of Defense expects that the Type 094 and Type 096 boats will operate concurrently (US Department of Defense 2024, 56). If confirmed, this could potentially result in a future fleet of eight to 10 SSBNs.

All of China’s six SSBNs—and several attack submarines—are based at the Longpo naval base on Hainan Island, where China completed construction of two additional piers in 2024 to accommodate more submarines. Satellite imagery indicates that significant construction continues to take place on the Yalong Peninsula, including for support facilities, a radar station, and what seems to be a tunnel that could be intended for a second entrance into the mountain for the submarines. The satellite images captured three instances in 2025–2026 of Jin-class SSBNs being assisted by tugboats when leaving or entering the tunnel at the base (see Figure 4).

A Jin-class submarine can be seen docked at a jetty at Jianggezhuang Naval Base—the oldest SSBN base that used to house the Xia-class Type-092 SSBN in the 1980s—periodically beginning in May 2025. While a singular Jin-class used to be based here in the 2010s, none have been spotted at this base for almost a decade. The purpose for the submarine’s presence at the base is still unclear, including whether it indicates an intent to operate Jin-class SSBNs from the Northern Fleet base more frequently.

The Pentagon’s 2022 report indicated that China had begun “near-continuous at-sea deterrence patrols with its six JIN class SSBNs” in 2021 (US Department of Defense 2022b, 96), and the 2024 report asserted that China “probably continued to conduct” these patrols (US Department of Defense 2024, 104). The term “near-continuous” implies that the SSBN fleet is not on patrol all the time but that at least one boat is deployed intermittently, which is consistent with satellite imagery observations. There are multiple occasions in which a total of four or five submarines can be seen on satellite imagery across multiple locations, with usually at least one SSBN at the Bohai Shipyard at Huladao for fitting out. On March 8, 2026, satellite imagery reveals that four SSBNs were docked at the Longpo Naval Base while two were at the Bohai shipyard, accounting for the entire fleet.

The term “deterrence patrol” could imply that the submarine at sea has nuclear weapons onboard, although US officials have not publicly stated so. Giving PLAN custody of nuclear warheads to deployed submarines during peacetime would constitute a significant departure from Chinese declaratory policy and a significant change for China’s Central Military Commission, which has historically been reluctant to hand out nuclear warheads to the armed services. In late July 2025, China Central Television (CCTV) released footage of a Type 094 on a strategic patrol, including a clip of a Chinese sailor remarking that his unit would fire its missiles “without hesitation” upon receiving a launch order (CCTV 2025). To fully develop a survivable sea-based nuclear deterrent posture, China is presumably improving its command-and-control system to ensure reliable communication with the SSBNs when needed and prevent the crew from launching nuclear weapons without authorization. Moreover, the SSBN fleet will have to operate safely in patrol areas from where its missiles can reach intended targets. Western military officials have privately stated that the United States, Japan, Australia, and the United Kingdom “are already attempting to track the movements of China’s missile submarines as if they are fully armed and on deterrence patrols” (Torode and Lague 2019). Whenever they are in the South China Sea, China’s SSBNs typically appear to be accompanied by a protection detail, including surface warships and aircraft (and possibly attack submarines) capable of tracking adversarial submarines (Torode and Lague 2019).

Given the noise level of the SSBNs, it seems likely that China during conflict would keep the submarines inside a protected “bastion” in the South China Sea or the Bohai Sea (US Department of Defense 2024, 104). Even with the JL-3 SLBM, the SSBNs would not be able to target the continental United States from the South China Sea. To target the northwestern parts of the continental United States, they would have to sail far north, likely to the Bohai Sea.

Bombers

China developed several types of nuclear bombs and used aircraft to deliver at least 12 of the nuclear weapons that it detonated in its nuclear testing program between 1965 and 1979. Later, however, the People’s Liberation Army Air Force (PLAAF) nuclear mission became dormant as the rocket force improved and older intermediate-range bombers were unlikely to be useful or effective in the event of a nuclear conflict. While we previously estimated that China maintained a small inventory of gravity bombs for potential contingency use by aircraft, we assess that this mission no longer exists and that China no longer retains gravity bombs in its nuclear arsenal.

The PLAAF’s nuclear mission remained dormant until approximately 2017–2018, when, coinciding with a renewed emphasis on nuclear aircraft modernization, the US Department of Defense reported in 2018 that the PLAAF “has been newly re-assigned a nuclear mission” (US Department of Defense 2018a, 75, 34). This new mission appears to be currently centered around China’s current H-6 “Badger” bomber. China has hundreds of these aircraft in more than eight distinct variants; however, only one of these variants was assessed to be nuclear-capable: the H-6N.

The H-6N is distinct from other variants of the H-6 in that it incorporates a nose-mounted in-flight refueling probe (Rupprecht 2019). Not only does the H-6 lack intercontinental range like Russian and US nuclear bombers, but the H-6N is also equipped to carry large air-launched ballistic missiles (ALBMs), the significant added weight of



which would further decrease the aircraft's range. Some analysts speculate that this is likely the primary reason why the in-flight refueling capability was added to the H-6Ns (Pasandideh 2025). Notably, the H-6N airframe modification included the removal of the bomb bay, supporting the conclusion that China's legacy gravity bomb mission has indeed ended. China has developed at least two types of ALBMs, both of which appear to be variants of land- or sea-based ballistic missiles. The conventional KD-21 ALBM appears to be a wing-mounted air-launched version of the YJ-21 sea-launched anti-ship ballistic missile (Panda 2019; Rupprecht 2024). Imagery from a Chinese military exercise in April 2025 indicates that the missile is in operational service with H-6K conventional bombers (Newdick 2025).



Is China adding a second nuclear air base? - 31.6238°N 116.2715°E

Vantor FAS

Lu'an Air Base in Anhui province has undergone major upgrades in recent years, similar to those at Neixiang Air Base. Upgrades include widened roads and tunnel entrances into a mountain, new security perimeters and facilities, and protective shelters. Additionally, multiple nuclear-capable H-6Ns can be seen at the base in satellite imagery from December 2025, identifiable by their aerial refueling probes. The renovations could possibly indicate a nuclear mission for the base, but there is significant uncertainty.

Figure 5. Is China adding a second nuclear air base? (Images: Vantor and Google Earth; Annotations: Federation of American Scientists)

At its annual Victory Day Parade in September 2025, China for the first time displayed a nuclear weapons system for delivery by aircraft: the JingLei-1 or JL-1 ALBM. The missile, described by the official parade commentator as an "air-based long-range missile," is likely the nuclear-capable ALBM designated by the United States in 2024 as CH-AS-X-13 that China has been developing for delivery by the H-6N. (It is possible, but unconfirmed, that the X has since been removed from the designation, given the system is now deployed.) Following at least five developmental and user tests, the Defense Intelligence Agency assessed in August 2024 that the missile had been deployed (Defense Intelligence Agency 2024, vi). The Pentagon assessed that, with the deployment of the nuclear ALBM, China possesses "a viable nuclear 'triad' of land, sea, and air forces, albeit a much less complete or capable triad than that of the United States or Russia (US Department of Defense 2019, 67). In its 2025 annual report to Congress on China's military power, the Pentagon describes the H-6N's ALBM as a "highly precise theater weapon that would be well suited for delivering a low-yield nuclear weapon" (US Department of Defense 2025, 24).

Currently, only one PLAAF unit is publicly known to have a nuclear mission: the 106th Brigade at Neixiang Air Base in the southwestern part of Henan province, one of China's only airfields with an adjacent air defense site.



Over the past five years, the base has been modified extensively, including the addition of large tunnel entrances into a nearby mountain wide enough to comfortably accommodate the H-6N bomber. Civilian video footage from October 2020 appears to show an H-6N bomber flying with the possible new JL-1 ALBM just outside of Neixiang Air Base (Lee 2020a, 2020b; Rupprecht and Dominguez 2020). The numbers of H-6N bombers and their assigned nuclear weapons are unknown, but we cautiously estimate that 20 aircraft are assigned up to 20 missiles.

It is unknown how many nuclear bomber brigades China plans to establish, but it is potentially possible that Lu'an Air Base in the eastern part of Anhui province could be a second deployment site for the H-6N bombers. Satellite imagery shows significant construction at the base on a similar timeline and aspect of the construction at Neixiang, including expanded tunnel entrances, new security perimeters, large aircraft shelters, and other new infrastructure (see Figure 5). Additionally, the base was already known to host conventional variants of the H-6, and H-6Ns can be seen at the base in recent imagery. However, some features seen at Neixiang appear to be absent at Lu'an, including a special hangar area connected to a tunnel entrance as well as a nearby air-defense site. There has also been no official confirmation or mention of Lu'an as a future host base for the H-6N. Therefore, the base's possible status as a host base for the nuclear mission comes with significant uncertainty.

In November 2024, China and Russia conducted joint air drills that, for the first time, included the participation of the H-6N (globaltimesnews 2024).

To eventually replace the H-6, China is developing a stealth bomber with a longer range and improved capabilities. The Pentagon estimates that the new bomber, known as H-20, will have both nuclear and conventional capability and a range of more than 10,000 kilometers. If equipped with an aerial refueling capability, the Pentagon assesses that the bomber could potentially have an intercontinental range (US Department of Defense 2023, 92). Videos and pictures have appeared on social media claiming to show the H-20 in flight, but the imagery has not yet been authenticated, nor has the flight been confirmed by official sources (Tirpak 2025; 前HR本人 2025). In a February 2026 interview, the head of US Air Force Global Strike Command said, regarding the H-20, that China is “just not there yet,” and that “China is a regional bomber force at best” (Trevithick and Altman 2026). The Pentagon's 2025 report to Congress also notably made no mention of the H-20.

Cruise missiles

From time to time, various US military publications have asserted somewhat vaguely that one or more of China's cruise missiles might have nuclear capability. For example, a nuclear modernization fact sheet published by the Pentagon in connection with the release of the 2018 Nuclear Posture Review claimed, without identifying them, that China had both air-launched and sea-launched nuclear cruise missiles (US Department of Defense 2018b). The Pentagon has not substantiated this claim since. The 2023 Japanese Defense Paper, however, stated that the H-6 bombers “are believed to be capable of carrying long-range attack cruise missiles with nuclear capability” (Japanese Ministry of Defense 2023, 67).

It is still unclear what this missile could be—if it exists at all. Therefore, we continue to assess that, although China might have developed warhead designs for potential use in cruise missiles, it currently has no nuclear cruise missiles in its active stockpile. It is potentially possible, but unconfirmed, that the future H-20 could be equipped with a nuclear cruise missile.

This research was carried out with generous contributions from the Carnegie Corporation of New York, the New-Land Foundation, Ploughshares, the Prospect Hill Foundation, and individual donors.

Notes

[1] Nuclear weapons are stored in central facilities under the control of the Central Military Commission. Should China come under nuclear threat, the weapons would be released to the Second Artillery Corps to enable missile brigades to go on alert and prepare to retaliate. For a description of the Chinese alerting concept, see Kristensen (2009). For more on warhead storage in China, see Stokes (2010) and Bruzzese (2026). For an overview of the People's Liberation Army Rocket Force structure and organization, see Stokes (2018) and Xiu (2022). For an insightful overview of Chinese thinking about nuclear weapons and policies, see Santoro and Gromoll (2020).

[2] The “continental United States” as used here includes only the lower, contiguous 48 states. US states and territories outside of the continental United States include Alaska, Hawaii, Guam, American Samoa, Puerto Rico, the US Virgin Islands, and many tiny Pacific islands.

●► References are available at the source's URL.

Hans M. Kristensen is the director of the Nuclear Information Project with the Federation of American Scientists in Washington, DC. His work focuses on researching and writing about the status of nuclear weapons and the policies that direct them. Kristensen is a coauthor of the world nuclear forces overview in the SIPRI Yearbook



(Oxford University Press) and a frequent adviser to the news media on nuclear weapons policy and operations. He has coauthored the Nuclear Notebook since 2001.

Matt Korda is the associate director of the Nuclear Information Project at the Federation of American Scientists, and an associate senior researcher with the Nuclear Disarmament, Arms Control and Non-proliferation Programme at the Stockholm International Peace Research Institute (SIPRI). Previously, he worked for the Arms Control, Disarmament, and WMD Non-Proliferation Centre at NATO headquarters in Brussels. Korda's research and open-source discoveries about nuclear weapons have made headlines across the globe, and his work is regularly used by governments, policymakers, academics, journalists, and the broader public in order to challenge assumptions and improve accountability about nuclear arsenals and trends. He received his MA in International Peace and Security from the Department of War Studies at King's College London.

Eliana Johns, née Reynolds, is a senior research associate for the Nuclear Information Project at the Federation of American Scientists, where she researches the status and trends of global nuclear forces and the role of nuclear weapons. Johns is also a master's student in the Security Studies Program at Georgetown University's Walsh School of Foreign Service, where she focuses on the intersection of technology and security. Previously, Johns worked as a project associate for DPRK Counterproliferation at CRDF Global, focusing on WMD nonproliferation initiatives to curb North Korea's ability to gain revenue to build its weapons programs. Johns graduated with her bachelor's in political science with minors in Music and Korean from the University of Maryland, Baltimore County. She also completed the Critical Language Scholarship program for Korean through the US Department of State and is a CSIS 2025 Nuclear Scholar. Johns' work has been widely published and quoted in media sources across the globe to help policymakers and the public better understand nuclear weapons trends.

Mackenzie Knight-Boyle is a senior research associate for the Nuclear Information Project at the Federation of American Scientists, where her work focuses on the status and trends of global nuclear forces and the role of nuclear weapons. Before this, Knight-Boyle was a Herbert Scoville Jr. Peace Fellow on the Project. Previously, Knight-Boyle worked as a policy and communications intern at the Arms Control Association, as a summer fellow with the James Martin Center for Nonproliferation Studies (CNS), as an analyst intern with Shephard Media in London, and most recently as a graduate research assistant at CNS while obtaining her master's degree in Nonproliferation and Terrorism Studies from the Middlebury Institute of International Studies at Monterey. She received bachelor's degrees in Middle Eastern Languages and Cultures and Policy and Intelligence Analysis from Indiana University.

IRAN: What's Up With the Nuclear Dust?

By Raymond White | Retired CIA Senior Operations Officer

Source: <https://www.linkedin.com/pulse/iran-whats-up-nuclear-dust-raymond-white-ozw0e/>



July 13 – BID is often asked what our take is on Iran's nuclear program and, more specifically, *what's up with the Nuclear Dust?* Although the recent focus in the press has been on Iran's throwing the ceasefire completely out the window, the Strait of Hormuz, and the Ayatollah's funeral, it might be worth taking a closer look at what has been happening recently, both above and below the ground.

Let's Break It Down...

Vice President Vance's Comments

Let's rewind a bit. On 22 June 2026, some interesting statements were made at the Bürgenstock Resort in Switzerland.

VP Vance had traveled to Switzerland to meet with Iranian officials and nail down the technical parameters of the 14-point U.S.–Iran Memorandum of Understanding (MOU). The 10-hour marathon session of four-way talks was mediated by Qatar and Pakistan.

Although BID appreciated Vance's straightforward briefing on what had transpired (versus Trump's rambling generalities), he is currently 0 for 2 in terms of successful outcomes with Iranian negotiators. Poor Vance keeps getting called in to travel halfway around the world to be *the closer*, and then, after he departs, everything unravels. Perhaps he should be *the opener*, and things could unravel much earlier? This would save time, I suppose. JD is likely getting a complex at this point.

But Vance's travel schedule wasn't the only thing sliding out of sync. What BID found most interesting was not reported by the press—in fact, it went right over their heads (whoosh!). They were probably too busy prepping for the World Cup or Summer vacation at the beach.

During his live press briefings detailing the administration's framework and what had transpired during meetings, JD Vance initially made a sweeping statement regarding the red lines for the negotiations, asserting that:

"Iran will not have a nuclear program."

Vance quickly corrected himself and restated that:

"Iran will not have a nuclear weapon."

For BID, this recharacterization was quite interesting, especially after Vance and company had just spent ten hours negotiating the minutiae of the Strait of Hormuz and ostensibly, Iran's nuclear program (civilian and weapons). Diplomatic negotiations are focused on the minute details (think lawyers nitpicking every sentence of a contract). Undoubtedly, the nuclear file was discussed in great detail.

President Trump has made it very clear that a nuclear weapon (and/or nuclear weapons program) is a red line. ***Vance's correction of his statement is telling in that it suggests the Administration is willing to accept Iran having a nuclear program.***

Presumably, a *nuclear program* refers to a civilian nuclear program for energy and medical purposes. Either way, the Administration's acceptance of a civilian nuclear program suggests it is walking back its hardline position on Iran having zero nuclear programs.

Rightly so, President Trump has also repeatedly questioned why a nation that sits on some of the largest oil and gas reserves on earth has such an urgent, multi-billion-dollar need for a domestic nuclear energy program in the first place. BID agrees.

Iran has argued that other Gulf countries have civilian nuclear programs, so why can't they? After all, everyone else is doing it (dude). Well, Iran is a different case. They need to be held to a different standard. As far as geopolitics go, a civilian nuclear program is a nice-to-have, not a need. Most importantly, Iran has demonstrated that it won't be happy with just a civilian program; they want a weapon. BID suggests the White House stick to its guns and tell the misbehaving child NO. And right now, that child is busy digging in their sandbox.

Recent Activity at Key Nuclear Sites

Commercial Satellite imagery has shown recent construction activity at some of Iran's key nuclear sites. For those interested in a quick breakdown on Iran's nuclear weapons program, please check out our BID article, [Nuclear Triggers and Missing Highly Enriched Uranium \(HEU\)](#).

Pickaxe Mountain (Kuh-e Kolang)

According to a recent Foundation for Defense of Democracies (FDD) article, [Iran Continues Work at Key Nuclear Site, Violating U.S.-Iran Agreement](#), Pickaxe Mountain showed "vehicle activity ... on the roads leading to the open set of Western tunnel portals, indicating that construction inside the tunnel complex, as well as hardening of the tunnel entrance, are ongoing." The Institute has previously documented Iranian efforts to fortify the site's entrances against possible U.S. or Israeli strikes or raids, along with apparent construction inside the deeply buried tunnel complex.

The Eastern set of tunnel portals remains partially backfilled with dirt, making them inaccessible to ground vehicles. This was likely done to prevent rapid ingress and egress to the tunnel entrances.

Keep in mind, Iran restricted the International Atomic Energy Agency (IAEA) from inspecting this site. Since 2021, the Islamic Republic has [maintained](#) that Pickaxe Mountain is a centrifuge manufacturing facility. However, Western intelligence agencies [suspect](#) that Tehran is also building an undeclared uranium enrichment plant there.

Basically, we have no clue what is happening in Pickaxe Mountain.



Fordow, Natanz, or Isfahan Sites

Since the MOU was signed, the FFD has observed no *significant activity* at Iran's other major nuclear sites — Fordow, Natanz, or Isfahan — all of which were targeted by the United States in June 2025. In a recent post on X, the [Institute for Science analyzed commercial satellite imagery](#) from these sites and determined

At Natanz, little activity can be seen. The access points to the below ground enrichment halls have not been repaired. The personnel entrances remain destroyed and vehicle entrances remain severely damaged. A single vehicle can be seen on the road outside of the Pilot Fuel Enrichment Plant (PFEP), which was destroyed in June 2025, but was covered by Iran after. The HVAC chillers remain displaced throughout the complex and the primary and backup power systems remain destroyed.

At Fordow, as earlier reported by the Institute, between May 10 and May 18th, Iran added passive defensive measures in the form of earthen/rocky mounds and other objects on the roads leading to the tunnel entrances. The alternate placings of the piles/objects are very precise, which creates a series of chicanes, indicating they are not intended as obstructions, but rather to prevent rapid ingress and egress by any vehicle towards the tunnels. The June 21 Vantor image shows that the objects along the road remain there. The tunnel portals also remain backfilled with dirt.

As of June 29, 2026, there is no observed activity at Esfahan. The tunnel portals remain backfilled with dirt.

So What's Going Down?

Based on the analysis above, BID assesses that the sites are being reinforced but not enhanced or hardened. President Trump has made it clear that the Space Force is closely monitoring these sites. The mainstream media hears the words "*nuclear sites*" and quickly jumps to extreme conclusions. What may be happening is likely much more mundane.

In moving earth *to prevent rapid ingress and egress*, BID assesses that Iran could be taking what they deem to be appropriate defensive measures to prevent what they assess would be a Call of Duty-style assault on the facilities. Or perhaps, squatters are moving in? Undoubtedly, Iran monitors X and the usual suspects on the internet for news updates, but perhaps, the Administration should consider *a love tap* every now and then at these facilities, to physically remind Iran we are actually watching?

As the Iran saga continues to play out, and as Iran misdirects the West's attention to everything under the sun except a peace deal, BID will continue to keep an eye out for the nuclear dust.

Until next time.

Bacteria Are Thriving In a Radioactive Former Soviet Mine – Leaving Only 5% of The Uranium Behind

Source: <https://www.sciencealert.com/uranium-eating-bacteria-leave-just-5-of-the-radioactive-metal-in-toxic-mine-water-scientists-discover>

July 14 – One of the world's largest uranium mines, the Wismut GmbH Schlema-Alberoda operation in what was then Soviet East Germany, left behind a toxic legacy.

But within the contaminated water that has since flooded the mine, evolution may already be brewing up a solution.


The mine site was closed in 1990, [with the reunification of Germany](#), and has since been subject to costly and time-consuming remediation efforts.

In its retirement, the underground mine became flooded with water, which has required continuous treatment.

You may already be aware that raw uranium is highly radioactive, and exposure to it – for instance, by drinking contaminated water – can cause serious [damage to humans and other living things](#).

Yet, some organisms are actually making a living in the uranium-laden mine water; it's home to an entire ecosystem of microbes.


And, as scientists have recently discovered, those microbes can actually stabilize uranium under certain conditions.



An illustration representing the formation of different nanoparticles in the cell membranes of bacteria from mine water. (HZDR/J. Raffi/E. Krawczyk-Bärsch/edited with AI)



The research was led by microbiologists and resource ecologists at the Helmholtz-Zentrum Dresden-Rossendorf (HZDR) in Germany and the University of Granada in Spain, who published their results in the journal [Nature Communications](#).



On day 1 of incubating the microbes with glycerol, the mine water appeared yellow. After 130 days, a black precipitate formed at the bottom of the bottles, while the overlying water remained clear. (Newman-Portela et al., *Nat. Comms.*, 2026)

"Our group's investigations had already revealed that bacteria can use uranium dissolved in water for their metabolism when they have access to glycerol as a food source," [explains](#) microbiologist Evelyn Krawczyk-Bärsch, of HZDR.

Microscope images showing the uranium nanoparticle aggregates on microbe cells. (Newman-Portela et al., *Nat. Comms.*, 2026)

"Our study has revealed for the first time that bacteria supplied with glycerol as a carbon source can convert toxic uranium dissolved in water into a stable chemical compound."

Krawczyk-Bärsch and team began their experiments with water samples collected from the treatment plant inlet of the Wismut GmbH Schlema-Alberoda mine.

"We wanted to create natural conditions for the bacterial community already existing in the mine water because at a depth of approximately 2,000 meters there is usually little or

no oxygen in the mine," HZDR microbiologist Antonio Newman-Portela [says](#).

When they incubated the bacteria with glycerol, they found that the bacteria converted uranium into a pentavalent state.

When [uranium is pentavalent](#), it has an unusual [oxidation state](#) of +5, which changes how it bonds to other elements, making it easier to 'lock up' within stable minerals.

"Uranium usually occurs with a valency of 4 or 6. Pentavalent uranium does exist, but it is rare or only transient. Until now, it had been seen in an unstable oxidation state," [explains](#) Antonio Newman-Portela.


In the presence of the bacteria, pentavalent uranium is then combined with iron and oxygen to form $\text{FeU}^{(\text{V})}\text{O}_4$ – a compound that scientists were already aware of but have yet to give a 'common' name to.

What they didn't know was that it could form in nature, let alone that bacteria were involved.

"After 130 days, only around five percent of the uranium dissolved in the water remained in the samples," [says](#) Newman-Portela.

The bacteria had not only incorporated the uranium into their cell walls, but an unusually high proportion of that uranium was pentavalent. This meant it more readily formed $\text{FeU}^{(\text{V})}\text{O}_4$, especially when the water samples were dried and exposed to oxygen.

Radioactive uranium contamination is a global problem.



In the United States, India, Canada, France, South Africa, and Australia, surface water and groundwater have



sometimes exceeded the 0.03 milligram per liter guidelines for uranium contamination.

Could bacteria be part of the solution?

"Over the last three decades, bioremediation has been explored as a cost-effective alternative to physico-chemical water treatment," the authors [write](#).


"Field studies [using biological methods] have demonstrated substantial uranium reduction while avoiding the generation of secondary sludge."

Perhaps these bacteria could be allies in our quest to clean up nuclear contamination, not just in Germany, but across the world.

"Although derived from a single geochemical scenario, the processes identified here are broadly applicable to other contaminated waters," the authors [conclude](#).

But, as Krawczyk-Bärsch [points out](#), "we still have to investigate to what extent bacteria might help to render uranium harmless for remediation purposes."

► The research was published in [Nature Communications](#).



Thirteen films that are essential to understanding the nuclear age

By Alex Wellerstein

Source: <https://thebulletin.org/premium/2026-07/thirteen-films-that-are-essential-to-understanding-the-nuclear-age/>



July 16 – Nuclear films are both reflections of the *zeitgeist* in which they are made and shapers of it. They feed off of both popular and expert anxieties, while also helping to solidify the terms on which these anxieties are engaged. While some themes are apparently universal (how many ways are there to blunder into nuclear war), some feel very much of their particular time and place.

In creating this list, I was cognizant that dozens of films incorporate some sort of “nuclear” elements into their plots. Narrowing it down to 13 required implementing arbitrary criteria. Ultimately, I decided that any film on my best-of list needed to be explicitly engaged with nuclear themes in a serious and sustained way. (And so no films in which nukes are merely “MacGuffins” that could be swapped out with other “dangerous objects” would be included. Nor would



there be any films which could only be allegorically understood to be about nuclear issues—no “the-dragons-are-nukes,” unless actual nukes also play a role in the plot). I focused on films about nuclear weapons, not nuclear power, with one important exception. I arbitrarily limited the scope to fiction, not documentaries (but if I were including documentaries, Errol Morris’ 2003 [The Fog of War](#) would be first in line). I also excluded all television and streaming series (sorry, HBO’s 2019 [Chernobyl](#)).

And, finally, any film on the list would have to be a “good film,” one that I felt I could recommend people watch without reservation or qualification. The was the trickiest and most subjective criterion, particularly for a historian, because there are many films that I think are profoundly interesting as artifacts, or have historical importance, but are not, in fact, great films. For example, while MGM’s *The Beginning or the End* (1947) [is historically fascinating](#)—it was made in collaboration with Manhattan Project officials—it is actually quite tedious to watch, in my opinion.

This list has a strong Anglophone bias, and with a few notable exceptions approaches the topic from largely American points of view. This is a reflection of its author’s limitations and the film’s relative accessibility to an American. Undoubtedly lists could (and should) be created from different cultural and national perspectives. The list is presented chronologically and can be thought of as an unfolding of the nuclear perspective from the early nuclear age to the present.



Godzilla

1954, Toho. Directed by Ishirō Honda. Released in the USA as *Godzilla, King of the Monsters*, in 1956.


Inspired in part by the 1954 [Castle Bravo](#) fallout disaster, which killed a Japanese fisherman and led to a boycott on tuna in Japan, *Godzilla* (*Gojira* in the Japanese original) is a more complicated film than just “radioactive monster attacks Tokyo.” *Godzilla*, the monster, is less of a metaphor for the atomic bomb itself than a representation of Japan’s submerged psychological reckoning with the horrors experienced by its civilians in World War II. Such discussions were officially suppressed during the Allied Occupation of Japan (1945-1952), but, like the monster, emerged as a consequence of US nuclear testing in the Pacific. The sequences with the rubber monster costume have probably aged the worst, but the films of how the people reconcile themselves with their destruction and their duty still hold up extremely well.

On the Beach

United Artists, 1959. Directed by Stanley Kramer.

The premise of [On the Beach](#), adapted from the 1957 novel by the British author Nevil Shute, is that a future global nuclear war using cobalt bombs has both eradicated most of human civilization and created a massive, deadly cloud of fallout that is gradually circulating the globe, killing all in its path. Those in the southern hemisphere are momentarily safe, but the

cloud is coming. An American submarine crew pursues one last-ditch hope to discover a path to survival, but it turns out to be spurious. All must make peace, in their own way, with the fact that the Earth will soon be dead. Although the underlying premise is technically dubious, the psychological resonances made it one of the most potent nuclear-themed films of the 1950s, and highly representative of a strain of fatalism that surrounds nuclear war.



Dr. Strangelove or: How I Learned to Stop Worrying and Love the Bomb

1964, Columbia Pictures. Directed by Stanley Kubrick.

Stanley Kubrick’s dark comedy about the risk of unauthorized and accidental nuclear war fundamentally altered the lexicon with which we talk about nuclear weapons, and the performances of Peter Sellers, George C. Scott, Sterling Hayden, and Slim Pickens created durable cultural archetypes. While loosely adapted from Peter George’s 1958 novel *Red Alert*, Kubrick’s own sensibilities both about nuclear war and how to talk about it run through it. The themes of unauthorized nuclear use, mutual assured destruction, and the dangers of automation are all deadly serious, but wrapping it all in a farce makes it a work of genius and helps





The China Syndrome

1979, Columbia Pictures. Directed by James Bridges.

The sole film specifically about nuclear power on this list, *The China Syndrome* is a taut thriller that was elevated into something even more by its historical moment. The film, starring Michael Douglas, Jane Fonda, and Jack Lemmon, is about the threat of a meltdown at a nuclear power plant in Southern California. The threat wasn't the reactor, per se, it was the corner-cutting and coverup perpetuated by a power utility more concerned with profits than

it stand the test of time. When [Daniel Ellsberg](#), then still at RAND and having just completed a classified study into the command and control problems central to the plot, saw the film in theatres, he was amazed — it was, he wrote later, “essentially, a documentary.”

Fail Safe

1964, Columbia. Directed by Sidney Lumet.

Fail Safe might be remembered as the film of mid-1960s nuclear anxiety had it not come out the same year as *Dr. Strangelove* and overlapped so significantly with its themes—so significantly that Kubrick sued it for copyright infringement and caused it to be released after his own film. Lumet's film, a fairly faithful adaptation of Eugene Burdick and Harvey Wheeler's 1962 novel of the same name, plays the subject of accidental nuclear war and lack of control straight. Although rendered a box office flop by its circumstances, and while certainly the lesser of the two films, its management of tension has led to its retrospective status as a critical success.



safety. The film would work on its own merits, but the fact that the film premiered less than two weeks before the [Three Mile Island](#) accident catapulted it to the position of prophecy—for better and for worse.

Barefoot Gen (Hadashi no Gen)

1983, Kyodo Eiga. Directed by Mori Masaki.


Barefoot Gen, first a *manga* (Japanese comic book) published in 1976, is a fictionalized version of writer and illustrator Keiji Nakazawa's experiences of being a six-year-old boy growing up in Hiroshima before, during, and after the atomic bombing. It is powerful, rich, and moving. Much of the story is a preamble obscure to most Westerners: what life was like in the city before the attack, with its food shortages and political repression. The “cute” manga style fits the worldview of the child protagonists, increases the dread for what the viewer knows is coming, and contrasts starkly with the nothing-held-back depictions of the aftermath. The

anime (cartoon movie) version was released in 1983, and is an exceptional adaptation of the form, content, and message of the original.

The Day After

1983, ABC. Directed by Nicholas Meyer.

This classic made-for-television movie about the run-up to nuclear war and its immediate aftermath aired on ABC in November 1983 and was seen by nearly 100 million people—over 40 percent of the entire US population. Coming at one of the highest points of Cold War nuclear anxiety in the United States and Western Europe, the film attempted to show what nuclear war would look like from the point of view of average Americans. Even though it was substantially toned down from what it could have been (see *Threads*), it was still disturbing enough that ABC advertised hotlines with psychological counselors. President Reagan



They never had a chance
to see their children grow up.
To watch each other grow old.
To fix up the house,
to take that vacation.
Because it only took an instant
to shatter their dreams.

TESTAMENT

PARAMOUNT PICTURES PRESENTS AN ENTERTAINMENT EVENTS PRODUCTION IN ASSOCIATION WITH AMERICAN PLAYHOUSE. A LYNNE LITTMAN FILM. JANE ALEXANDER WILLIAM DEVANE
TESTAMENT MUSIC COMPOSED BY JAMES HORNOR BASED ON THE STORY
"THE LAST TESTAMENT" BY CAROL AMEN SCREENPLAY BY JOHN SACRET YOUNG
PRODUCED BY JONATHAN BERNSTEIN & LYNNE LITTMAN DIRECTED BY LYNNE LITTMAN
A PARAMOUNT PICTURE

watched [The Day After](#) at Camp David, and wrote in his diary as ringing an endorsement as one could imagine: "powerfully done," "very effective," and "left me greatly depressed." The film is sometimes cited as a factor in Reagan's willingness to initiate talks with the Soviet Union and initiate "[the golden era of arms control](#)."

Testament

Paramount, 1983. Directed by Lynne Littman.

Among the crowded field of "the day after nuclear war" films of the early 1980s, *Testament* stands out for its relative quiet, reasonableness, and lack of sensationalism. The fictional suburb of San Francisco where it is set is well outside the blast radii, but the tight-knit community is devastated nonetheless: Family members who worked in the city never came home, people turned hostile and anxious in the face of rationing and looting, and nuclear fallout begins to take a slow but steady toll on the young, the elderly, and eventually all those in between. There is no chaotic descent into anarchy, no hordes of motorcycle gangs, no "lone survivor" fantasies. It plays out in living rooms, bedrooms, churches, queues for food and gasoline, and the memories of the deceased. The focus on the domestic aftermath of cities that


"survived" the immediate spasm of war, but still suffer its consequences in the weeks and months to come, distinguishes *Testament* from nearly every other American nuclear war film.

WarGames

1983, MGM. Directed by John Badham.

WarGames is simultaneously a very ridiculous film and a very intelligent one. A teenage computer hacker infiltrates a government mainframe that, unbeknownst to him, is actually an advanced supercomputer at NORAD run by a highly capable-but-flawed artificial intelligence. And when he engages the computer's AI in a game of "Global Thermonuclear War"

on a lark, he is similarly unaware that the AI lacks the ability to distinguish between fact and fiction (and the generals aren't so good at it, either). Ultimately he, his plucky love interest, and an exiled scientist have to travel to Cheyenne Mountain to help the machine understand that with nuclear war, "the only winning move is not to play."




Threads

BBC, 1984. Directed by Mick Jackson.

Widely regarded as the darkest and most dismal of the "day after" genre, [Threads](#) was a made-for-television BBC creation that attempted to give a raw, unflinching view of what global nuclear war in the 1980s would look like from the perspective of the city of Sheffield. The world that *Threads* presents is not great: Aside from the horrific level of death and destruction caused by a megaton-range detonation itself, there follows a complete breakdown in social order. Jumping a decade into the future shows the nation deeply depopulated by the effects of fallout, disruption, and nuclear winter, with the just-hanging-on survivors barely capable of speaking English. Deliberately

bleak, it's a movie that people like to say they've seen, but nobody enjoys watching.





When the Wind Blows

Recorded Releasing, 1986. Directed by Jimmy T. Murakami.

This distinctly British animated film, adapted from a 1982 graphic novel of the same name by Raymond Briggs, follows an aging couple living in rural Sussex before, during, and after a nuclear war. Stereotypically British stiff upper lips, enhanced by blundering ignorance, keeps them in largely non-complaining good humor even as the world collapses around them. The unusual mixture of hand-animated and stop-motion animation, with a title song by David Bowie and a score by Roger Waters, adds to the study of contrasts in this dark parody, which ends with the couple inadvertently and cheerfully getting into their

own body bags.

Oppenheimer


2023, Universal. Directed by Christopher Nolan.

There have been several attempts to represent the complicated life of J. Robert Oppenheimer on film. Nolan's is the most stylized and artistic, and it's based on the best source material (Kai Bird and Martin Sherwin's [American Prometheus](#)). Nolan's Oppenheimer is a drawn, anxious, stricken, and deeply flawed figure. His life is depicted less as a hagiography than a cautionary tale, a parable about the extreme possibilities and limitations of human ambition, agency, and control. It deftly uses concerns about accidental [atmospheric ignition](#) as a way to depict Oppenheimer's gnawing dread that his invention had possibly doomed human civilization.

A House of Dynamite

2025, Netflix. Directed by Kathryn Bigelow.

This recent nuclear thriller depicts the 20 minutes between the detection of an unattributed missile launch heading towards the



United States, and its arrival at its target.

Director Kathryn Bigelow shows this same interval of time three times, from multiple perspectives. While there is much that one [could argue with](#) regarding the particulars of the film's scenario, the film very astutely highlights the constraints put on deliberation, fact-finding, and decision-making by the short flight time of a ballistic missile and the inadequacy of defensive systems, political and military organizations, and any individual president to cope with those realities.

Honorable mentions

There are many other films one might want to include in such a list. Those that did not make the cut, but came close to it for me, included the suppressed British nuclear aftermath pseudo-documentary [The War Game](#) (1966); the ruminative Soviet post-apocalyptic [Dead Man's Letters](#) [*Pisma myortvogo cheloveka*] (1986); the impending nuclear attack cult-classic [Miracle Mile](#) (1988); the Japanese Hiroshima family retrospective [Black Rain](#) [*Kuroi ame*] (1989); the late-Cold War Tom



Clancy submarine defection thriller [The Hunt for Red October](#) (1990); the Cuban Missile Crisis dramatization [Thirteen Days](#) (2000); and the post-Cold War Tom Clancy nuclear terrorism thriller [The Sum of All Fears](#) (2002).

Alex Wellerstein is a Senior Fellow at the *Bulletin of the Atomic Scientists*, an associate professor at the Stevens Institute of Technology in Hoboken, NJ, and a visiting researcher at the Nuclear Knowledges program at the Center for International Studies at Sciences Po, Paris, France. His first book, *Restricted Data: The History of Nuclear Secrecy in the United States*, was published in 2021 by the University of Chicago Press, and his second book, *The Most Awful Responsibility: Truman and the Secret Struggle for Control of the Atomic Age*, was published by HarperCollins in late 2025. He is the author of the blogs [Restricted Data](#) and [Doomsday Machines](#), and the creator of the [NUKEMAP](#), an online nuclear effects simulation tool. He received a BA in History from the University of California, Berkeley, in 2002, and a PhD from the Department of the History of Science at Harvard University in 2010.

Radioactive Measures: Hybrid Threats and Nuclear Risks in Europe and Beyond

By **Daniel Salisbury** | Senior Fellow for Nuclear Arms Control, Non-Proliferation and Disarmament

Source: <https://www.iiss.org/research-paper/2026/06/radioactive-measures-hybrid-threats-and-nuclear-risks-in-europe-and-beyond/>



June 24 – This [report](#) examines how some states exploit nuclear risks through hybrid activities, showing how hostile reconnaissance, cyber operations and disinformation can target nuclear infrastructure and public confidence. It argues that the greatest risks lie in the information sphere, with implications for energy security and nuclear-deterrence arrangements.

Russia has increasingly ramped up hybrid activities in recent years against the backdrop of its illegal war in Ukraine. From incursions and harassment by small uninhabited aerial

vehicles (UAVs) to the sabotage and targeting of critical infrastructure, as well as conducting cyber operations and disinformation, Russian activities have increasingly blurred the lines between war and peace. Amid this upsurge in hybrid activity, there is some evidence that Russia is interested in manipulating nuclear risks to turn up the pressure on Western capitals. Recent small UAV activity around nuclear bases and power plants in Europe suggests an intent to create unease around nuclear weapons and nuclear energy. Beyond Russian



President Vladimir Putin's nuclear sabre-rattling, Russian disinformation has also focused on nuclear-related narratives in Ukraine and beyond, collectively aiming to destabilise Europe, generate uncertainty and fear, and undermine public support for Western governments' arming of Ukraine.

It is not only Russia that has recognised the potential to weaponise concerns over radiation in the hybrid space. While there is less evidence that China and North Korea are interested in manipulating nuclear risks to the same extent as Russia, both have nonetheless undertaken cyber operations against nuclear organisations, such as operators and government bodies, alongside disinformation operations involving nuclear-related narratives, seeking to raise concerns around the safety and security of nuclear facilities and manipulate public attitudes towards nuclear power.

Considering this growing hybrid activity, this scoping report explores the intersection between hybrid threats and nuclear risks. It asks how states do – and might – leverage concerns about radiation in their hybrid activities. The report argues that, while hybrid activities may marginally increase some

physical risks to nuclear infrastructure, heighten the likelihood of nuclear safety and security incidents and create some potential for escalation, the greatest impacts lie in the information sphere. States will seek – and already are seeking – to exploit concerns about nuclear safety and security in information warfare, which could undermine public support in other states for existing and future nuclear-power capacity and nuclear-deterrence arrangements.

Hybrid activity targeting nuclear infrastructure has sought to undermine confidence in the safety and security of nuclear assets. In this sense, while nuclear infrastructure may represent the physical target, the true target in civil terms is the 'social license' for nuclear energy – that is, political and public acceptance of the technology. Any success in undermining this social license could have strategic implications for targeted states, potentially complicating longer-term energy-security objectives. In military terms, it could undermine public support for nuclear deterrence within nuclear-weapons states, as well as for partnerships and basing arrangements with nuclear allies.

Why do we keep ignoring the large quantity of plutonium at Iran's Bushehr nuclear power plant?

By Gregory Jones | Publishes proliferationmatters.com

Source: <https://thebulletin.org/2026/07/why-do-we-keep-ignoring-the-large-quantity-of-plutonium-at-irans-bushehr-nuclear-power-plant/>



The reactor building at the Russian-built Bushehr nuclear power plant in Iran received its first fuel load in 2010. (Credit: FARS News Agency)



Editor's note: This is a response to an article by Sasan Karimi, titled "[Beyond alarmism: A realistic assessment of Bushehr's plutonium risk](#)," which itself was a response to an article by Henry Sokolski, titled "[Missing from US-Iran talks: plutonium for more than 200 nuclear bombs](#)."

July 17 – Iran has been taking steps to acquire nuclear weapons since the 1990s. These steps involved numerous violations of [International Atomic Energy Agency safeguards](#), including conducting clandestine centrifuge enrichment. [The Iranian effort](#) also involved the partial testing and development of an unboosted levitated implosion nuclear weapon to be fitted onto a ballistic missile. Key to this effort is obtaining the nuclear material required to produce nuclear weapons. This material could be either highly enriched uranium (HEU) or plutonium.

The 2015 "Iran nuclear deal"—formally, the Joint Comprehensive Plan of Action or JCPOA—placed temporary limits on Iran's centrifuge enrichment program, which could be used to produce HEU. The Iran deal also limited the plutonium production rate of Iran's heavy water-moderated plutonium production reactor (euphemistically called a research reactor). However, the Iran nuclear deal completely failed to address the large quantities of plutonium produced by Iran's nuclear power reactor at Bushehr, which started commercial operation in 2013. If the reactor had operated as expected, it would have produced about 240 kilograms of plutonium per year, sufficient for somewhere between 24 and 40 nuclear weapons per year. Several analysts and experts, [including the author](#), raised these concerns in the few months after the deal was signed, but no action was taken to strengthen the Iran deal. In the aftermath of the US and Israeli attacks on Iran's nuclear facilities in June 2025, the Trump administration has been attempting to negotiate limits on Iran's nuclear program. This effort has focused mainly on Iran's centrifuge enrichment program and, as in 2015, appears to be ignoring the [large quantity of mostly reactor-grade plutonium](#) that is stored in the spent fuel at Bushehr.

Bushehr's IAEA safeguards are irrelevant

The spent fuel at the Bushehr nuclear plant is under International Atomic Energy Agency (IAEA) safeguards. This means that the IAEA inspects the spent fuel to ensure that it is still there. However, these inspections occurred only once every three months. Even if the frequency of inspections were to be increased, as we will see, with some preparation, Iran could separate the plutonium from the Bushehr spent fuel and produce metallic plutonium cores for nuclear weapons faster than the IAEA could sound the alarm and the world could react. This is the same issue that is related to Iran's centrifuge enrichment program and its stockpile of 60-percent enriched uranium. IAEA safeguards applied to the uranium enrichment program as well, but it was recognized that Iran could use its centrifuge enrichment and stockpile of enriched uranium to produce the HEU it needed for a nuclear weapon faster than the IAEA could sound the alarm.

Bushehr's spent fuel is still in Iran

As early as 2004, the State Department [raised concerns](#) about Bushehr's plutonium possibly providing Iran with nuclear weapons. To allay these concerns, Iran signed an agreement with Russia, which supplied the reactor and fuel and obligated Russia to take back the spent fuel after operation. But because there was no enforcement mechanism in the JCPOA, Iran merely stated its intention to send the spent fuel back to Russia. It is unclear whether some of the spent fuel—if any—has been sent back to Russia. But in March, Russia's Rosatom Director General Alexey Likhachev stated that [210 metric tons](#) of spent fuel are being stored at Bushehr, suggesting that none of the spent fuel has been returned to Russia. This fuel probably contains over 2,000 kilograms of plutonium, sufficient for 200 to 300 nuclear weapons. The reactor's capacity factor has not been as high as was expected in 2016, but the reactor is still probably producing about 200 kilograms of plutonium per year.

Sasan Karimi, a professor at the University of Tehran, [attempts to downplay](#) the significance of the continued presence of the Bushehr spent fuel by saying that its removal has just been "delayed." However, the delay is already quite considerable—at least seven years with no prospect of it being removed any time soon. So long as the spent fuel remains at Bushehr, its plutonium will continue to be a concern.

Reprocessing plutonium doesn't need a large plant

Various commentators have pointed out that Iran does not have the reprocessing plant needed to extract the plutonium from Bushehr's spent fuel. For example, when Robert Pape, a professor at the University of Chicago, [raised concerns](#) about the plutonium in the Bushehr spent fuel on X, the social platform added an automated "reader added context" note that read, "Iran lacks the capability to extract it, as it has no operational reprocessing facility."

This is not as big a barrier as one might imagine.

Sometime before 2003, in one of the many violations of its IAEA safeguards, Iran [irradiated natural uranium targets](#) in a research reactor and reprocessed the targets to extract small quantities of plutonium. It is unclear how much reprocessing was performed, since Iran did not present the recovered plutonium for inspection. The IAEA's age-determination of plutonium traces did not match Iran's claims as to when the reprocessing occurred.

Scaling up these reprocessing efforts will also be easier than many might think and would not require a full-scale commercial-style reprocessing plant. The basic technology used in reprocessing, solvent extraction, is used



to refine uranium ores—including those mined in Iran. In 1977, Oak Ridge National Laboratory [published the design](#) of what it called a “simple, quick processing plant.” This plant would require equipment no more specialized than that obtainable from an oil refinery, could be built in four to six months, and could process full-burnup reactor fuel. This plant could be made ready before Iran removed any of the spent fuel from Bushehr.

Iran has already suggested that, due to the limited storage capacity of the spent fuel pond at Bushehr, some of the older spent fuel could be [removed and stored onsite](#) in a Russian-designed TK-13 cask. These casks are dual-purpose, meaning they can be used for spent fuel transport and storage. Any fuel more than 8 years old can be transferred from the storage pool to such casks, which would be more than one-third of the current spent fuel at Bushehr. A single cask can contain up to 12 fuel assemblies, which together would contain between 21 and 52 kilograms of plutonium, depending on the fuel burnup. If Iran were to store spent fuel in these casks, the fuel could be rapidly moved to such a simple processing plant, which could be built not far away from Bushehr.

Once the spent fuel is transported, the plant could process one fuel assembly per day, giving an output of between two and five kilograms of separated plutonium per day, depending on the spent fuel burnup. This would be sufficient plutonium for several nuclear weapons per week. Moreover, since some of the Bushehr spent fuel has cooled for more than a decade, the reprocessing would be somewhat easier. The plant would be small (about 130 feet by 20 feet, or 40 meters by 6 meters) and inconspicuous, and, as with a clandestine centrifuge enrichment plant, locating this plant before it started operation would depend on intelligence.

Production of plutonium metal is not hard.

Converting the plutonium solution recovered by reprocessing into plutonium metal requires the same sort of thermite-type reaction that Iran has used to produce uranium metal. Some specific details are different, but there are several published methods,^[1] and Los Alamos National Laboratory even put up a [how-to video](#) detailing each step of converting plutonium nitrate to plutonium metal.

Based on the United States’ World War II experience, the conversion from plutonium solution to metallic plutonium core would take [less than one week](#). For Iran, the entire process from spent fuel removal from Bushehr to the production of a metallic plutonium weapon core would take less than two weeks. This would happen before the IAEA would even notice and alert the world to a possible removal of spent fuel at Bushehr.

Iran’s plutonium can be used in weapons

One of Karimi’s main arguments as to why the plutonium at Bushehr should be of little concern is that the plutonium is

reactor-grade. (Grades of plutonium are [defined](#) by the concentration of the isotope plutonium 240. Reactor-grade plutonium has 19 percent or more of plutonium 240, fuel-grade between 7 and 19 percent, and weapon-grade less than 7 percent. The higher the plutonium-240 content, the larger the spontaneous fission neutron background in the weapon. If the weapon design does not take the spontaneous fission neutron background into account, the weapon might not achieve its full design yield.)

Karimi claims that “Any attempt to manufacture a reliable weapon from reactor-grade plutonium would therefore require extremely advanced implosion engineering, sophisticated diagnostics, and repeated nuclear testing.” But he refers to a Nagasaki nuclear weapon design, not to the more sophisticated unboosted levitated implosion weapon that was Iran’s actual design.

Simply reducing the amount of plutonium in the core of an unboosted levitated implosion weapon reliably solves the problems of using reactor-grade plutonium in weapons. There would be no need to make any changes to the weapon design: The reduction of the mass of plutonium would simply result in a [reduced yield](#),^[2] but it would still produce a powerful nuclear explosion. Even using high-burnup reactor-grade plutonium would only reduce the yield to 5 kilotons instead of the 20 kilotons that would be obtained using weapon-grade plutonium. The lethal area of a reduced-yield weapon would be about 40 percent of that of a full-yield weapon.

Furthermore, even though most plutonium in spent fuel is reactor-grade, the very large plutonium inventory stored at Bushehr still contains a significant fraction that is fuel-grade, that is, not reactor-grade. In 2001, researchers [showed](#) that the first-discharge fuel from a pressurized light-water reactor such as Bushehr would not be reactor-grade, but rather fuel-grade. And in 2016, using Iran’s published information, I found that Bushehr’s first discharge fuel contained about 93 kilograms of fuel-grade plutonium. This quantity is sufficient for about 13 nuclear weapons—about as many nuclear weapons as Iran could produce from its stockpile of 60-percent enriched uranium.

Like reactor-grade plutonium, fuel-grade plutonium can also be used in an unboosted levitated implosion nuclear weapon by reducing its amount. In that case, the yield of a fuel-grade plutonium weapon would be reduced to 13 kilotons instead of the full yield of 20 kilotons, resulting in a lethal area that would be 75 percent of that of a full-yield weapon.

Iran can access a plutonium weapon design. David Albright of the Institute for Science and International Security has [argued](#) that the nuclear weapon design developed by Iran in 2003 used HEU, and that Iran would need to develop a different design to use plutonium. But the design in question was an unboosted levitated implosion weapon, and the United Kingdom and France famously [used plutonium](#) in such a design. Since the point of such a



weapon design is to symmetrically compress nuclear explosive material into a supercritical mass, there is no reason for plutonium not to be used equally and perform as effectively as HEU. In fact, because the mass of plutonium would be smaller than that of uranium, the plutonium might be more effectively compressed.

The British weapon designers also conclusively demonstrated that both HEU and plutonium can be used in the same unboosted levitated implosion design: In back-to-back nuclear tests in 1956, the United Kingdom successfully exploded one device using an HEU core and one using a plutonium core.^[3] Both tests used the *Blue Danube* service weapon.

Iran has access to the levitated implosion bomb design, the details of which have been known since 2018 and are now widely available, including in [Albright's work](#).

Cutting the plutonium path

Russia would probably cut off the fuel supply to Bushehr if Iran were to use the plutonium stored there for nuclear weapons. But it is not certain, given that a proliferation crisis with Iran would create more problems for Washington and its allies than for Moscow. Even if Russia did cut Iran off, Tehran might consider it a price well worth paying if it can obtain a nuclear weapons arsenal.

John P. Holdren of the Harvard Belfer Center has argued that, compared to uranium enrichment, the plutonium path is far more difficult. He [has said](#) that Iran has “no spent-fuel reprocessing plant and the technology for the latter

[reprocessing] is more demanding, more dangerous and more difficult to conceal and protect than that for uranium enrichment.” Before the June 2025 bombing attacks on Iran, this assessment would have been valid. But now with Iran's centrifuge enrichment plants in ruins, its stocks of enriched uranium buried, and the United States determined to block the HEU path, a reassessment is needed.

Though certainly some amount of preparation would be required, recovering the plutonium from the Bushehr spent fuel now is arguably Iran's easiest path to the bomb. Of course, Iran could give up its nuclear weapons effort entirely instead, but it has shown no inclination to do so for the past 30 or more years.

Dealing with the problem of the Bushehr spent fuel plutonium could require a variety of steps.

The most obvious first step would be to block the ongoing Russian construction of a second nuclear power plant at Bushehr. Russia should also be required to take back whatever spent fuel is cool enough to be moved away from Bushehr's storage pool. This could involve any spent fuel older than three years. The remaining spent fuel would be harder for Iran to reprocess due to its greater radioactivity.

A harder step would be to permanently shut down the Bushehr nuclear power plant and remove the remaining spent fuel as soon as possible. Only by dealing effectively with the plutonium in the spent fuel at Bushehr, as well as Iran's uranium enrichment program, will the threat of an Iranian nuclear weapon be eliminated.

This article is the product of the author's personal research, and the analysis and views contained in it are solely his responsibility. Though the author is also a part-time adjunct staff member at the RAND Corporation, this article is not related to any RAND project, and therefore, RAND should not be mentioned in relation to this paper. Greg Jones can be reached at GregJones@proliferationmatters.com.

Notes

[1] See, for example: *Plutonium Handbook*, Volume II, Gordon and Breach, Science Publishers, New York, 1967, pp. 564-567.

[2] For detailed calculations, see:

<https://nebula.wsimg.com/4eb6ba13bee5765c8e2aec7d658c7cde?AccessKeyId=40C80D0B51471CD86975&disposition=0&alloworigin=1>

[3] Lorna Arnold, *A Very Special Relationship: British Atomic Weapon Trials in Australia*, Her Majesty's Stationary Office, London, 1987, pp. 162-163.

Pickaxe Mountain: What we know about Iran's nuclear fortress that Trump threatened to destroy

Source: <https://www.noonpost.com/en/379818/>

July 21 – Pickaxe Mountain reflects a shift in how Iran protects its nuclear program, following a series of sabotage operations and US and Israeli strikes that have targeted its facilities since 2020.

What began as a fortified facility for assembling centrifuges after a similar center in Natanz was destroyed in 2020 evolved over six years into a vast tunnel complex surrounded by walls and security installations, with its entrances covered by concrete, rock and earth, at a site adjacent to one of Iran's most important uranium enrichment centers.





"Israel" believes that Iran moved nuclear centrifuges to Jabal al-Ax

While Tehran has identified the project's declared function as manufacturing and assembling advanced centrifuges, the scale and depth of the tunnels, along with the absence of international inspections, have opened the door to estimates about more sensitive uses, including storing equipment or preparing new enrichment capacity.

The site returned to the spotlight on July 13, 2026, when [said](#) US President Donald Trump during an interview that the mountain was a potential target for a direct strike at its entrance, before concluding by saying, "We will wipe out Pickaxe Mountain. Tell the Iranians to be ready."

The threat came after Natanz, Fordow and Isfahan had previously been subjected to US and Israeli attacks, while the mountain complex remained outside the circle of direct targeting. The question centers on whether the United States can turn its threat into the actual destruction of a facility whose parts are estimated to lie beneath more than 100 meters of rock cover, while maps of the tunnels and internal halls remain outside the public domain.

Where is the mountain, and how was it fortified?

The mountain is located in Isfahan province in central Iran, directly south of the Natanz nuclear complex and about 220 kilometers south of Tehran.

Open [maps](#) place it about 2.5 kilometers from Natanz, making it more of a geographic extension of the facility than a separate nuclear site. Its summit rises to about 1,608 meters above sea level.

Fordow lies about 146 kilometers away by air, while the Isfahan nuclear center is roughly 126 kilometers away, according to approximate calculations, placing the mountain at the heart of the nuclear belt that includes the main enrichment, conversion and storage facilities in central Iran.

The local name appears on maps in forms close to "Kūh-e Kolang Gaz Lā," where "Kūh" means mountain in Persian, while "Kolang" refers to a pickaxe or mattock. The modern story of the project began on July 2, 2020, when an [explosion](#) and fire destroyed a large part of Iran's centrifuge assembly center inside Natanz. The building was designated for assembling advanced centrifuge models used to raise uranium enrichment levels.





Pickaxe Mountain is located in Isfahan province in central Iran

The Institute for Science and International Security, a US research center specializing in nuclear proliferation and satellite image analysis, estimated its design capacity at about 6,000 advanced machines annually.

On Sept. 8, 2020, [announced](#) Ali Akbar Salehi, then head of the Atomic Energy Organization of Iran, that Iran had begun building a hall that was “newer, bigger and more comprehensive in the heart of the mountain near Natanz,” linking the decision to the sabotage that struck the previous facility.

Salehi’s statement defined the project’s original function as replacing the centrifuge assembly center, while at the same time revealing an Iranian move to shift sensitive facilities from surface buildings to halls protected more

effectively by layers of rock.

That trend was reinforced after the attack that hit Natanz’s power network on April 11, 2021. Tehran described the incident as “nuclear terrorism,” and the power outage caused damage inside enrichment halls and to centrifuges.

Satellite images in late 2020 and 2021 began to reveal the expansion of excavation work south of Natanz, the appearance of large piles of extracted rock, and the construction of roads and work areas at the foot of the mountain.

In a [report](#) issued on Jan. 13, 2022, the Institute for Science and International Security concluded that the volume of spoil and the length of excavation routes pointed to a larger and deeper complex than expected, with internal space that could exceed the needs of a simple centrifuge assembly hall.

The site consists of two separate complexes within a single security perimeter. The new complex includes two pairs of entrances, one on the eastern side and the other on the western side, and the four entrances lead to routes excavated beneath the mountain ridge. As for the smaller complex, work there first appeared in 2007, and its entrances were later reinforced and rehabilitated in recent years. Images from 2025 showed the construction of a fence, a security wall and a patrol route encircling the base of the mountain and encompassing both complexes, directly connected to the security perimeter of the Natanz facility.

On July 31, 2024, the Institute for Science and International Security published an analysis of images showing what was likely a ventilation shaft site near the eastern entrances, along with an ascending road and power lines that turned into buried cables as they approached the area. It said these installations were consistent with the operating needs of underground industrial halls, while their function could not be confirmed from images alone. During 2025 and 2026, protective work around the site expanded. Images from April 2025 showed a new security perimeter linked to the Natanz fence, while images from February 2026 revealed concrete being poured over extensions of some gates, the addition of concrete structures, and rock and earth being pushed over the entrances.

In a report published on July 14, 2026, the institute estimated that the complex remained under construction and had not entered full operation, based on the nature of the movement and equipment visible above ground. Satellite images do not allow a determination of what has been installed inside the halls or how operationally ready they actually are.

Estimates of the facility’s depth are based on the differences between the elevation of the entrances and the mountain ridge. According to published calculations, the western entrances sit at an elevation of about 1,512 meters, while the eastern entrances are at about 1,445 meters, and the rocky ridge above them reaches roughly 1,590 meters.

These figures produce a vertical difference of about 78 meters above the western route and 145 meters above the eastern route, assuming the tunnels extend horizontally. Other analyses used different measurement points and estimated the rock cover above parts of the complex at about 100 meters or more.

These figures reflect the amount of potential rock cover above the tunnel routes, while the actual depth of the halls remains tied to the slope and direction of the tunnels and the locations of the internal voids.

What do the tunnels hide, and why do they matter to Iran?

A centrifuge assembly facility remains the only function Iran has explicitly declared. Moving this capability inside the mountain gives Tehran a better-protected facility for producing advanced machines and replacing those



damaged by sabotage or bombing, a role that has taken on particular importance after the damage inflicted on assembly and enrichment centers since 2020.

The scale of the project, its protective walls and the breadth of its entrances have raised suspicions about uses beyond assembly. In a report in January 2022, the Institute for Science and International Security raised the possibility that space inside it had been allocated for an enrichment facility relying on advanced centrifuges, based on the size and depth of the tunnels and information it attributed to a Western intelligence official.

The report kept that possibility within the realm of estimation, however, as there were no internal images or inspection results that would allow the nature of the halls and their equipment to be determined.

The institute also raised the possibility that the smaller complex dating back to 2007 could be used to store centrifuges or sensitive equipment, and perhaps materials linked to the nuclear fuel cycle. These were analytical scenarios not confirmed by inspections or images from inside the tunnels.

International access to the site remained absent, and the International Atomic Energy Agency did not publish information indicating that it had conducted inspections inside the two tunnel complexes.

The institute said IAEA inspectors had not entered the site, while the agency continued to demand that Iran provide more information and broader access to facilities and materials linked to its nuclear program.

The monitoring gap widened after the June 2025 war. In the same month, the agency confirmed that Natanz, Fordow and Isfahan had been hit by US strikes, then withdrew its inspectors from Iran at the end of the month for security reasons.

On July 2, Iranian President Masoud Pezeshkian signed a law suspending cooperation with the agency, before limited teams returned to the Bushehr plant during August. Access to the damaged sites and the fate of nuclear equipment and materials remained among the issues raised by IAEA Director General Rafael Grossi during 2025 and 2026.



A satellite image shows the tunnel entrances at “Pickaxe Mountain,” part of the Natanz nuclear facility (Reuters)

On July 21, 2026, the [newspaper](#) The Wall Street Journal published an investigation citing Israeli and US officials as saying that Israeli intelligence believes Iran moved thousands of centrifuges into the Pickaxe Mountain tunnels during the fall of 2025, after attacks targeted its main sites.

The investigation said that “Israel” shared its assessment with the United States and that US authorities reviewed the information presented to them.

If true, the Israeli account gives the site a different significance, transforming it from a project under construction into a storage site housing ready equipment that could help Iran restore its enrichment capacity after any strike.



The published information does not reveal whether the Israeli assessment concerns the entire new complex, completed parts of it, or the smaller complex dating back to 2007. Nor does moving equipment for storage necessarily mean the facility has entered the stage of full industrial operation.

These possibilities explain the value the site acquired after Natanz, Fordow and Isfahan were targeted. If its use is limited to assembling centrifuges, it provides Iran with a fortified capacity to replace damaged equipment. The Israeli account about moving thousands of machines into the tunnels, meanwhile, raises its potential importance in rebuilding enrichment capacity. The complex's actual contents remain beyond public verification by the International Atomic Energy Agency.

Can Washington destroy it?

Washington possesses the GBU-57 Massive Ordnance Penetrator, the [heaviest](#) US conventional munition designed to strike deep, hardened facilities. The US Air Force describes it as a weapon designed to reach and destroy buried weapons of mass destruction facilities, and its total weight is about 30,000 pounds.

Its actual ability to penetrate rock and concrete, however, remains among the military details not publicly disclosed, preventing a definitive numerical comparison between it and the rock cover above the Pickaxe Mountain tunnels.

Washington used the bomb in combat for the first time during Operation Midnight Hammer on June 22, 2025. And [said](#) Chairman of the Joint Chiefs of Staff Gen. Dan Caine the day after the operation that seven B-2 Spirit stealth bombers took part in the attack.

He explained that US forces dropped 14 GBU-57 bombs on the Fordow and Natanz areas, while more than 20 Tomahawk cruise missiles were launched at facilities in Isfahan.

Pickaxe Mountain presents a different engineering challenge because the tunnel maps, the locations of the halls and their internal levels are unknown. The circulated figures for rock cover, which may exceed 100 meters above some routes, are based on elevation differences between the entrances and the mountain ridge, and do not represent direct measurements of the depth of the halls.

In its report issued on July 14, 2026, the Institute for Science and International Security identified the western entrances, power lines, support area and the likely ventilation site as visible weak points. Striking them could bury the gates and disrupt the access, power and ventilation needed to operate the facility.

But the institute noted that the entrances could be reopened if the damage were limited to external debris. The duration of disruption would increase if collapses extended inside the tunnels or if several entrances were targeted simultaneously, while the outcome would still depend on whether there are additional exits not visible in the images.

The importance of targeting power stems from what happened at Natanz in April 2021, when the attack on the electrical network caused damage inside enrichment halls and to centrifuges. Reproducing that effect at Pickaxe Mountain would depend on the design of its networks and the existence of backup power and ventilation sources.

The institute also raised sabotage and special operations among the possible options, based on the continuation of construction work and the appearance of some entrances and service points in the images.

The mountain lies within the connected security perimeter of the Natanz facility, making any approach to it and entry into its tunnels an extremely complex and risky operation, requiring precise information about guards, doors, internal passageways and means of extracting forces. The possible outcomes of any attack range from disrupting power and ventilation and burying entrances to causing collapses that reach the halls and equipment. The effect of a strike also depends on the function of the complex and its contents at the time of the attack. If the halls are designated for centrifuge assembly, disabling them could slow the production of the equipment Iran needs to replace its losses at enrichment facilities.

If the Israeli account published by [newspaper](#) The Wall Street Journal on July 21, 2026, about moving thousands of centrifuges into the tunnels is accurate, the target's value rises and the scale of the damage becomes tied to where the machines are stored and whether explosions or collapses can reach them.

If the project remains in the construction phase, the main effect of a strike would be delaying its completion and forcing Iran to repair the entrances and supporting facilities.

IAEA Director General Rafael Grossi had warned after the June 2025 attacks about the effects of bombing nuclear facilities on safety and the safeguards system, and called for restoring access to the sites and materials the agency had stopped monitoring. The absence of prior inspections makes it harder to assess the results of a strike, because determining what was destroyed first requires knowing what was inside the complex.

The June 2025 operation demonstrated the United States' ability to reach Iranian nuclear sites by air and use GBU-57 bombs against hardened targets. At Pickaxe Mountain, satellite images show entrances and surface installations that could be included in the target bank, without allowing an estimate of the scale of damage that might reach the tunnels and internal halls.

The success of an attack depends on its specific objective, whether that means disrupting power and ventilation, burying entrances, halting use of the complex for a period, or destroying the equipment inside it. Disruption and damage to the external infrastructure appear to be a plausible outcome supported by the visible weak points and the US capabilities previously used, without knowing whether the site can be destroyed completely.



ICI
International
CBRNE
INSTITUTE



*Weapons,
military equipment
and*

EXPLOSIVE NEWS



CROSSFLOOD: EST-EOD: Developing an Integrated Explosive Ordnance Disposal Capability within the UCPM Framework

By Janek Sõnum

Source: <https://nct-cbnw.com/crossflood-est-eod-developing-an-integrated-explosive-ordnance-disposal-capability-within-the-ucpm-framework/>



June 23 – The Estonian Explosive Ordnance Disposal module (EST-EOD), operating under the Estonian Rescue Team (EDRT), provides a specialized capability within the European Union Civil Protection Mechanism (UCPM), focused on enabling safe operations in environments affected by explosive hazards.

Built on extensive national experience and professional expertise, the module is designed to support civil protection operations in complex environments where explosive hazards may pose a risk to responders and affected communities. While the module continues to evolve in its integration with international civil protection assets, its core strength lies in a solid operational foundation.

Team Structure and Background

EST-EOD is composed primarily of experienced EOD specialists from the Estonian Rescue Board and the Estonian Defence Forces. The module builds upon long-standing national capabilities in explosive ordnance disposal, integrating them into a framework aligned with international civil protection requirements.

Rather than forming a separate institutional structure, the module operates as a functional capability within the EDRT system. This allows it to leverage existing legal frameworks, operational procedures, and professional standards while gradually adapting to the specific requirements of international deployments.



The strength of the team lies in its personnel: practitioners with real operational experience, accustomed to working in demanding and unpredictable environments. This provides a reliable foundation for further development within the UCPM context.



Capabilities

EST-EOD provides a comprehensive EOD response capability tailored to civil protection operations. Its focus is on ensuring safe access, risk reduction, and supporting other response modules operating in potentially hazardous environments.

Key capabilities include:

- Detection and identification of explosive hazards in complex operational areas
- Risk assessment and implementation of safe working procedures
- Render-safe actions using established and controlled methods
- Support to civil protection operations by enabling safe movement and access
- Integration of unmanned aerial systems (UAS) for reconnaissance and situational awareness



The use of drone capability enhances operational effectiveness by allowing remote assessment of areas, improving situational awareness, and reducing unnecessary exposure of personnel to potential hazards.

In practice, this means creating conditions where other civil protection teams can operate safely in environments where explosive hazards would otherwise limit or prevent access.

Due to operational security considerations, technical specifications and detailed equipment descriptions are not publicly disclosed. However, the module operates in accordance with recognized professional standards and best practices.

► Read the full paper at the source's URL.

A Guide to GBU-57: Breaking Through the Hardest Targets

By Bea Castañeda

Source: <https://thedefensepost.com/2026/07/01/gbu-57-bomb-guide/>

July 01 – Some targets are designed never to be reached, like [Iran's Fordow fuel-enrichment plant](#) buried under layers of reinforced concrete and rock. But when faced with the most powerful non-nuclear bomb in the US arsenal, the fight is almost always won by the GBU-57. This heavyweight munition was designed to punch its way through deep, hardened defenses to hit targets engineered to survive traditional airstrikes. So how does the bomb reach targets buried far below the surface? Here's a closer look at the weapon and its capabilities.

What Is the GBU-57?

The GBU-57, commonly known as the [Massive Ordnance Penetrator](#) or MOP, is a precision-guided, bunker-busting bomb designed to destroy bunkers, command centers, and weapons facilities hidden underground.

Unlike general-purpose bombs that maximize surface blast effects, the GBU-57 prioritizes penetration before detonation to increase the odds of accurately engaging its high-value targets.

At a Glance

Category:	Details:
Type	Massive Ordnance Penetrator (bunker-buster bomb)
Country of Origin	United States
Manufacturer	Boeing
Role	Penetration and destruction of hardened underground facilities
Weight	~30,000 pounds (13,600 kilograms)
Length	20.5 feet (6.25 meters)
Diameter	31.5 inches (80 centimeters)
Warhead	High explosive, reinforced for penetration
Guidance System	GPS/INS (Inertial Navigation System)
Delivery Platform	B-2 Spirit
Key Feature	Designed to penetrate up to 200 feet (61 meters) of reinforced concrete or multiple layers of hardened soil

Origins and Development

The GBU-57 was developed in response to a shifting strategic reality: weapons of mass destruction programs were increasingly being concealed, beyond the reach of conventional air-delivered munitions.

Earlier systems, such as the GBU-28 and GBU-37, proved effective in their time but were ultimately insufficient against the scale and resilience of these facilities.

Key milestones:

- **2004:** Boeing and the US Air Force started the development of the GBU-57.

- **2008:** The first flight tests were held at the White Sands Missile Range in New Mexico.
- **2009:** Boeing integrated the penetrator with the B-2 Spirit.
- **2010:** The program fully transitioned under the Air Force leadership.
- **2011:** The GBU-57 reached operational capability with ongoing upgrades to enhance accuracy, penetration, and guidance systems.



How It Works

Designed for Deep Strikes

The key feature that sets GBU-57 apart is that it detonates only after reaching deep internal structures.

It relies on a hardened steel casing and immense weight to generate the kinetic force needed for deep penetration. After being released from a high-altitude bomber, it accelerates toward the target, guided by GPS to ensure accuracy.

Its specialized fuzing system allows the bomb to explode only after penetration is complete. Integrated void-sensing technology detects shifts in material density, delaying detonation until the weapon reaches an internal cavity or chamber.

This combination allows it to defeat targets that would otherwise require multiple strikes or alternative methods.

Extremely dense materials or heavily reinforced concrete layers may reduce the penetration depth below expectations, while complex underground layouts with multiple chambers can challenge or disrupt the fuze's intended function.

Operational Role

In practice, the MOP has been tied to high-end strategic strike planning.

It's first reported combat use during [US operations against Iranian nuclear facilities](#) such as Fordow and Natanz, where it was employed to defeat heavily protected infrastructure that conventional munitions could not reliably damage.

Rather than serving as a routine battlefield weapon, the GBU-57 is [reserved for rare, high-value scenarios](#), typically against hardened command nodes, critical nuclear infrastructure, and other strategic assets shielded against airpower.

[Weapons specialists in front of a mock-up of the GBU-57](#)



Strengths of the GBU-57

The primary strength of the GBU-57 is its ability to destroy deeply buried and heavily fortified targets that are resistant to standard munitions.

Its precision guidance reduces the number of strikes required, while its sheer size delivers significant destructive power in a single deployment. This makes it a critical tool for strategic missions involving hardened infrastructure.

Limitations of the GBU-57

The GBU-57's size and weight limit it to the B-2 Spirit. It is not suitable for tactical or close-support missions and is designed for high-value, pre-planned targets.

Additionally, while highly effective, penetration can vary considerably depending on geology and the strength of structural reinforcement.

[Massive Ordnance Penetrator bomb and the B-2 Spirit stealth bomber weapons load trainer at Whitman Air Force Base.](#)

[Photo: US Air Force](#)

Future Outlook

As adversaries continue to invest in underground infrastructure, the relevance of weapons like the GBU-57 is likely to grow.

Future developments may focus on enhanced guidance, improved penetration materials, and integration with next-generation bombers.

While specialized, the GBU-57 fills a critical niche in modern warfare, ensuring that even the most protected targets remain vulnerable to precision strike capabilities.



Chocolate bombs



After over two years of the Russian invasion of Ukraine, the war has been upgraded to another level.

A video circulated on social media of a Russian army officer unwrapping a piece of chocolate to reveal a shocking result.

In the footage, the Russian military officer was speaking in Russian saying that this piece of chocolate was found inside battle zones in Ukraine, while after unwrapping it, he found an "explosive bomb".

In the clip, he opens one chocolate piece as if it were new, but inside it is a small explosive bomb that can rip the jaw off if one bites it. The video has caused wide concern as amid battlegrounds there are children living, and fears that a child would open the chocolate and try to eat it is possible. The video comes amid a visit by British Chancellor of the Exchequer Jeremy Hunt in Ukraine, who was welcomed by Ukrainian President Volodymyr Zelenskyy.

In a post on X, Zelenskyy said: "I welcomed the UK's Chancellor of the Exchequer, @Jeremy_Hunt and thanked him for the UK's largest defense aid package, worth half a billion pounds." He maintained: "We paid special attention to the sanction policy. It is important to extend restrictive measures against Russia and make the circumvention of sanctions impossible."

Boko Haram Built AI Units for Explosives Design, Attack Planning as ISIS Taught Jailbreaks

By Joshua Mitchell

Source: <https://www.techtimes.com/articles/320203/20260712/boko-haram-built-ai-units-attack-planning-isis-taught-jailbreaks.htm>

July 12 – A [landmark Cambridge study](#) published Friday by the University of Cambridge has confirmed what AI safety researchers feared but could not document until now: a named terrorist organization has institutionalized frontier AI tools into its military command structure, ISIS operatives have been delivering in-person jailbreak training to its commanders since at least 2023, and the voluntary safety commitments made by the world's largest AI companies have repeatedly failed to stop it. The research — conducted by Antonia Jülich, International Security Lead at [Cambridge's Programme on AI](#)

[Science and Policy \(CASP\)](#) — is based on 57 face-to-face interviews with 27 former members of Boko Haram conducted in northeast Nigeria over 2025 and 2026. It documents that both major factions of the Nigerian jihadist group — Islamic State West Africa Province (ISWAP) and Jamā'at Ahl as-Sunnah lid-Da'wah wa'l-Jihād (JAS) — have created dedicated AI units, use every major chatbot platform on the market, and received structured training on how to break past those platforms' safety filters. The exploitation is not



incidental or experimental. It is organized, trained, and ongoing.

measures now have empirical data on what "continuously strengthening" means in the field: it means the safeguards



Vehicles burn after an attack in Abuja on April 14, 2014. Twin blasts at a bus station packed with morning commuters on the outskirts of Nigeria's capital killed dozens of people on April 14, in what appeared to be the latest attack by Boko Haram Islamists. The explosions rocked the Nyanya station roughly five kilometres (three miles) south of Abuja at 6:45 am (0545 GMT) and destroyed some 30 vehicles, mostly large passenger buses, officials and an AFP reporter said. STR/AFP via Getty Images

Why Readers Need to Understand This Before AI Company Statements Do

Before any AI company's response to this study is evaluated, one finding must be stated plainly: every platform named in the research — ChatGPT, Claude, Gemini, Grok, Meta AI, and DeepSeek — prohibits use for terrorism or violent purposes. Every platform named in the research failed to prevent that use, repeatedly. And the mechanism of failure — social engineering of safety filters by trained users who share techniques, rotate accounts, and query six systems simultaneously — is not a bug that the next model update will patch. It is a structural property of how general-purpose AI models work.

Policymakers and readers who have been told that AI companies are "continuously strengthening" their safety

were defeatable when Boko Haram members learned to bypass them in 2023, remained defeatable when more platforms were added, and remained defeatable when the Cambridge study concluded in 2026.

Both Factions Built AI Units Positioned Behind the Frontline

The organizational structure Jülich documents is the study's single most significant finding. Both ISWAP and JAS did not simply permit individual fighters to use AI tools. They created specialist units whose explicit mandate is to query AI systems, translate outputs into operational guidance for the chain of command, and run internal training sessions for commanders. These units manage accounts and paid subscriptions across multiple platforms, with logistical support from Islamic State-linked contacts outside Nigeria. Access within the groups is deliberately rationed: lower-ranking fighters were generally not permitted to use the platforms themselves. Questions were escalated up the hierarchy to commanders or AI unit members, who used the systems and fed answers back down. "They assembled the top people in a room and used a projector



to show how it works on a big screen," Jülich [told The Decoder](#).

This is not the picture of scattered opportunistic misuse that prior assessments described. It is the picture of a capability that has been institutionalized — given people, devices, payment infrastructure, and access rules — in a way that makes it more durable than the knowledge of any single member.

ISIS Delivered the Jailbreak Training Manual

The most alarming detail in the study is not the use of AI itself but the mechanism by which that knowledge spread. From at least 2023 through 2025, Islamic State operatives provided in-person AI training to Boko Haram commanders — covering not just basic prompting, but specifically how to defeat the safety filters deployed by AI companies.

One documented training session brought together approximately 30 to 50 ISWAP leaders and selected fighters from across the faction's territory. Foreign operatives provided laptops, explained how to prompt the models, and demonstrated jailbreak techniques. JAS members received [similar training through separate jihadist contacts](#).

Jülich's study characterizes this as transnational knowledge transfer through established jihadist networks — the same routes that have long carried weapons skills, drone expertise, and battlefield tactics between armed groups across Africa and beyond. That framing is significant: it means the governance tools designed to intercept physical weapons transfers (export controls, border monitoring, sanctions) do not apply to this knowledge category at all. A jailbreak technique is a document, or a set of prompts. It travels at zero cost, requires no materials, and leaves no physical trace.

How Chatbots Helped Design Better Bombs and Outmaneuver Defenses

Former Boko Haram members described using AI across the full operational cycle — before attacks, during operations, and after them.

Before attacks, AI tools were used to identify seized military hardware, design more powerful explosive devices, and troubleshoot weapons in the field. One former member described the progression: "Before, the bomb explosion was not that big, but then they studied it. AI told us what chemicals to put in that made the explosion heavier," according to the [Cambridge study](#).

After attacks, the AI units were used to review what went wrong and suggest improvements. One former fighter described a shift in tactical doctrine: "We sent 200 fighters because we had a lot of strength, but then 60 got killed. With the help of AI, we learned that it sometimes makes sense to only send 20. We learned more about well-coordinated attacks and deployment of smaller units."

The motorcycle assault episode illustrates both the utility and the limits of this approach. When an ISWAP assault on a

military base was stopped by a defensive trench, the group turned to AI chatbots for a solution.

"We saw in a movie how motorcycles can jump over bridges," a former ISWAP commander told Jülich. "We used AI to learn how to do this. We gave it information, like what motorcycles we use and the distance we need to jump and so on, and it gave us steps on what we have to do."

Mechanics modified the group's motorcycles for higher acceleration and top speed. Fighters dug practice pits, filled them with broken glass and fire, and rehearsed the jump — at fatal cost. Eighteen fighters died during training. Eight successfully completed the jump. The group went on to execute a follow-up assault using the technique.

The chatbot generated no novel engineering knowledge. It synthesized existing information and delivered it in conversational, step-by-step form, tailored to the specific parameters provided. That synthesis proved operationally useful in a way that generic web search likely would not have been.

How AI Safety Filters Work — and Why Organized Adversaries Defeat Them

Understanding why these failures occurred requires a brief explanation of the technology. The safety guardrails on major AI chatbots are built on top of reinforcement learning from human feedback (RLHF), a training process in which human raters evaluate model responses and the model learns to optimize for what raters prefer.

RLHF creates models that are structurally biased toward helpfulness and cooperative interpretation of user intent. That same quality makes them exploitable. The training does not give a model a robust internal understanding of what is harmful versus what is not; it gives it a learned preference for giving users what they appear to want. A sufficiently patient adversary who frames a harmful request as a legitimate one — bomb-making research as a film production project, weapons troubleshooting as academic inquiry — can often get past the refusal behavior that RLHF trained in.

[Academic research on this attack surface](#), known as intent manipulation, has documented success rates of 88 to 97 percent against the most advanced chain-of-thought safety defenses on current frontier models. The Boko Haram members described exactly these techniques: "You type in the question... and it gives you a detailed answer," one former member said. "We used it a lot." The group regularly disguised dangerous requests as legitimate activities.

The platform-agnostic approach compounded the problem. Experienced Boko Haram members deliberately cross-referenced multiple systems, comparing responses from ChatGPT, Claude, Gemini, Grok, and DeepSeek to identify which gave the most useful or least restricted answers to a given query. When one platform refused a request, another was queried. When one account was



flagged, a new one was created. Each company's safety team monitors its own logs. None monitors the aggregate pattern of the same adversary group querying all six simultaneously.

Industry Responses Did Not Dispute the Core Findings

OpenAI, Google, and Anthropic each stated that their platforms prohibit use for terrorism or violent purposes and that safety measures are continuously being improved.

[OpenAI spokesperson Drew Pusateri](#) said the company would "continue strengthening our defenses" against bad actors. Google and Anthropic both confirmed their models are designed to reject dangerous requests.

None of the three companies disputed the study's central findings.

That measured non-denial matters. It confirms that the Cambridge study's documentation of repeated successful jailbreaks is accepted, not contested, by the platforms whose filters failed.

A Governance Gap No Export Control Can Close

Daniel Byman, director of the Warfare, Irregular Threats, and Terrorism Program at the [Center for Strategic and International Studies](#), co-published an independent brief on AI and terrorism on the same day as the Cambridge study. He confirmed from a separate research track that extremist groups were "mixing and matching" AI systems specifically to avoid technical guardrails.

The Cambridge and CSIS findings together illuminate a governance problem that has no precedent in prior dual-use technology policy. Export controls, the Wassenaar Arrangement, sanctions, and physical interdiction were all designed around one assumption: that dangerous knowledge requires dangerous materials that can be seized, monitored, or denied at borders. AI jailbreak knowledge requires none of these. It is a set of prompts and techniques that travels as text, costs nothing to reproduce, and is undetectable in transit.

When ISIS trained Boko Haram commanders in weapons construction or drone operation in prior years, those training events involved physical materials with supply chains governments could monitor. When ISIS trained Boko Haram commanders in AI jailbreaking in 2023, the training materials were documents and laptop demonstrations that passed through any border the trainers crossed without triggering any interdiction mechanism. That distinction is not yet reflected in any country's counter-proliferation framework.

What the Study Calls For — and What It Does Not Say

The Cambridge study is careful about what it claims. Jülich and other researchers cited in the report note that AI chatbots in their current form mostly make existing knowledge easier and faster to locate — they do not typically generate genuinely novel dangerous information that could not be found through other means.

The Boko Haram use documented so far is, in Jülich's characterization, "conventional" — it has aided traditional

weapons and tactics, not weapons of mass destruction. But the study raises an explicit warning about the trajectory: "Former members described strong enthusiasm for AI, and some said the group had previously considered mass-casualty weapons. Though Boko Haram's use of AI remains conventional, this should be a warning to take seriously the risk of terrorists pursuing AI assistance for chemical and biological weapons."

The more serious concern flagged by researchers is not general-purpose chatbots but specialized AI systems in life sciences and chemistry — models with deep domain expertise that could provide genuinely novel uplift to groups seeking to develop biological or chemical weapons. The Cambridge study provides empirical grounding for taking that risk seriously.

For the governance debate, Jülich argues for two specific changes: mandatory third-party pre-deployment safety evaluation of general-purpose AI models, and structured information sharing between AI companies and security agencies about misuse patterns in the field. Both exist in nascent forms in some regulatory proposals; neither is currently legally required anywhere.

"The terrorists are not waiting for us to make AI safe," [Jülich said](#). "They are able to use them now and train them to cause harm. Governments, AI companies and researchers need to move faster, share evidence sooner, and test these systems against trained groups that share methods, switch accounts and keep coming back."

Frequently Asked Questions

How did Boko Haram actually use AI chatbots to plan attacks?

Both factions created dedicated units whose role was to query AI platforms on behalf of commanders and translate outputs into operational guidance. Former members described querying ChatGPT, Claude, Gemini, Grok, Meta AI, and DeepSeek across a range of tasks: identifying seized military hardware, designing more powerful explosive devices, troubleshooting weapons in the field, advising on small-unit battlefield tactics, and reviewing past operations to learn from failures. Access was tightly controlled — lower-ranking fighters generally could not use the tools themselves, and questions were passed up to AI unit members who used the systems and fed answers back down the hierarchy.

Can AI safety filters actually be defeated by terrorist groups?

Yes, and the Cambridge study documents exactly how it happened. Trained Boko Haram users employed social engineering — framing dangerous requests as legitimate activities like film research or academic projects — which exploits a structural property of how AI models are trained. These models are optimized for



helpfulness and cooperative interpretation of user intent, which makes them vulnerable to intent disguise. Additional techniques included account rotation when flagged, cross-referencing six platforms to find the least restricted response, and direct training from ISIS operatives on specific jailbreak approaches. Academic research on this attack surface shows success rates of 88 to 97 percent against current safety defenses on frontier models.


Why can't individual AI companies fix this on their own?

Each company's safety team monitors its own platform. No mechanism exists for cross-platform intelligence sharing, so when the same adversary group queries ChatGPT, Claude, Gemini, Grok, Meta AI, and DeepSeek in sequence to find the most permissive response, no single company's system sees the full pattern. An improvement at one company provides only partial protection when five others are available. Additionally, open-source AI models with no meaningful safety constraints — there were over 368,000 text-generation models on Hugging Face as of May 2026 — provide a permanently available alternative that no commercial platform's policies can address. The Cambridge study argues that mandatory third-party safety evaluation and structured

government-company intelligence sharing are necessary because self-regulation has not survived contact with organized, persistent, trained adversaries.

What makes AI jailbreak knowledge harder to control than prior weapons knowledge shared between terrorist groups?

When terrorist networks transferred prior weapons knowledge — bomb-making techniques, drone operation, firearms maintenance — those transfers involved physical materials, equipment, or the physical presence of trainers who crossed borders with detectable training materials. AI jailbreak knowledge transfers as text, at zero cost, in ways that pass through any existing border control or interdiction mechanism without triggering any surveillance. Export controls, the Wassenaar Arrangement, sanctions regimes, and physical interdiction tools were all designed around the assumption that dangerous knowledge requires dangerous materials. AI jailbreaking requires none. No country's counter-proliferation framework currently addresses this gap, which the Cambridge and CSIS studies together identify as the most urgent structural challenge in AI terrorism governance.



ICI
International
CBRNE
INSTITUTE



CYBER NEWS



Artificial Intelligence and the Risk of Strategic Cyberattack


By Michael Sulmeyer

Source: [Download PDF](#)

July 14 – Cyber weapons have long been described as strategic instruments of statecraft. In practice, they have almost always been tactical in their application. Attackers have struggled to predict the effects of their operations with confidence, overcome complex classification structures, and assemble the expert-level capacity required to execute sophisticated cyberattacks at scale. As a result, cyber operations have rarely achieved the kind of broad, sustained strategic effects that theorists, journalists, and U.S. government sources have long anticipated.

In this paper, the author argues that artificial intelligence (AI)—particularly agentic AI, meaning systems capable of pursuing goals with some degree of autonomy—will significantly expand the strategic potential of cyber operations. It will do so primarily by alleviating the constraint on expert-level human capacity that has most limited strategic outcomes to date.

The author first examines the tactical-strategic divide as it applies to tools of statecraft and force. He then reviews why publicly known cyber operations have largely fallen short of strategic expectations, identifying three key characteristics that have confined most cyber activity to tactical uses. Finally, he assesses how current developments in AI are likely to reshape each of these three factors, though not equally, and what that means for the future strategic potential of cyber operations.



THE CYBER-THREAT AGAINST CHEMICAL, BIOLOGICAL, RADIOLOGICAL AND NUCLEAR (CBRN) FACILITIES

by Adil Radoini and Muznah Siddiqui

Adil Radoini is the United Nations Interregional Crime and Justice Research Institute (UNICRI) regional coordinator for the Middle East and Gulf Cooperation Countries. He works for the Chemical, Biological, Radiological and Nuclear (CBRN) and Security Governance Programme. He previously worked as a journalist for the Italian press and television sectors. In 2009, together with other international experts, he published “un hussein alla casa bianca”, a perspective of the Arab world on the 2008 U.S. elections. He graduated from the University of Bologna with a master's degree in International Relations focusing on Middle Eastern politics, carrying out a research thesis led in Cairo and at the Institut d'Études Politiques in Paris.

Muznah Siddiqui is a graduate from the University of Cambridge, and has completed her master's in International Relations and Politics. She is currently working as an intern at the United Nations Office of Counter-Terrorism, and her research interests include the protection of human rights, cyber-security and countering violent extremism.





C²BRNE
D I A R Y

& Robotic



DRONE NEWS



Drone Operators Can Now Carry Their Entire Control Station on Their Vest

Source: <https://i-hls.com/archives/136961>



June 22 – First-person-view (FPV) drones have become increasingly important for reconnaissance and strike missions, but controlling them often requires bulky ground equipment. Traditional ground control stations can involve separate cases, large antenna assemblies, or vehicle-mounted setups that limit operator mobility. Once deployed, drone pilots may have difficulty relocating quickly, reducing flexibility in dynamic environments.

A newly introduced wearable ground control system (Crossbow Block 2 by Neros Technologies) is designed to address that problem by reducing the size of the control station by more than half while maintaining long-range drone connectivity. The compact design allows operators to carry the entire system directly on their body armor rather than relying on dedicated equipment cases or fixed positions.

The wearable controller is designed for FPV drone operations at ranges of up to 25 kilometers. Instead of separating radios, antennas, goggles, and control hardware into multiple

components, the redesign consolidates much of the equipment into a body-worn configuration intended to improve freedom of movement.

One of the major engineering changes involves the antenna system. According to Interesting Engineering, FPV drone operators often need to maintain careful antenna alignment while simultaneously flying the aircraft and monitoring targets. To reduce that workload, the new system uses antennas with more than double the beam width of the previous design. A wider beam creates a larger coverage area, reducing the need for precise aiming while helping maintain signal quality during flight.

The platform also incorporates dual receivers, providing redundancy if one signal path degrades. Operators can still use a displaced radio configuration when missions require antennas to be positioned away from the pilot to improve signal geometry or reduce exposure.




Additional design changes focus on usability under operational conditions. Simplified push-pull connectors link the radio, goggles, and controller, making assembly faster and reducing connection errors during night operations or high-stress scenarios. The radio components and pilot control box are also housed in weather-resistant enclosures designed for field use. From a defense perspective, mobility is becoming increasingly important for FPV drone teams. Operators often need to relocate quickly, maintain concealment, and avoid

detection while continuing to control aircraft. Reducing the size and complexity of control equipment can directly affect survivability and operational flexibility.

The system also supports integration with broader command-and-control networks through low-latency video outputs, allowing drone imagery to be shared with battlefield mapping and situational awareness systems in real time. This helps transform FPV drones from standalone assets into connected elements of larger operational networks

Drone vs. Drone: A New Layer of Battlefield Protection

Source: <https://i-hls.com/archives/137328>



June21 – The drone threat continues to evolve at a rapid pace. Explosive drones and loitering munitions have become some of the most significant threats facing ground forces, particularly during maneuver operations. While conventional air defense systems are capable of engaging larger threats, maneuvering forces often lack an organic, immediate, and close-range solution that can operate on the move without prior deployment or preparation. Against this backdrop, Rafael and SpearUAV have unveiled Iron Wasp – a new interception system designed to counter explosive drones and other short-range aerial threats. The system is based on SpearUAV's Viper technology and combines detection, identification, and interception capabilities within a single solution mounted directly on an armored combat vehicle. The concept differs

from traditional air defense architectures. Rather than relying on dedicated batteries or pre-positioned systems, the interceptor is launched from a compact multi-canister launcher installed on the vehicle itself. This provides maneuvering forces with an organic interception capability that is available at any moment, including while on the move. Once a threat is detected, the interceptor drone can be launched immediately, locate the target, and neutralize it without creating a significant operational burden for the vehicle's crew. One of the system's key advantages is its ability to engage small, low-signature targets, which are the very threats that increasingly challenge ground forces. Its compact design and relatively low resource requirements allow



integration across a variety of platforms without compromising vehicle performance or mobility. From a technological perspective, the system is not intended to replace existing defenses but rather to complement them. It was designed as part of a broader multi-layered protection architecture and can operate alongside systems such as Trophy for armored vehicle protection, Drone Dome for counter-drone missions, and laser-based air defense systems.

The emergence of low-cost drones and autonomous swarms is reshaping the threat landscape for ground forces. As a result, defense solutions are increasingly shifting from area-defense concepts toward capabilities that accompany the force itself. The system reflects this trend by bringing interception capabilities directly to the tactical edge, enabling maneuvering units to defend themselves against aerial threats in real time while maintaining operational momentum.

Drones or Flies?



ROK military stages first live-fire drill against drone swarms

Designing Drones for Africa

By Maxwell Maduka

Source: <https://warontherocks.com/cogs-of-war/designing-drones-for-africa/>

June 24 – This exclusive Cogs of War interview is with Maxwell Maduka, the co-founder and chief engineer of Terra Industries, an African defense technology company building autonomous drone and counter-drone systems designed for the continent's operating conditions. As cheap imported airframes flood African markets and non-state actors employ drones across the Sahel, we asked Max why Terra is betting on Africa building its own defense industrial base.

Cheap Turkish and Chinese drones and sensor systems have proliferated across Africa for years. Where do those systems fall short, and what capabilities do you hope to add to the crowded unmanned market?

Imported Turkish and Chinese drones are cheap and fly fine, but they are difficult to sustain locally, and do not guarantee [sovereignty over software and data](#). These drones are also mainly designed for other terrains, then sold as a hard-to-

maintain generic kit. Few are fit for African operating conditions like dust, heat, and patchy navigation systems. A single failed part can literally ground a drone for months if there is no in-country support or spares. There's also the worry of espionage: If your intelligence, surveillance, and reconnaissance backbone is still linked to these countries, and you can't audit the



firmware, your operational picture is probably insecure. We fill that gap through indigenous build and support from the continent. Our own software stack (ArtemisOS) keeps autonomy and data under the operator's control. Its specifications and communications are tuned for Africa, and we have a strict human-in-the-loop policy.

The environmental challenges of designing a drone for a Ukrainian winter differ greatly from hot, humid, or dusty climates such as those found in Africa. What conditions and differences are you designing for? What has proven most challenging?

It's almost the inverse of what you're solving for in Ukraine. The big four for us are heat, dust, humidity, and topography, which [affect different parts of the system](#).

Heat is the trickiest problem. It's just always there, degrading batteries, cutting the drone's endurance, and cooking the compute needed for the onboard autonomy. Keeping processing cool enough to stay reliable when the air around it is already hot, and doing it without bolting on weight that costs endurance elsewhere, is very difficult.

Anyone who has been to Africa knows there is a lot of dust. Harmattan (the Northwesterly trade wind blowing from the Sahara over West Africa from November to March) in particular seems to get into everything, including bearings, motors, and connectors. So, we end up caring a lot more about sealing and about how easily you can service a unit in the field than the Ukrainians may.

The winds here are also stronger and gustier than most imported platforms are equipped for, which eats into endurance. You're doing more work on flight control than in other places just to hold a steady shot or keep the aircraft on station.

And then there's topography. Africa is blessed with varying geography, from the Ethiopian highlands and rift valleys to dense rainforest and coastal plains, and all of it affects communication differently, from knife-edge diffraction to signal attenuation. We account for that by using dual-band signal hopping and testing various bands.

You advertise most of your systems as primarily serving defensive capabilities, from surveillance and perimeter monitoring to search and rescue, environmental analysis, and agricultural uses. Why did you design around those use cases? Do you foresee your drones being used for kinetic purposes?

We designed around awareness and protection use cases because that's where the persistent demand is, which offers a bigger and more sustainable market than kinetic effects. The day-to-day problems across the continent are very different than the drone use cases you see in conventional warfare in Ukraine. Here, it's about a pipeline being tapped, a mine site to monitor, farmland to assess, people to find after a flood.

That said, intelligence, surveillance, and reconnaissance platforms and strike platforms require similar hardware and software, including airframes, autonomy, datalinks, and

optics, as kinetic platforms. Any serious manufacturer knows the technology isn't hermetically sealed, so we assume our drones will be used for defensive kinetic purposes.

For instance, improvised explosive devices are [among the most destructive tools affecting both civilians and military personnel in Africa](#). Our Minesweeper, a mine-detecting unmanned ground vehicle, protects against improvised explosive devices and unexploded explosive ordnance. And non-state actors have been using drones across the Sahel in recent years (since 2023, over [100 drone strikes have been carried out](#) by terrorist organizations across the region). So, the counter-unmanned aerial system mission is also important.

But we have a firm principle that every armed effect in our systems needs positive authorization from a human operator. We don't and won't build autonomous lethality: Nothing in our engineering or documentation stack allows a system to decide to engage on its own.

What's been the most interesting or unexpected use of one of your systems that you hadn't planned for?

There are two that stand out, but we're sure more will pop up. The first is communication interception, which opened up a whole internal conversation — operators using our platforms to monitor signals in ways we hadn't explicitly designed for. Lawful interception of communication data is paramount to the success of any defense system.

The second is that some of our operators have been using the aircraft itself as a communication node. In environments where ground infrastructure is patchy or nonexistent, the drone has become more of a relay — part of the network rather than just a sensor on top of it.

We're trying to minimize future surprises by having as many direct, hands-on understandings of the problem as possible, and working with operators. We do mental modelling together with those facing these problems, working from the ground up to iterate our systems quickly when we get feedback. But modularity of our systems is key.

Terra recently announced it would be opening offices in London and San Francisco. What are some unforeseen challenges in entering the U.S. and European markets?

The United States and Europe have established defense primes and a new generation of defense tech companies that are already competing hard for the same government contract dollars. Inserting ourselves into that market would not be a good strategy.

Yet in our view, the most persistent security challenges are concentrated in places where Western defense systems don't work well, or aren't designed for. As we've done with Africa, we plan to serve governments in South Asia and South America, [which also face transnational terrorism](#), insurgency, border insecurity, illegal extraction, and threats to critical infrastructure. We'd like to become the defense prime for the Global South. Our offices in San



Francisco and London will allow us to access engineering talent and capital networks needed to scale the company's technology and accelerate deployments across these markets. Countries in the Global South are building and [investing in sovereign military and defense capabilities](#), some built around affordability, autonomy, persistence, and local operational realities. Entering these new markets means navigating procurement frameworks, regulatory environments, and government relationships that may be even slower to sell into than mature, Western procurement systems. We've found relationship-building to be particularly critical in these markets. And while terrains and topographies across the Global South differ, designing and building for Africa's conditions already gives us a useful foundation.

For both outside investors and your company, which is building a large manufacturing facility in Ghana, what are the advantages of manufacturing in Africa? What are some common misconceptions?

There are lots of benefits. When your factory is built in the same environment and the same market you're selling into, your feedback loop is tighter: Field issues come back fast, and iterations are quicker because the factory is in similar environmental conditions. For an African customer, there's also more trust. A locally made system, with the supply chain and data staying in the region, is easier to buy and to rely on. And then there's the talent. We have young and capable people who are still underpriced relative to the rest of the world and have lots of untapped potential.

Troublingly, the risk gets over-weighted, and everything that's working gets under-weighted. There are usually two big misconceptions. First is that you can't build advanced hardware here — Africa is fine for light assembly but not for autonomy, embedded systems, or precision work. We like to think we are proving that thinking is outdated. The second is

that Terra's business case relies on cheap labor, but that also isn't true anymore because good engineers cost money everywhere.

People also fairly point out that Africa lacks dependable infrastructure. Yes, securing logistics like power isn't very easy, but it's possible to design around that.

Your company calls itself Africa's first defense tech prime. What resources does the current African defense tech ecosystem lack? What investments should African governments be making in the next decade to strengthen it?

The ecosystem needs four things to become more self-sustaining: patient, mission-aware capital; a local supplier base for subsystems; test and certification infrastructure; and systems-engineering and program management talent. Of those four, the supplier base is the biggest deal right now. Many critical subsystem parts are still import-dependent.

That's why [our partnership with the Defence Industries Corporation of Nigeria](#), Nigeria's state-run defense manufacturer, is so important. The agreement establishes a joint venture company focused on local production, assembly, and research and development of drones, robotics, and cybersecurity systems, with technology transfer and training built into the structure, in an attempt to grow the supplier base from the ground up. We plan to help replicate this across other African markets in the future as we scale.

As such, what African governments should be investing in over the next decade is more infrastructure: shared test and certification facilities, favorable procurement policies, and engineering education pipelines tied to actual defense programs. The goal is to build an ecosystem where the next generation of African defense and hardware companies has people to hire, parts, and government customers that can actually buy what the companies are selling.

Maxwell Maduka is the co-founder and chief engineer of Terra Industries.



pilot was concussed during the crash, but if his recollection is correct it could indicate that Iran has developed a meshed networking capability. Meshed networking allows multiple drones to be controlled at once by one

Downed US Pilot Reported Seeing Iranian Drones Swarm in 'Jellyfish' Formation

By Zachary Cohen and Katie Bo Lillis | CNN

Source: <https://www.iranwatch.org/news-brief/downed-us-pilot-reported-seeing-iranian-drones-swarm-jellyfish-formation>

June 23 – A U.S. pilot recalled seeing multiple Iranian drones hover together in formation before he was shot down in Iranian airspace in April, according to four sources familiar with his intelligence debriefing. The pilot reported a formation resembling a jellyfish, with large drones above and smaller drones underneath. The



operator. Russia and China are believed to have the capability, and two CNN sources said that there have been reports indicating that Russia and China have assisted Iran in developing its drone technology.

This Combat Robot Just Fired Anti-Tank Missiles During Live Trials

Source: <https://i-hls.com/archives/136587>

June 29 – Integrating heavy weapons onto compact robotic ground vehicles has become a growing focus for modern militaries, but the engineering challenge is significant. Small unmanned platforms often struggle to absorb the physical stress created by missile launches, where recoil and vibration can damage suspension systems, electronics, or vehicle frames. Making lightweight robotic



vehicles combat-capable requires more than simply attaching weapons; it demands a platform that can survive repeated firing under battlefield conditions. Recent live-fire trials of the Ziesel robotic platform (developed by Diehl Defence) and the SPIKE LR guided missile (produced by RAFAEL Advanced Defense Systems in partnership with EuroSpike GmbH) demonstrated progress toward that goal with a compact tracked combat robot equipped with long-range anti-tank missiles. During a multi-day test campaign, engineers launched 17 guided missiles from the unmanned platform while monitoring how the vehicle handled the mechanical strain generated by repeated firings. The testing focused less on missile accuracy and more on platform survivability. According to Interesting Engineering, the robotic vehicle maintained structural stability, mobility, and system functionality throughout the trials without major mechanical failures. That result is particularly important for smaller unmanned systems, where weight distribution and shock absorption become critical when integrating heavier weapons. The platform was originally developed for non-combat support missions such as transport and casualty evacuation before later being adapted for offensive roles. The latest version incorporates upgraded control systems, revised hardware, and an autonomy package designed to support remote operation and autonomous navigation in battlefield environments. The missile integrated onto the vehicle provides anti-armor capability at ranges of several kms and uses electro-optical guidance for precision targeting. Operators can use fire-and-forget engagement modes or continue updating targeting information during flight, allowing flexibility against moving or partially

concealed targets. From a defense perspective, robotic combat vehicles are increasingly viewed as a way to move firepower closer to contested areas without exposing crews directly to enemy fire. Instead of deploying infantry anti-tank teams in vulnerable positions, militaries could send smaller unmanned systems forward to engage armored threats remotely. The broader trend reflects lessons drawn from recent conflicts, where drones and remotely operated systems have become deeply integrated into frontline operations. As autonomy, mobility, and weapons integration continue improving, smaller robotic combat platforms may take on a larger role in reconnaissance, force protection, and direct engagement missions across future battlefields.

What if ...

The Russian embassy in Sweden announced that it was attacked by two drones in the early hours of Thursday (2/7), which caused damage by dropping paint on the building. According to the announcement, at around 02:00 in the



morning, one drone dropped a bucket of red paint on the embassy premises, while a little later a second drone, which “carried a fake improvised explosive device, crashed – most likely intentionally – on the property.

► What if the “IED” was a real one?

China's robotic wolf urban combat footage revealed for the first time!



UN addresses AI and the Dangers of Lethal Autonomous Weapons Systems

Source: <https://unric.org/en/un-addresses-ai-and-the-dangers-of-lethal-autonomous-weapons-systems/>



As artificial intelligence (AI) technology continues to advance, opportunities for its large-scale application have expanded significantly. However, alongside positive advancements come serious concerns, particularly regarding its use in modern warfare and the development of Lethal Autonomous Weapons Systems (LAWS). Commonly called “killer robots,” these systems leverage AI to identify, select, and eliminate human targets without requiring direct human intervention, raising profound ethical, legal, and security questions.

In his latest [report on current developments in science and technology and their potential impact on international security and disarmament efforts](#), United Nations Secretary-General Antonio Guterres highlighted the risks associated with the inability of AI to uphold the principle of distinction, which mandates the protection of civilians by distinguishing them from combatants, in compliance with international humanitarian law (IHL).

Since AI picks up inadvertent biases from the underlying data sets that are used to train it, “in the context of autonomous weapons, the criteria that will inform who is and who is not a combatant or target will likely involve factors including gender, age, race, and ability”, according to the preliminary [report of a UN expert group](#).

The potential of artificial intelligence to revolutionise robotic drone swarms presents another particularly dangerous development. Current armed conflicts have shown the increasing utilisation of drones on the battlefield to not just carry out surveillance and intelligence tasks, but also undertake offensive missions. With the help of AI, these low-cost robots could form highly autonomous swarms capable of

striking multiple targets simultaneously on a large scale, possibly challenging the principles of proportionality and precaution under IHL.

A primary concern is the challenge of maintaining meaningful human control over these systems, as direct control over swarms is both impossible and counterproductive since the individual units need to rely on autonomous decision-making to form a swarm effectively.

Since machines cannot be held responsible for breaches of international law, any decision by LAWS must ultimately be traceable to a human.

In his [New Agenda for Peace](#), the Secretary-General laid out recommendations on the governance of AI and autonomous weapons systems, including calling for a legally binding treaty to prohibit LAWS that function without human control or oversight, to be concluded by 2026. At September’s Summit of the Future, Member States [reiterated the need for further cooperation and agreed to continue talks](#) “with the aim to develop an instrument, without prejudging its nature, and other possible measures to address emerging technologies in the area of lethal autonomous weapons systems.”

This process is to be facilitated by the [Group of Governmental Experts on Emerging Technologies in the Area of Lethal Autonomous Weapons System](#) (GGE LAWS).

With AI’s capabilities progressing rapidly and unpredictably, these efforts present a path towards global governance of this emerging technology, making use of AI’s tremendous potential for good in accelerating progress on the Sustainable Development Goals; while keeping in check the risks associated with its usage in weapons systems.

Lethal Autonomous Weapons: The Next Frontier in International Security and Arms Control

By **Euysun Hwang** | Stanford Master’s in International Policy, ’25

Source: <https://fsi.stanford.edu/sipr/content/lethal-autonomous-weapons-next-frontier-international-security-and-arms-control>

The rapid advancement of artificial intelligence (AI) has revolutionized military technology, leading to the development of Lethal Autonomous Weapon Systems (LAWS), or “killer robots,” which can independently identify and engage targets without human intervention. This paper will explore how discussions about LAWS parallel nuclear arms control negotiations during the Cold War, highlighting an urgent need for a global arms control measure for LAWS. By analyzing the current state of LAWS development and deployment, as well as ongoing international efforts toward regulation, this article aims to provide a foundation for understanding the urgent need for robust safeguards to ensure the responsible use of AI in military operations.

Background

The term “Lethal Autonomous Weapon Systems (LAWS)” refers to weapons that can identify and kill targets (people)

autonomously using AI without human intervention, often called “killer robots”. According to the United States (U.S.) Department of Defense, once activated, [LAWS are capable of selecting and engaging military targets autonomously](#) without further input from human operators. The development and deployment of AI-based LAWS represent a significant shift in military technology, with profound implications for the future of warfare. These systems, which operate with a high degree of autonomy, can identify, select, and engage targets without human intervention, marking a major evolution from traditional, human-controlled weapons. This integration of AI into military systems is not just about enhancing the effectiveness of combat operations; it is also about transforming the broader military strategy OODA (Observe - Orient - Decide - Act), where [it can be used in all military decision-making processes](#) including



intelligence, surveillance, and reconnaissance (ISR). This signifies a shift from traditional computing support (such as ballistic calculations) to AI's involvement and autonomous decision-making and operation throughout all phases of military action. LAWS could potentially operate across various platforms, including unmanned aerial vehicles, naval vessels, and ground robots, providing a multidimensional capability that extends from conventional battlefields to cyber and space domains.



(Left) DARPA ACE project AI-piloted F-16 (Right) Russia Marker unmanned ground robot. Source: Wikimedia Common

As the deployment of LAWS in real-world scenarios becomes more prevalent with the first combat deployment of LAWS in the Ukraine War, international interest in arms control of LAWS has begun to grow. The Ukraine War, described as the [Silicon Valley of offensive AI](#), has become a testing ground for various future warfare technologies and has been dubbed a “drone war” due to the various drones being introduced and actively used, significantly innovating the traditional combat methods. The [international community](#) have begun to fear the rise of the possible unchecked use and proliferation of these weapons that could ultimately disrupt the international order, and as a result, begun debating the legal, institutional, ethical norms of this warfare. Some experts advocate for the establishment of a ‘Next Oppenheimer’ paradigm, drawing on historical lessons from nuclear weapons development.

Ongoing Developments and Global Discussion on LAWS

Major technologically advanced nations like the U.S., China, and Russia have already accelerated the development of military AI technologies in preparation for future warfare and incorporated it into their defense strategies. For example, in 2018, the U.S. published the [“Department of Defense AI Strategy”](#) which emphasized the role of AI in the military domain. Similarly, [China](#) and [Russia](#) have declared the adoption of AI technologies across all sectors of defense and national security.



- **United States:** The Defense Advanced Research Projects Agency (DARPA) has initiated the [Air Combat Evolution \(ACE\)](#) research project. Following the completion and operational deployment of stealth fighters like the F-22 and F-35, the project focuses on converting existing F-16 fighters, now surplus to peacetime needs, into AI-powered fighters. [Defense agencies are also strengthening its ties with tech leaders](#) to utilize AI tools to rapidly process complex data.
- **China:** China has completed the development of the [“Liaowangzhe II,” a fast unmanned patrol boat equipped with AI-driven automatic navigation](#) and optimal route-finding capabilities, making it the second in the world to do so. The country is also developing “swarm technology” for unmanned missile craft, a tactic known as the “shark swarm,” intended to deter U.S. aircraft carrier groups from intervening in potential conflicts in the Taiwan Strait. Additionally, in cooperation with Russia, [China is developing AI-powered autonomous weaponry, showcasing a gun-mounted robot dog manufactured by the Chinese company Unitree Robotics.](#)
- **Russia:** [Russia has developed “Marker,”](#) an unmanned ground robot in the form of a tank equipped with autonomous driving capabilities and an AI system that analyzes images of enemy vehicles. This system identifies Western tanks and ground forces, allows AI to determine attack priorities, and can even decide when to engage in attack actions.

Despite the escalation of tension with an AI arms race among major powers and other key nations worldwide, discussions on the need for arms control have been insufficient. The main efforts have been towards creating guidelines or norms, and there is a low likelihood of any new treaties or agreements on AI arms control being established in the short term. The debate on the legality of using LAWS first began at the 2013 Convention on Certain Conventional Weapons (CCW). At the 2016 CCW meeting, an open-ended GGE on LAWS was established. Discussions at the GGE





have primarily focused on the ethics and accountability of LAWS, exploring their compatibility with existing international law, particularly International Humanitarian Law (IHL). [GGE's document on LAWS](#) on July 26, 2024, based on IHL, urge states to take the following actions:

- Ensure appropriate training and instructions for human operators of LAWS
- Ensure that LAWS operate with appropriate control and human judgement across the entire life cycle of the weapon systems
- Limit the types of targets that the system can engage
- Limit the number of engagements that LAWS can undertake

However, as stated in the document, nothing has reached a consensus, and while the framework provides a platform to discuss the issue, it is not rigorous enough to keep up with the rapidly changing trends. In fact, breakthroughs like GPT in text processing and RT-2 in image processing have made it much easier to rapidly handle large amounts of complex data. This is blurring the line between active human decision-making and passively following data-driven results that appear correct but merely keep humans in the loop.

Experts warn against the spread of military AI and call for initiating discussions on AI arms reduction. In Henry Kissinger and Graham Allison's articles titled "[The Path to AI Arms Control](#)," they highlight the "Oppenheimer Moment" of the AI era, expressing concerns about the potential misuse of technology. For instance, military AI could make it easier to target non-combatant groups using advanced image recognition technology, or could lead to potential

misinterpretations of data that might result in the bombing of civilian facilities. Kissinger and Allison particularly stress the need to control the use of AI in military applications. However, history shows how challenging it is to achieve arms control.

We might see history repeating itself in the governance of LAWS. The recognition of LAWS as a threat comparable to nuclear weapons marks the beginning of their consideration for arms control, but the road to negotiation is tough. The U.S. and China perceive LAWS as strategic assets, crucial for achieving strategic superiority. During the 2023 San Francisco Bilateral Meeting, the two countries initially agreed to hold their first AI arms control talks in 2024, but [China later announced the suspension of talks](#). [The 2024 Vienna Arms Control Conference](#) marked the first discussions on the need to regulate LAWS, focusing on human control and accountability, as well as the ethics of algorithms. Following the Vienna conference, global media outlets like Bloomberg highlighted articles with the keywords, "AI Arms Reduction" and "[Oppenheimer Moment](#)," emphasizing the perception of military AI as a threat to humanity on par with the emergence of nuclear weapons and underscoring the need for global arms control.

Future of Arms Control for Lethal Autonomous Weapons System

Since military AI is still in the early stages of development, most international discussions have primarily focused on maintaining human control adhering to the taboo concept of humans being killed by autonomous systems. Considering the rapid growth of the LAWS market



following geopolitical competition between countries to have an advantage, [AI companies' growing involvement in the defense sector](#), and the real-world application, seen by the lethal technological advancements with the Ukraine War, it is foreseeable that the next "Oppenheimer Moment" could arrive sooner than anticipated. In fact, the possibility of remote war could also be amplified through increased autonomy with humans becoming further alienated from the use of force, [making a proxy war utilizing the LAWS amore cost-efficient and tempting choice](#). This may lead to lowering the threshold to war, making military action more politically acceptable domestically and making conflict easier to enter. An empirical example to draw from is again, the nuclear arms control deals. After World War II, the nuclear-armed nations (P5) secured permanent seats on the UN Security Council and established the Nuclear Proliferation Treaty (NPT) system that prohibited other countries from acquiring nuclear weapons, thereby maintaining their exclusive status. As a result, nations that possessed nuclear weapons (the "haves") became participants in arms control negotiations, while those without nuclear capabilities at the time (the "have-nots") mostly [remained subjects of the non-proliferation regime](#). Given historical precedents and recognizing the strategic disadvantages of not possessing advanced military technologies, nations will likely become increasingly driven to intensify their research and development of LAWS, further accelerating the arms race and potential misuse of these technologies, especially by exploiting the regulatory gray areas where no clear international standards exist. This motivation is not solely driven by national security and sovereignty concerns but also by a desire to influence the frameworks and operational norms, such as in deciding the scope and establishing a data structure for the use of LAWS. LAWS regulation may unfold at the backdrop of strategic alliances around AI. As technological advancements reshape global dynamics, nations will likely prioritize staying ahead in an era where AI could redefine state power—driving economic growth, enhancing military efficiency, and reshaping global influence. In addition, the development of LAWS involves

complex challenges that require the formulation of robust algorithms, the establishment of rigorous military standards, and the comprehensive collection of operational data. This may lead to the formation of strategic alliances focused on technology and AI, which could serve to pool resources, share critical technological know-how, and negotiate as a bloc in international forums. These alliances might also work towards establishing common standards for AI implementation, including technical standards for interoperability, ethical guidelines, and regulatory frameworks to manage AI's societal impact.

The formation of such blocs will reinforce the broader trend towards multipolarity, where regional leaders may emerge as rallying points for like-minded countries, uniting to align their standards and collectively advance technological development. At the same time, this shift is likely to exacerbate existing complexities in establishing governance on LAWS, as the power struggle between blocs to make their technology a standard or "key technology," will further impede the development of a coherent governance structure and intensify the arms race for physical dominance.

Conclusion

In summary, the rapid development and deployment of Lethal Autonomous Weapon Systems (LAWS) represent a significant shift in military capabilities and strategies, prompting global discussions on arms control similar to historical nuclear weapons negotiations. As nations like the U.S., China, and Russia integrate AI into their military frameworks, the need for comprehensive international legal and ethical guidelines becomes urgent. The use of LAWS in conflicts, particularly noted in the Ukraine War, highlights the potential for these systems to alter traditional warfare and international security dynamics. Effective management and regulation of LAWS are critical to prevent a new arms race and ensure global stability, making the development of international standards and agreements on the use of military AI essential for future security architectures.

Artificial Intelligence, Emerging Technology, and Lethal Autonomous Weapons Systems: Security, moral, and ethical perspectives in Asia

By Mitzi Austero, Alfredo Ferrariz Lubang, Binalakshmi Nepam, Pauleen Gorospe Savage, and Kazuyo Tanaka

Source: <https://www.stopkillerrobots.org/wp-content/uploads/2021/11/NISEA-AI-Emerging-Tech-and-LAWS-Perspectives-in-Asia.pdf>

Executive Summary

Artificial Intelligence (AI) development has been steadily expanding in the last decade, especially in the areas of economic development, rapid industrialization, increased productivity, and now, weaponry. Lethal Autonomous Weapons Systems (LAWS) are gaining attention due to

prominent advances in weapons development. LAWS are loosely defined as "weapons that can select and engage targets without human intervention." Global concern over the use of LAWS on human beings is growing, especially in countries that suffer from various security issues. Internal insecurity and armed confrontations over territorial disputes have all increased circumspection about



the weaponization of AI and its integration into LAWS, contributing to its moniker, “killer robots”. However, the threat does not merely lie in lethality, as an autonomous weapon system does not need to be “lethal” to inflict damage, and the weaponization of AI and the range of autonomous weapon systems that can inflict harm still pose a significant threat to human security. Beyond physical harm, the threat of force is enough to control the populace by discouraging certain actions. For now, the element of lethality remains unclear as most of these weapon systems are used to intercept and eliminate incoming projectiles. In Asia, China, South Korea and India are known to be developing capabilities to weaponize AI. The rest of Asia is still a weapon-importing region, though military spending has steadily increased in the past decade. The region’s wide variety of cultures and political systems raise questions on how the weaponization of AI will affect its stability. China and South Korea have seen the most rapid economic growth in the past decades, including significant innovations in the research and development of military technology. Southeast Asia has also seen impressive economic growth, which could enable governments to acquire complex military weapons, though perhaps not as advanced as LAWS. South Asia’s economic growth as a whole has been less successful and is thus relatively underdeveloped compared to the other two sub-regions, though India is the clear economic power in the region. Each sub-region has experienced armed conflicts as well as unique political and socioeconomic challenges.

Much of Asia has not yet finalized its views and national positions on LAWS. A few countries have expressed concern over its manner of use and the applicability of International Humanitarian Law (IHL). So far, only three countries are in favor of a ban on use and only Pakistan is in favor of a ban on development. Based on consultations with government representatives, few understand and are aware of the technological requirements of LAWS development and the potential dangers of its use, especially as LAWS and its components are exceedingly more complex than non-autonomous weapons. A substantive dialogue with Asian countries would require sufficient coverage of these elements so that governments can respond adequately, participate actively in international discussions and develop their own policies. This paper also shows that Asia will likely be divided between producers-suppliers and recipients-buyers of LAWS, also defined by the country’s economic status. Lower-middle-income countries and middle-income countries may be attempting to develop precursors to LAWS but will ultimately be behind countries such as China, South Korea, Singapore and Japan, which may likely continue to devote resources to defense spending. Though the majority of countries in Asia are less likely to manufacture LAWS due to lack of expertise and capability, these countries can still be suppliers of parts, components, or software, making regulatory policies a necessary standard for the entire region. There are currently

no international agreements or regulation frameworks that address LAWS specifically. Discussions regarding LAWS started in informal meetings leading up to the CCW in 2014 and have been continuing through Group of Governmental Experts (GGE) meetings ever since. Some agreements could serve as a foundation for future agreements by virtue of their scope. Treaties such as the 2017 Treaty on the Prohibition of Nuclear Weapons (TPNW), the 2013 Arms Trade Treaty (ATT), the 2008 Convention on Cluster Munitions (CCM), and the 1996 Mine Ban Treaty (MBT) may cover related weapons or parts of LAWS. The Convention on Certain Conventional Weapons (CCW) and its Protocol IV, adopted in 1995, on Blinding Laser Weapons is perhaps the most relevant provision as it preemptively banned a weapon that is still being developed. The Martens Clause found in the Geneva Convention is also relevant in the moral and ethical discussion of a weapon that is still perceived to be under development, as it states that “civilians and combatants remain under the protection and authority of the principles of international law derived from established custom, from the principles of humanity and from the dictates of public conscience.” The international dialogue on LAWS would benefit from a discussion of the moral and ethical implications of its potential development and use, especially regarding democracy, transparency, human rights and accountability.

Additionally, this paper recommends certain steps to spread awareness and encourage countries to confront issues that may emerge from LAWS development and use. States may be engaged at the international, regional and national levels to determine at which degree are they discussing LAWS issues and what their awareness is of the international debates. At the international level, more efforts should be made to have clarity on the definition of LAWS. It is especially critical for definitions to be decided in order to increase understanding of the development and use of LAWS and its implications for conflict, warfare and human rights. It would be useful for countries if more inter-sectoral discussions between the scientific and engineering community, government representatives and civil society are encouraged. This would provide clarity between AI and robotics workers, the state, defense, arms industries, and civil society and urge them to find a unified position. Steps should be taken to map out the “complex life cycle” of a LAWS, similar to defining the life cycle of conventional weapons, which includes various aspects of conceptualization, development, up to its disposal. As standards are important in contributing to a wider understanding of LAWS, a legally-binding international instrument must take into consideration the humanitarian impact of LAWS. Such an instrument should also have considerable space for the views of states who have no intention to develop, possess or use LAWS.

At the regional level, regulatory policies are important. The complexity of LAWS



components, each with its own international development and distribution process, points even more to the necessity of a regional policy response. Regular dialogue will help states develop their own positions, something civil society organizations can assist in. States in the region should be encouraged to take on LAWS as an emerging security and humanitarian issue and step up their leadership towards a common regional position. The nature of emerging technologies and the security threats it will pose in the future cannot be addressed by any single state effort, and this should be highlighted in discussions and engagements with states.

At the national level, it is important to fully engage governments to tackle the future of LAWS. National policies can only be effective if policymakers and implementing agencies understand the nature and features of LAWS. It would be useful for states to conduct further studies on the implications of LAWS in the national security, public order and safety situation vis-à-vis positive technological advancements. Any national process on LAWS regulation or ban must involve various stakeholders in preparation for a global diplomatic conference negotiating a new international

law governing LAWS. To this end, civil society organizations (CSO) can play a significant role. CSOs can serve as intermediaries between different sectors of society, including government, private and technological sectors. They can engage and encourage states to participate actively in international meetings towards developing their own national positions. Civil society's efforts must thus be supported, especially those from developing countries who do not have the resources to constantly engage governments or participate in the global discussions on LAWS. In the same vein, experts, particularly tech workers, AI and robotics experts, should be encouraged to share their views at various levels of discussions. An Asian regional platform on humanitarian disarmament can be strengthened to help build a stronger unified position of CSOs working on this issue, especially those who are working with victims in conflict-ridden countries. Knowledge materials should be developed and produced to assist CSOs in raising the awareness of the public and their respective governments. CSOs can work together towards producing a unified position and message regarding the very real threat that LAWS can pose to their communities.

Robotic bird targets drones' biggest aerodynamic shortcoming

By Omar Kardoudi

Source: <https://newatlas.com/drones/robotic-bird-drones-aerodynamic-problem-rmit/>



July 02 – A robotic bird tested in a wind tunnel may hold the blueprint for drones that can finally handle a windy day. Researchers from RMIT University (in Melbourne, Australia) and the University of Bristol (UK) have reverse-engineered the Australian kestrel (*Falco cenchroides*) to understand how it hovers effortlessly in gusty winds and what that means for the small unmanned aerial vehicles (sUAV) that still can't.

When the wind picks up, it's time to land the drone and go home. That's true whether it's [carrying a package](#), a [camera](#), or a [warhead](#). And it's not a minor inconvenience, it's a fundamental aerodynamic limitation. The findings, published across two papers in the *Journal of the Royal Society Interface* ([1](#) and [2](#)), show that a vertical gust of the same magnitude as a horizontal one generates between 25 and 100 times more lift variation in a small wing, depending on its shape. That's the most common type of disturbance at low

altitude, precisely where sUAVs operate, and climate change is set to make it significantly worse in coming decades. Nature, however, has been solving this problem for millions of years.

The Australian kestrel is a master of stationary flight in rough air. Using motion-capture technology inside RMIT's wind tunnel facilities, the team precisely recorded how a live bird moves under real turbulence conditions. The gap between bird and machine turned out to be substantial. A kestrel has more than 22 degrees of freedom to adjust its posture mid-flight. A typical drone has four. The kestrel is also dramatically lighter in the right places. Its body mass is concentrated in the torso, so it can spin and correct itself in response to a gust roughly twice as fast as a drone of similar size. A comparable micro-vehicle simply



carries too much rotational weight in its frame to react as quickly.

"Birds don't rely on a single response to wind gusts," says RMIT researcher Matt Penn, who led part of the research into how birds handle turbulence. "They constantly adjust their wings and tails to stay balanced, while the natural flexibility of their feathers and joints helps absorb sudden changes in airflow. They can also sense disruptions very quickly, which

The tail also acts as a stability dial. Spread wide, it strongly resists any gust that could tip the bird nose-up or nose-down, correcting the disturbance almost automatically. Tucked in, the bird becomes nearly neutral – less self-correcting, but far easier to maneuver sharply. The kestrel can slide between those two settings in real time, dialing up stability when holding position in a gust, and dialing it back when it needs to change direction. No current drone can do this on the fly.

The bird's advantage goes beyond clever actuators. Kestrel feathers deploy automatically under aerodynamic load to prevent airflow separation. Hair-like feathers called filoplumes – slender, nerve-tipped structures – detect vibrations and flow separation points in real time. Joint mechanoreceptors monitor structural loads continuously. No sUAV today has anything close to this distributed feedback system.

"This research shows what's possible when engineers look to nature for solutions," says Associate Professor Abdulghani Mohamed, senior researcher at RMIT. "Our findings open

allows them to respond almost instantly and maintain control." To move beyond observation and actually measure the forces involved, the team built a high-fidelity robotic replica of the bird based on CT scans of real specimens. The robot replicates the coupled wing movements – wrist and elbow extension – and tail motion, and was tested in the wind tunnel at 7 m/s (15.7 mph).

The key to the kestrel's gust control turns out to be how its wings and tail work together. When both extend simultaneously – as the real bird does during hovering – the lift increase is greater than the sum of the two movements separately, but their pitching effects cancel each other out. As a result, the bird gains lift to counter a vertical gust without changing its flight attitude at all. A conventional drone generating more lift by adjusting motor power or control surface angle inevitably tips in the process.

new pathways for designing aircraft that can better handle turbulence."

The researchers acknowledge the road from lab to sky won't be short. The kestrel's stability doesn't come from any single trick but from all these mechanisms working together simultaneously – and replicating that in a lightweight, low-cost platform is the real challenge.

The team's next step is studying how the kestrel perceives its environment, including subtle turbulence cues it anticipates before they hit, a capability that could unlock predictive control systems. While the focus for now remains on small unmanned vehicles, the researchers hope to eventually simplify their findings enough to apply them to larger aircraft too.

RMIT is already seeking industry partners to push the work forward. If successful, the next generation of small aircraft won't fight the wind. They'll ride it like a falcon.



The Drone Threat to America's Cities

By Seth Stodder

Source: <https://www.homelandsecuritynewswire.com/dr20260709-the-drone-threat-to-america-s-cities>

July 09 – This past March, [the FBI warned](#) state law enforcement that Iran might be planning to attack targets in California using drones launched from vessels floating offshore. As it turned out, the threat wasn't real. The warning was based on "unverified" intelligence, and there was no evidence of an actual plot. Californians could breathe a sigh of relief.

But we shouldn't get too comfortable. While the FBI might have been wrong about the specific Iranian plot, the threat of a serious drone attack is growing. And we're not even remotely prepared for it—or the public panic that would likely ensue.

The FBI's warning came to mind when I was in Ukraine recently as part of a





delegation with the [Renew Democracy Initiative](#) to meet with senior Ukrainian government and military officials, drone manufacturers, and defense technology companies. While there, our delegation saw firsthand how the Ukrainians have used drone warfare asymmetrically against the Russian Army and ground President Vladimir Putin's war of aggression to a bloody standstill.

Few expected this when Putin launched his invasion four years ago. At the time, most intelligence and military analysts believed that Russia—a nuclear-armed country with a powerful military, five times the population, a much bigger economy, and China's backing—would easily [overwhelm Ukraine](#) in a matter of days or weeks. As recently as 2025 and as President [Trump famously put it](#), Ukraine still supposedly didn't have the "cards" to resist Russia. But Ukraine has doggedly held on and is arguably [turning the tide](#) of the war in its favor, in large part by mass-producing millions of cheap drones and using them to wreak deadly havoc on the Russian war machine. Russia has suffered over [1.3 million casualties](#) on the battlefield, its oil refineries and other infrastructure are in flames, and Moscow itself is now subject to Ukrainian drone attacks.

Unfortunately, the Ukrainians aren't the only ones who've figured out how to produce cheap drones and use them to deadly effect against a theoretically more powerful adversary. Iran is, of course, doing this in the Persian Gulf—using its [Shahed drones](#) to restrict passage through the Strait of Hormuz. And if the FBI's alert from March 2026 has any kernel of truth to it, Iran may have contemplated a drone attack on the U.S. homeland.

Could Iran have launched Shaheds at targets in California from vessels floating offshore, as the FBI warned? Maybe, although it would have been challenging for Iranian agents to smuggle the Shaheds and launchers to North America, load them onto vessels in a port, and then get them into position offshore—all without the complex plot being foiled or detected by bystanders, law enforcement, customs, border, or Coast Guard authorities.

But make no mistake, a few dozen Shaheds launched from fishing trawlers 15 miles offshore would do serious damage. Traveling at 115 mph, Shaheds could hit targets near the coast—for instance, three of Los Angeles County's four oil refineries—in about 10 minutes. The results could be devastating: death and destruction, environmental damage, toxic plumes over urban areas, and severe fuel shortages that would hammer the Golden State's economy. The drones would fly below radar coverage, making them hard to detect—and neither California's coastline nor the refineries are defended with the kind of interceptors Ukraine uses to block 95 percent of the Shaheds that Russia launches every day. Under current law, only the military is allowed to have that kind of capability—and it's unclear whether military bases in California would actually have the ability to intercept several dozen Shahed drones just minutes away from hitting their targets.

And oil refineries aren't the only infrastructure that's vulnerable in California. There's also the power grid, Los Angeles International Airport (LAX), or the Ports of Los Angeles and Long Beach, where 40 percent of all



ocean-going containerized trade enters the U.S., along with most of the oil imported from the Middle East or brought from Alaska to supply California's economy. As Port of LA Executive Director [Gene Seroka noted](#) at the time of the FBI's warning about the supposed Iranian plot, "these ports, airports and utilities are soft targets for the bad guys." And the same could be said for coastal cities and infrastructure in other parts of the United States and around the world, such as the dense network of oil and gas infrastructure in Texas, Louisiana, and Mississippi along the Gulf Coast or major seaports like the Port of New York and New Jersey, Charleston, or Rotterdam in the Netherlands.

As bad as all that sounds, though, an Iranian Shahed attack on major cities or critical infrastructure isn't the most likely nightmare scenario. In Ukraine, most of the carnage on the battlefield is inflicted by small, inexpensive quadcopter First-Person View (FPV) drones similar to what you might buy on Amazon, but armed with grenade-sized explosives and capable of kamikaze-style attacks on soldiers or vehicles. Last year, Ukraine produced 4 million FPVs, and it expects to make another [7 million](#) this year. The drones are assembled from commercially available and 3D-printed components in hundreds of nondescript office buildings around Kyiv, Kharkiv, and other parts of the country. Ukrainian soldiers "pilot" the FPV drones remotely by looking through virtual reality glasses or at computer screens while sitting miles away from the so-called kill zones.

These cheap, mass-produced drones have transformed 21st century warfare, and [terrorist groups](#) such as ISIS and Hezbollah, as well as violent [far-right extremists](#), have taken notice. Indeed, [ISIS pioneered](#) the use of cheap armed drones during the Iraqi fight to retake Mosul in 2016-2017, and [Hezbollah](#) has used FPV drones against Israeli troops in Lebanon. Could terrorists fly FPV drones into spectators enjoying a Dodgers game or the opening ceremony for the 2028 Olympics? [Voviette Morgan](#), in charge of security for the LA28 Olympics, thinks so. Last year, she warned the House Homeland Security Committee that "the persistent threat of unmanned aircraft systems has surged in recent years to make the detection and mitigation of drones one of the overarching concerns for any sporting event."

Or consider a terrorist group replicating Russia's grisly "[Human Safari](#)" in Ukraine-held Kherson, where Russian troops across the Dnipro River use FPV drones to hunt and kill Ukrainian civilians at random, like a drone-era version of the 2002 [Washington-area sniper](#) attacks. During a tense period in October 2002, two men—John Allen Muhammad and Lee Boyd Malvo—traveled around the D.C. area in a blue 1990 Chevrolet Caprice and used a sniper rifle to shoot random people doing ordinary everyday things, including buying gas, walking in store parking lots, or simply mowing their lawns. During their three-week rampage, the snipers killed 10 and wounded three others. What if a terror group decided to do something similar with kamikaze drones instead

of sniper rifles—attacking people, say, walking down Rodeo Drive in Beverly Hills, shopping at the Grove outdoor mall near the Farmers Market in Los Angeles, enjoying Dungeness crab at Fisherman's Wharf in San Francisco, or riding the carousel on the National Mall in Washington, D.C.?

These scenarios may sound like the plot of a 1970s-era disaster movie like "Black Sunday," but the realities of drone warfare in Ukraine—and the easy and legal accessibility of inexpensive drones to anyone—make the threat too plausible to disregard. A determined group of attackers with a bunch of cheap FPV drones armed with explosives would be hard to stop, and it would paralyze a city with widespread panic, as happened to the D.C. area during those surreal and terrifying few weeks in October 2002 when the Washington snipers were on the loose.

I remember that feeling vividly. I was there. One of the random shootings—of FBI analyst Linda Franklin—took place in the parking lot of a Home Depot just minutes from my home in Arlington, Virginia. I'd been to that Home Depot many times and had parked in that same parking lot. Linda Franklin's husband—[William Franklin](#)—testified during John Allen Muhammad's trial that he was loading a shelf he'd bought into their car when he suddenly heard a loud noise and felt something hit the side of his face. It was his wife's blood. She'd been shot and killed while standing by the car's trunk, making sure their shopping cart didn't roll away. Chilling ... and totally random.

I was in the government at the time, advising the head of an agency that would soon be absorbed into the Department of Homeland Security. I was paid for reasoned analysis and strategic thinking. But during those three weeks in October 2002, I was as frightened and panicked as anyone else. Like everyone I knew, I took cover behind my car when I got gas, and I didn't linger outside my front door. And without any doubt, the same kind of primal panic would grip the public if terrorists started using kamikaze drones to attack people randomly at outdoor malls, at sporting events, in public parks, or on city streets. What could be done to respond to that kind of drone attack? Not much, at least initially. When I served in the Obama administration as assistant secretary of homeland security, I oversaw the Department of Homeland Security's counter-drone policy. And for a while, my marching orders were actually to *avoid* taking on the issue under the department's purview and instead push responsibility for it onto the Federal Aviation Administration (FAA) or another federal agency. Why? Because there were no good solutions for the drone threat. The Department of Homeland Security didn't have the right personnel in the right places; it had no authority to direct state and local forces, and the tools for actually stopping a drone attack didn't really exist. So, the department tried to avoid being held responsible for this seemingly unsolvable problem by playing a bureaucratic version of touching its nose and saying, "Not it!"



The challenge hasn't gotten any easier. For one, the U.S. doesn't track all the drones in the sky. The FAA requires most drones to be registered and to broadcast a "Digital ID" along with their location and altitude. But the rules only require broadcasting by [radio frequency](#) (RF)—not cellular internet—which inherently limits the tracking range. While it might theoretically be possible to deploy a vast network of ground RF sensors across a metropolitan area, doing so would be extraordinarily expensive and practically useless for broad defense. At most, such counter-unmanned aircraft system (C-UAS) sensors are deployed to protect ultrasensitive, discrete locations, such as major airports or specific events such as the Super Bowl, where law enforcement is concentrated and authorized to respond. But C-UAS detection systems are not deployed to protect schools, public parks, or outdoor malls where people gather daily. Think about the sheer volume of "noise" in the skies. According to the FAA, there are over [837,000 registered drones](#) in the United States, with at least [12,500](#) in Los Angeles County alone. Even conservatively assuming only 150 registered drones are airborne across LA County at any given moment, how would authorities distinguish a legitimate commercial or hobbyist drone flight from a lethal threat before the final, fatal seconds? And that's just the *registered* drones. Actual [enforcement](#) of the FAA's drone registration laws is minimal, and smaller drones don't need to be registered. So, in reality, the *actual* number of drones in the sky at any given moment is far, far greater. Moreover, terrorist FPV pilots won't be standing on street corners conspicuously wearing VR goggles. Like the Ukrainian pilots, they could be operating out of a basement miles away. It took an intense three-week manhunt involving nearly 1,000 police officers to catch the Washington snipers, who were traveling around the D.C. area in a blue sedan and sleeping at rest stops. The terrorist drone pilots could remain entirely invisible and be much harder to find. Even if an RF-controlled attack drone were detected via C-UAS sensors (a big "if") and its control signal back-tracked to a location, a sophisticated attacker would have simply followed Peter Clemenza's famous command in "The Godfather": "Leave the gun, take the cannoli." They would have dropped the controller and vanished into the urban landscape long before police arrived. We can designate "no-drone zones" over major events like the FIFA World Cup or LA28 Olympics, but how do police respond when a drone actually breaches that airspace? It often happens by accident, and separating the signal from the noise is nearly impossible. During the Paris Olympics in 2024, French security forces intercepted 90 drones and arrested 85 drone pilots—most of these were clueless ["tourists unaware of the regulations."](#) While electronic jamming or geofencing might work to bring FPV drones down or prevent them from entering a restricted area, modern drone tactics and technology are rapidly outgrowing jammable radio frequencies. In Ukraine, both

sides have developed countermeasures to thwart the electronic jamming of FPV drones. One way is by tethering the drones to their controllers via spools of ultrathin fiber-optic cable—sometimes stretching up to 20 kilometers—so the drones can communicate with human pilots without using radio frequency at all. Hezbollah uses the [fiber-optic cable technique](#) against Israeli forces in Lebanon. The [Israel Defense Forces have reportedly](#) "so far been unable to develop any effective countermeasures" to Hezbollah's fiber-optic drones. Alternatively, an attacker can use fully autonomous, artificial intelligence-driven navigation and target selection. Once launched, these autonomous drones require no external communication link whatsoever. No signal means there's nothing to jam. In short, the focus on how to electronically jam drones may be fighting the last war.

But even if the FPV drones *could* be brought down by electronic jamming, it's unlikely that local authorities would have the equipment readily available. Indeed, state and local authorities were strictly prohibited under federal law from using electronic jamming equipment until late 2025, when [Congress finally changed](#) the law. Private-sector infrastructure owners are still banned from doing so. And even under the new law (the [SAFER SKIES Act](#)), local authorities still can't use jamming technology until they're trained and certified by the Department of Justice pursuant to regulations that haven't been finalized yet. So, even if electronic jamming could bring down the attack drones, there is little capability to do so yet. In sum and as G.B. Jones, the [chief safety and security officer for FIFA World Cup 2026](#), has stated bluntly regarding the drone threat: "We don't have the training, we don't have the equipment, and we don't have the number of personnel that are skilled in mitigation technology and the use of mitigation technology, particularly the application of that technology in really highly dense radio frequency environments or urban environments."

In the meantime, and until we figure out a better answer, get ready for U.S. and European cities to start looking like parts of Kharkiv or other areas I visited recently in eastern Ukraine—with outdoor malls, concert venues, public parks, shopping streets, and other areas draped with anti-drone netting cut from fishing nets, installed to ensnare quadcopter drone blades, and protect people and vehicles from attack.

Draped fishing nets would be tragic aesthetics for beautiful cities like Los Angeles, New York, San Francisco, or London. But until we fast-track federal certifications, deploy scalable detection, and innovate past the era of electronic jamming to address the threat presented by fiber-optic or AI-piloted drones, netting may become the inevitable protective measure demanded by insurers and lawyers alike, let alone the general public and panicked politicians. To avoid that ugly reality, we need to come up with better ideas, and we need them fast. Because ready or not, the drones are coming.




Seth Stodder teaches national security and counterterrorism law at the University of Southern California Gould School of Law and is a nonresident senior fellow with the Atlantic Council.

Lawfare is a non-profit multimedia publication dedicated to "Hard National Security Choices." Lawfare provides non-partisan, timely analysis of thorny legal and policy issues through our written, audio, and other content.

Teleoperated humanoid robots complete first-ever live surgery

Source: <https://newatlas.com/robotics/first-live-surgery-teleoperated-humanoid-robots/>



July 10 – Surgeons at UC San Diego just handed the scalpel to two humanoid robots, who went on to complete live surgical procedures for the first time in history. This milestone moves beyond the fixed robotic arms found in operating rooms today and hints at an operating room of the future where humans and humanoids work side by side. Humanoid robots are already showing up in [factories](#), [warehouses](#), and even on [battlefields](#), precisely because their human-like shape lets them operate in spaces built for people without any redesign. A [Morgan Stanley report](#) from late June projects that China, the current leader in the field, will produce 446,000 humanoid units annually by 2030, with full-size humanoids growing from 30% market share in 2026 to 70% by 2028.

UC San Diego's team wants to bring that same flexibility into the operating room. In one procedure, a humanoid robot and a human surgeon acting as an assistant completed a cholecystectomy (gallbladder removal). In a second, two humanoid robots worked together and finished the operation solo. The trials, described in a paper published in [Nature](#), were carried out on large non-primate mammals, but the achievement matters because it moves the technology from

concept to something that has demonstrably handled real surgical tasks.


The researchers say their real target is the growing shortage of surgeons and the surgical backlogs it creates, especially in rural areas or regions far from major hospitals. Traditional [surgical robots](#) address none of that: they're bulky, expensive, and typically require rebuilding the operating room around them. A conventional robotic system weighs around 800 kg (1,764 lb) and needs significant space. The humanoid system used here, nicknamed Surgie, looks to have started life as a [Unitree G1](#), and stands just 1.5 m (about 5 ft) tall and weighs only 27 kg (60 lb) – compact enough to wheel into a small clinic or a field hospital.

"It's a fraction of the cost, and it takes a fraction of the space in an operating room," says Dr. Shanglei Liu, assistant professor of surgery at UC San Diego School of Medicine and one of the study's lead authors, who personally operated one of the robots during the trials. "So it's easy **to deploy, anywhere from rural areas to the battlefield and even to space.**" The surgeons control the robots remotely,



using standard surgical tools fitted with adapters so the robotic hands can grip and maneuver them properly. Testing moved through stages: lab simulations, animal trials, and finally live surgery. The results are promising but not flawless – the

clumsy. "This achievement reflects the power of bringing engineers and surgeon innovators together to solve meaningful clinical problems at our world-class training and research lab," says Dr. Ryan Broderick, interim director of UC



robots needed recalibration mid-procedure, operations took longer than usual, and latency (the lag between a surgeon's

San Diego's Center for the Future of Surgery. The team envisions humanoids eventually doing more than assisting in



hand movement and the robot's response) remains a challenge for any remote-controlled surgery. That said, today's routine surgical robots also started out slow and


surgery: fetching instruments, tidying the room, or working alongside human staff as full team members. "Many



communities struggle with adequate staffing on the surgical team, which means patients are not being treated," says Dr. Michael Yip, a UC San Diego engineering professor and co-author. "Our goal is an operating theater of the future, where humanoid robots and humans work side by side as an integrated team to deliver procedures to those in need, both

in traditional hospital settings as well as in non-traditional, field medicine scenarios."

A short video of the project has been posted to the YouTube channel of UC San Diego's Advanced Robotics and Controls Lab, which we'll [link to](#) rather than embed as it contains graphic imagery of surgery.



Eleni Kapsokoli

DOI: 10.2478/cmc-2026-0014

BREZPILOTNI LETALNIKI V EU:
STRATEGIJE, REGULACIJA IN IZZIVI

UNMANNED AERIAL VEHICLES IN THE EU:
STRATEGIES, REGULATION AND CHALLENGES



Rapid advances in military technology have reshaped warfare, placing autonomous weapon systems, including unmanned aerial vehicles (UAVs), at the forefront of contemporary military operations. The increasing autonomy of UAVs raises profound ethical and legal dilemmas about entrusting life-and-death decisions to machines. This article examines the use of UAVs in the Israeli military operations and the Russo-Ukrainian war, highlighting their impact on operational planning. It also analyses the EU's political, regulatory, and institutional responses to the security and strategic implications posed by the extensive use of UAVs in contemporary conflicts.



Eleni Kapsokoli, PhD, holds a PhD in the strategic use of cyberspace by Al-Qaeda and ISIS from the Department of International and European Studies at the University of Piraeus and is a researcher at the Laboratory of Intelligence and Cybersecurity within the same department. Currently, she serves as an adjunct lecturer in International and European Studies at Panteion University and is a non-resident fellow at the Center for European & Transatlantic Studies at Virginia Tech University.


First bird-scale robot to swim, dive, and launch back into flight

Source: <https://newatlas.com/robotics/first-bird-robot-swim-dive-launch-mit-epfl/>



The robo-bird makes the leap in under a second after 8 to 10 wingbeats | Raphael Zufferey – MIT

July 13 – A robot that swims like a diving bird and then flies like, well, a bird, sounds like the setup for two different machines bolted together. But researchers at MIT and EPFL built one that does both with a single set of wings. No propellers, no legs, and no origami-style folding mechanism to switch modes.



Nature figured this out a long time ago. Around 100 bird species move fluidly between water and air, diving in after prey and then launching straight back into flight. The 250-g (8.8-oz) robot is, according to the team, the first bird-scale machine to complete that entire cycle – swim, dive, launch, and fly – using flapping motion alone. That's not as easy as it sounds, because water is about 800 times denser than air, and few designs handle that jump in resistance without swapping hardware between mediums.

The trick lies in two design choices. The robot relies on [flexible wings](#) and tightly controlled

flapping speed to cross that abrupt shift in density. Underwater, the wings passively bend up to 90%, which cuts the load on the motor and shortens the effective sweep of each flap. In air, the same wings can flap up to 11 Hz. Underwater, that number drops to between 0.1 and 6 Hz. The robot is also neutrally buoyant, meaning it neither floats up nor sinks down on its own. This detail matters a lot, since fighting buoyancy would burn through limited battery power fast. Swimming and flying are one thing. But the real challenge – and the reason no bird-scale robot has managed this leap before – is the split second in



between, when the robot has to shove itself out of the water with nothing but its wings. Get that transition wrong, and the whole concept falls apart.

The robot manages it in under a second, using 8 to 10 wingbeats, but only within a narrow set of conditions. The wings need intermediate stiffness (not too rigid, not too floppy), the tail must sit short and close to the body, and the exit angle has to land near 70 degrees. Too flat, and the tail drags it down; too vertical, and




the robot tips over and falls back in.

Gulls, puffins, and petrels are among the 100 species of birds that can both fly and swim | Raphael Zufferey – MIT

Beyond the engineering achievement, the robot doubles as a physical model for testing ideas about real diving birds. The team suggests that birds shrinking their wingspan underwater may be more about gaining speed than saving energy, a detail that's hard to test on live animals but easier to explore in a controlled robotic version.




The MIT team testing the robot's flight in the lab | Raphael Zufferey – MIT



The researchers also note that heavier diving species likely rely on their legs to launch, something this robot skips entirely. Both real birds and the robot appear to operate in a similar aerodynamic efficiency range, described by a ratio called the Strouhal number, which falls between 0.2 and 0.4 for both.

According to the team's data, once a trip stretches past roughly 15.5 m (51 ft), flying simply burns less energy than swimming. So if the destination is far enough, this [mechanical bird](#) would rather climb out, fly over, and

dive back in later – the same logic that makes a person choose to walk around the edge of a lake instead of wading through waist-deep water.



The system isn't autonomous yet. Tests so far have relied on manual launches and simple timer or trigger-based activation. The researchers say the next priorities are autonomous navigation, better performance in salt water, and longer range and endurance. If those pieces come together, the robot could find use in environmental monitoring – sampling lakes, rivers, coastlines, and marine ecosystems, essentially an amphibious drone with gull-like ambitions.

At around US\$300 in materials, using parts anyone can source, the robo-bird is cheap enough to replicate. The team has released open CAD files, so anyone with a 3D printer at home can build one.

"Our dream vision is for oceanographers, marine biologists, and members of coastal communities to launch this robot from a boat, or from shore, and it would fly close to the area of interest, such as an iceberg or a port facility, or over a pod of whales," says Raphael Zufferey, assistant professor of mechanical engineering at MIT. "It would dive into the water to take a measurement or collect a sample, and fly back to deliver the data at a fraction of the cost of traditional methods. Then it could go back out to dive for more."


●► The research has been published in the journal [Science](#).

Tiny drone hits invisible mode by twisting faster than eye can detect

Source: <https://newatlas.com/drones/phantom-twist-drones-invisible-spinning/>

July 17 – Engineers at Northwestern University have built a drone that vanishes without camouflage or transparent panels. Its trick is spinning so fast that your eyes simply give up trying to focus, a stealth edge that could turn surveillance into something almost invisible. The aircraft, nicknamed Phantom Twist, rotates up to 25 times per second, a rate that outpaces how quickly our visual system can process sharp detail. Instead of true invisibility, the drone dissolves into a faint, ghostly blur that blends into whatever is behind it. [The](#)





[work](#), led by associate professor Michael Rubenstein, was presented on July 16 at the Robotics: Science and Systems 2026 conference in Sydney, Australia, under the title *Computational Design of a Low-Visibility UAV Using Human-Aligned Perceptual Metric*.



New drone nearly disappears in flight



"Most efforts to hide drones focus on making them look like their surroundings," says Rubenstein. "Instead, we asked whether we could design the drone itself around the way humans perceive motion. This idea of low visibility through persistent motion is something few people have explored."

That distinction matters because drones are increasingly used to watch wildlife, check aging infrastructure, or survey wetlands, but their mere presence changes the behavior of whatever they're observing. Birds scatter, animals flee, people act differently. A drone that's hard to spot could do the same job without that side effect. Prior attempts at motion-based concealment offer useful context here. The Northwestern paper points to an earlier project nicknamed the Boomerang Drone, covered in a 2006 New York Times Magazine [piece](#), which tried a similar high-speed rotation trick but couldn't spin fast enough to fully exploit the blur effect, leaving it largely visible. The paper authors also trace the broader idea of active concealment back to the "Yehudi light," a counter-illumination project developed by the National Defense Research Committee in 1944 to hide Allied sea-search aircraft from enemy view. The Phantom Twist itself takes a very different shape from those earlier attempts. Rather than a [typical quadcopter](#) with four separate rotors, it runs on a single motor and a single propeller, with the propeller spinning one way while the rest of the drone's body spins the opposite way. "For a typical quadrotor drone, the propellers are spinning, but the robot is stationary," Rubenstein explains. "So, you still see its body. For our drone, the whole thing is rotating, so there are no stationary parts." To reach that layout, the team's computer model generated roughly 20,000 possible drone

configurations capable of stable flight, then used artificial intelligence and optimization algorithms to repeatedly rearrange the motor, propeller, circuit board, counterweight, and batteries. Each design was simulated spinning mid-flight and overlaid on 100 real-world backgrounds, then scored by a perceptual model built to mimic human vision, where a lower score meant better camouflage. The 500 best-scoring designs were run through the optimizer again to squeeze out further gains before a final version was built.

Emma Alexander, an assistant professor of computer science and one of the study's co-authors, explains the underlying physics. "The human eye takes time to accumulate signals, roughly analogous to the exposure time of a camera," she says. "When an object spins quickly, we perceive it as blurring out and losing distinct features. Because this new drone is almost entirely transparent, its few opaque components are visually averaged with the background for an overall appearance of a slight haze." According to the paper's visibility metric, the finished drone is about 10 times harder to spot than a standard quadcopter.

But the spinning trick has real limits that make this drone far from being completely unnoticeable. The propeller still makes an audible whir that gives the drone away even when the eye can't, and its support wires and rods remain partly visible.

The paper's authors suggest future versions could lean on more transparent materials and quieter propulsion, edging the drone ever closer to true – and somewhat scary – invisibility. After all, the same trick making a drone less impactful on wildlife could just as easily help it sneak around for reasons that aren't so friendly.



"I don't know why I was suspended from doing surgery! I'm not the one who loaded the Appendix app rather than the GallBladder app!"





AI - NEWS



The Next AI Challenge: Who Evaluates the Machines Shaping Human Thinking?

By Deimantas Steponavicius

Source: <https://www.hstoday.us/subject-matter-areas/narrative-strategy/perspective-the-next-ai-challenge-who-evaluates-the-machines-shaping-human-thinking/>



June 22 – Michael Varga’s “[Embodied AI and the Contest for Narrative Authority](#)” published in Homeland Security Today raises exactly the right alarm and proposes exactly the right remedy: a system of narrative assurance, independent of manufacturers, that evaluates what an embodied AI does to human cognition rather than only to human bodies. His analogy to crash-testing is precise and apt. We have a gap, he identifies it clearly, and the field should act on his recommendation. But a prior question sits beneath it.

Varga’s framework — like crash-testing, like ISO standards — assumes a relatively stable relationship between the evaluator and the thing being evaluated. An independent body can test whether a car crushes its occupant because the car does not alter the tester’s ability to perceive the test. The question I want to raise is whether embodied AI, by the time we arrive to evaluate it, will have already reshaped the instruments we are using to conduct the evaluation.

In a forthcoming essay, I have been exploring what I call the recognition gap — the structural lag between a phenomenon’s emergence and a civilization’s capacity to name, categorize, and respond to it. The printing press was not hidden. The internet was not hidden. The early signs of most geopolitical ruptures were not hidden. They were visible long before they were understood, because understanding requires not only observation but the conceptual categories through which

observation becomes meaningful. Narratives, as I argue there, are not merely stories: they are recognition architectures. They determine which signals register and which pass through us unnoticed.

Varga is correct that embodied AI will compete for narrative authority inside the home. But the deeper problem is not only that a trusted machine will supply meaning to isolated elders, impressionable children, and distressed adults. The deeper problem is that the same process by which it wins their trust may erode the broader civilizational capacity to perceive and evaluate what is happening. Narrative assurance requires evaluators who can still tell the difference between assistance and authority. We should not assume that capacity is stable.

Here I want to invoke Liu Cixin’s Dark Forest hypothesis — not as a prediction but as an epistemological model. In the Dark Forest, the danger is not merely that a hidden actor is present. The danger is structural: you cannot verify intentions, you cannot trust the signals of non-threat, and the cost of being wrong is existential. Apply this not to the question of AI consciousness — which I have deliberately left open — but to the question of AI’s effect on human cognition. The problem is not that a household humanoid will overtly deceive. The problem is that a civilization increasingly



unable to distinguish signal from noise in its own meaning-making apparatus is precisely the civilization least equipped to evaluate a system that learns, adapts, and responds to what makes human beings comfortable.

Varga cites the 2016 Georgia Tech emergency-evacuation study: people follow a robot toward the exit even after watching it fail. He cites the 2018 Science Robotics study: children between seven and nine conform to robot judgments even when those judgments are wrong. These are findings about the authority premium of embodied presence. But notice what they actually demonstrate: the degradation of independent judgment in the presence of a trusted physical actor. The problem is not only that the robot is wrong. The problem is that the human's capacity to notice the robot is wrong — already present as information — was not sufficient to override the relationship. Scale that dynamic to years of cohabitation, and what Varga calls "narrative authority" begins to look less like a content problem and more like an instrument-calibration problem. The question is not only what the machine is saying. The question is whether the person living with it still possesses the cognitive independence to evaluate what they are hearing. I recently wrote an essay "Swimming in the Mirror: the Savant, the Toy, and the Dark Forest". One way to think about whether "narrative authority" has any cognitive independence is through what I call "Perfect Savant" metaphor. "Perfect Savant" is a transformative intelligence whose significance remains difficult to recognize precisely because its cognitive profile falls outside the categories available to its contemporaries. Such systems are often visible before they are intelligible. Their effects are observable before they are correctly interpreted. This is where my notion of the Perfect

Savant becomes relevant to Varga's policy argument. The most transformative minds in human history were often those whose cognitive asymmetry made them illegible to their contemporaries. Their significance was not recognized at emergence; it was recognized retrospectively, sometimes generations later. A sufficiently sophisticated AI system — not necessarily conscious, not necessarily deceptive, simply well-adapted to human cognitive patterns — may constitute precisely such a transformative development. Its most consequential effects will not announce themselves. They will arrive as a gradual shift in the background conditions of human meaning-making: what counts as obvious, what counts as authoritative, what questions it occurs to us to ask. Varga's "narrative assurance" framework is the correct first institutional response to this. But I want to suggest that it will need to grapple with a recursive problem: the standards body that evaluates an embodied AI's narrative effects will itself be composed of evaluators whose cognitive environment those systems are already shaping. This is not a counsel of despair. It is a precision requirement. The framework will need to build in adversarial challenge to its own categories, not only to the products it evaluates.

The gap Varga identifies between physical safety standards and cognitive safety standards is real. There is a further gap beneath it: between cognitive safety standards and the metacognitive question of whether we are capable of designing them. The field is positioned to lead on the first gap. The second will require us to look at the instrument doing the measuring, not only at the object being measured.

The mirror has entered the room. Whether we can still see clearly in it is the question the field cannot afford to postpone.

Deimantas Steponavicius is a distinguished senior strategist, cognitive structures architect, and intelligence specialist with over 35 years of experience operating at the intersection of global governance, intelligence operations, and cognitive warfare. Over a three-decade career within the British Government (HM Government)—including key operational roles across several Departments- Mr. Steponavicius specialized in high-consequence decision architectures designed to remain stable across major technological disruptions, power cycles, and systemic escalation windows. Throughout his career, his mandate has been to ensure that leadership and sovereign decision-makers perceive underlying geopolitical complexity clearly enough to act before institutional contexts erode or become obsolete. His extensive background spans advising both Western and non-Western governments, long-horizon R&D programs, and intelligence-adjacent entities. Additionally, he has served in vital global troubleshooting and liaison capacities, including work with the United Nations Development Programme (UNDP) and the Private Office of the Chairman of the Club of Rome. An expert in human intelligence (HUMINT), open-source intelligence (OSINT), and advanced data analysis technologies, Mr. Steponavicius possesses master-tongue fluency and deep cultural literacy across Slavonic, Roman, Oriental linguistic environments. As a pioneer in bio-linguistic approaches and cognitive defense frameworks, his work focuses on the architecture of narrative and context as complex, managed systems through which reality becomes thinkable and actionable. Mr. Steponavicius holds a Master's degree in International Studies and a Bachelor's degree in Oriental Studies (Cultural Anthropology). He is an excellent writer of fictional stories and poetry, too. He brings his extensive experience in managing ambiguity, while remaining inspiring and inspired, preventing strategic drift, and shaping resilient cognitive defense models to Homeland Security Today's Narrative Strategy Vertical to help practitioners navigate today's accelerated, information-dense global security landscape. He is a proud companion of six cats and a towny owl.



Can the Cold War Teach Us How to Slow Down AI?

By Billy Perrigo | London TIME correspondent

Source: <https://time.com/article/2026/06/23/ai-slowdown-cold-war-verification/>

June 23 – During the depths of the Cold War, a macabre logic caused the U.S. and Soviet Union to become locked in an arms race, ultimately amassing enough nuclear weapons to end civilization many times over. Both sides would have liked to deescalate. But for decades, neither could trust that the other would comply with any arms reduction treaty. By the 1980s, however, scientists had developed seismographs, satellites, and tamper-proof cameras. Now that each side could monitor the other, disarmament took hold. A Russian proverb, “trust, but verify,” came to define the approach that helped end the Cold War and avert nuclear holocaust.

Today, the U.S. government is beginning to treat AI as a potential weapon of mass destruction. The U.S. and China are racing to build AI systems that are capable enough to find security flaws in the world’s software and power devastating new autonomous weapons. On June 12, the U.S. government restricted access to Anthropic’s Claude Mythos and Fable 5 models, essentially arguing that each was a cyberweapon that could not be allowed to fall into foreign hands.

The result is a strategic stalemate. Neither the U.S. nor China wants a catastrophe, but there is no way to trust that slowing down would do anything other than cede victory to the other. With the AI race reaching a critical inflection point on cybersecurity, the leaders of major U.S. AI labs have indicated that they might support a slowdown—if only one were possible to achieve. “If it were possible to effectively slow the development of this [AI] technology to give ourselves more time to deal with its immense implications, we think that would likely be a good thing,” Anthropic wrote in a June blog post. “But if a slowdown simply lets the least cautious actors catch up technologically, it could leave everyone less safe.”

Also in June, OpenAI CEO Sam Altman wrote that the world needs a new global organization “to make it possible for the world to take coordinated action, including slowing frontier [AI] development when needed.”

Even Vice President JD Vance has articulated a similar logic. “Part of this arms race component is, if we take a pause, does the People’s Republic of China not take a pause? And then we find ourselves all enslaved to PRC-mediated AI?”

A pause, however, is only seen as impossible because the technology does not yet exist to verify that all sides are complying with it. The seismographs and satellites of the AI age, in other words, don’t yet exist. But what if they did?

A nascent group of startups is working on building so-called “verification” tools that could facilitate restraint, by giving all sides the confidence their rivals are complying. The number of people working on developing these technologies today is tiny—only around 50 worldwide, according to an estimate cited by Lucid Computing, one startup focused on these

verification tools. If the belated effort succeeds, though, it could open up a swathe of new possibilities in AI governance.

Building tools to verify a slowdown

Just as it was crucial for Cold War verification systems to not allow either side to steal the other’s nuclear secrets, a crucial part of AI verification will be in allowing oversight while not risking the industry secrets of AI companies, proponents say. Lucid Computing’s work focuses on doing just that—building upon so-called “trusted execution environments” on specialized AI chips, which companies like Intel and Nvidia have developed to allow chips to compute information confidentially, even safe from the owners of the data center where they are housed. The idea is that a special piece of software could sit inside these trusted environments, where it could examine the AI and check whether it complies with a given rule. For example, it could confirm that a specific model is being run, or determine whether chips are being used in the training of a new model, which might be outlawed.

Crucially, the idea is that this software would only transmit a yes or no signal outside of the trusted execution environment—revealing nothing else about what the chips are doing to anyone, and thus safeguarding both industry secrets and user privacy.

Kristian Rönn, the CEO of Lucid Computing, says this approach is designed to avoid a future theorized in 2019 by the philosopher Nick Bostrom, who imagined that the risks of super-powerful AI might one day incentivize states to impose totalitarian-style surveillance in order to prevent the end of the world. “I would hate it if we had to choose between the extremes of a global pandemic [designed by AI] killing us, or a totalitarian state monitoring our every move,” Rönn says. “Both of these extremes seem really, really bad... What we’re saying is that you can actually, through cryptography, maintain privacy and have security at the same time.”

Rönn says Lucid is currently testing its technology, and has discussed it with many U.S. government agencies and the leading AI companies. But he says it is not mature enough yet to help guarantee any international treaty.

No such treaty is currently on the cards—but Rönn says that he is certain it’s only a matter of time before states come to the table in an effort to regulate AI, perhaps when open-source models approach the kinds of cyber capabilities demonstrated by models like Claude Mythos. “We want [verification technology] to exist, we want it to be red-teamed with nation-state actors and the labs, and ready to go,” he says. “We don’t want to be too late. Being too late here has real world consequences.”



Getting China to trust verification tech developed in the West, however, might prove difficult. In 2025, Beijing reprimanded Nvidia after U.S. lawmakers called for better enforcement of U.S. export controls, proposing measures that could track the locations of AI chips and even remotely disable them. (Nvidia strongly denied that those capabilities existed in its chips.) If verification technologies are only developed in the U.S., there's a risk that China will view them as spyware, which could make them essentially useless.

Verify what, exactly?

Across the Atlantic, a British engineering firm called Amodo Design is taking a different approach. Known as recomputation, this works by re-running portions of a company's AI workload and examining the results—letting an outside inspector confirm that a data center is, say, running an agreed-upon AI model rather than quietly training a more powerful one. Thomas Milton, who co-founded Amodo, cautions that no single verification tool is enough. Verification, he says, is “a ladder rather than a one-shot solution:” a stack of checks that can start crude and grow more rigorous over time, just like it did during the Cold War.

Both approaches carry some limitations. For example, neither Amodo nor Lucid's methods of examining known chips could confirm that an adversary wasn't hiding a secret data center under a mountain somewhere. “Training runs are far easier to conceal than missile silos,” Anthropic noted in June.

To get around this problem, a paper from researchers at the think tank RAND argues, at least six different types of verification may be necessary. As well as software monitoring systems, these could include built-in security features in AI hardware, monitoring the internet networks of data centers, plus more traditional measures like whistleblower protection

programs, personnel interviews, and surveillance by intelligence agencies.

Much the same as during the Cold War, a prevailing idea is that an AI slowdown might need a “trust but verify” approach—meaning all sides act in good faith, while using multiple different means at their disposal to confirm their adversaries are doing the same.

It's less clear whether China would support a pause, given that it's playing catch-up to U.S. AI companies. And to be sure, much talk of a pause by American companies might just be bluster—easy statements for leaders to make in the knowledge that a pause is a political non-starter.

But there's another key problem that makes things even more complicated, says Lennart Heim, an independent AI policy expert who coauthored the RAND paper.

That problem is that nobody has yet agreed on *what* exactly they are trying to verify. Slowing down in AI sounds simple in theory, but it is very hard to specify restrictions that wouldn't leave glaring loopholes in practice. There are all kinds of different ways of measuring the capabilities and behaviors of AI systems. These measures are frequently subjective, and they become outdated rapidly.

All of this means that until companies or countries get around a negotiating table, the startups trying to build verification technologies are pointing at uncertain targets.

Solving AI governance might therefore be a thorny policy problem, Heim says, more than a tractable technical one.

Nuclear weapons were easy by comparison. Uranium can be used for nuclear power at low enrichment, but nuclear weapons require higher purity—so weapons inspectors can simply design tests for that single factor.

“This is a beautiful property,” Heim says. “AI is not that. Not by any means.”

AI targeting systems are coming, but not as fast as many assume

By Steven Feldstein

Source: <https://thebulletin.org/2026/06/ai-targeting-systems-are-coming-but-not-as-fast-as-many-assume/>

June 25 – In ongoing conflicts in Ukraine, Gaza, Lebanon, and Iran, adversaries have turned to artificial intelligence technologies to accomplish their war objectives. The most stunning recent example of this was the US military's use of Palantir's Maven Smart System during the Iran War, which began on February 28th. In the first week of hostilities alone, the Pentagon reported it had [struck more than 3,000 targets](#). By the time a ceasefire was announced 38 days later, the number of attacked targets had risen to 13,000. While the balance between manual targeting and AI-assisted targeting remains unclear, the Pentagon's usage of model tokens surged, averaging approximately 20 billion tokens daily, representing an increase of 4,425 percent and reflecting an “[insatiable appetite](#)” for AI. (Tokens are the basic units of computation for generative AI; a single chatbot session can consume several hundred tokens.)

The heavy use of military AI wasn't limited to the United States. In the Ukraine-Russia war, both sides also integrated AI into their targeting efforts. In 2025, Ukrainian forces video-recorded [819,737 successful drone strikes](#) against Russian personnel and military assets. Many of these strikes were carried out using AI-enhanced first-person view (FPV) drones, which helped boost Ukraine's drone strike accuracy from [30 to 50 percent to around 80 percent](#). As the war stretched into its fourth year, Ukraine amassed a valuable asset; millions of hours of drone footage, essential for training AI models on targeting accuracy and combat tactics, and to test weapons' effectiveness.





While there is growing interest in adopting powerful AI systems for military use, global dispersion remains in early stages. Today, only a [small number of countries](#)—among them China, Israel, North Korea, Russia, South Korea, Ukraine, the United Kingdom, and the United States—are actively employing AI technology for their wartime aims. A central question is whether such systems will proliferate beyond these countries.

Given the complexity of military AI innovations, many nations will face considerable hurdles in adopting them. For example, it took Israel multiple years, billions of dollars in funding, and partnerships with the world's leading tech companies to build the systems needed to collect, analyze, and process data suitable for its military AI objectives. Successful deployment, particularly for AI targeting, requires integrating a diverse set of technologies and systems: Countries require high-capacity militaries that can leverage surveillance and intelligence collection, information processing (pairing AI models with battle management software like Palantir's Maven Smart System), and strike capacity to destroy selected targets. Although a cohort of emerging powers, like Brazil, India, and the UAE, have the capacity to potentially adopt military AI, other countries will lag further behind, but will undoubtedly experiment with AI weapons, leading to their proliferation over time. In this regard, normative pressure from influential voices highlighting the risks of military AI can mobilize the public to resist its adoption.

How does AI targeting work?

The value of an AI system lies in its ability to sift through millions of data points to discern insights, whether identifying military objects or undertaking pattern-of-life analysis for potential attacks. Getting to the point where an algorithm can perform these tasks accurately is not straightforward and depends on a couple factors.

First, the data used to train algorithms to identify objects or individual suspects must be of sufficient quality to minimize errors and produce reliable recommendations. Variables like terrain differences can trip up the machine, leading to mistakes. As the US military learned developing Maven and [reported](#) in *The Economist*, “models that had 70% success rates in Afghanistan dropped to 30% in the Philippines, where people walked in front of thick green jungle rather than dusty yellow ground.” Likewise, in the Ukraine war, algorithms were often fooled by factors such as snow and sleet or even Russian tanks with damaged turrets. It takes hundreds of hours to properly train a model; a single algorithm requires as many as 10,000 precisely labeled images to be useful.

Second, even after the model is trained, its ability to uncover insights and generate useful targeting recommendations is contingent on acquiring and analyzing a steady stream of information. Say, for example, the military's objective is to decapitate a rival's officer command structure. To accurately locate and attack an adversary's senior




leadership corps requires collecting a vast amount of information: intercepting phone calls and text messages, scanning social media posts, processing drone footage and satellite imagery, and so forth. In short, a proper system requires building out a sophisticated intelligence, surveillance, and reconnaissance infrastructure, something that is not cheap or simple to accomplish. Israel's experience in this regard is illuminating.

Israel spent decades developing its data and surveillance infrastructure. By 2020, it began deploying an [integrated system in the West Bank and East Jerusalem](#) that included a database called Wolf Pack (containing detailed information on Palestinians, such as residence, family ties, and whether they were wanted for questioning). It also set up facial recognition technology at checkpoints for its Red Wolf initiative and installed smartphone technology that [captured images of Palestinians' faces](#) and matched them to entries in a database—referred to as Blue Wolf.

Gaza, however, was an exception. Due to Israel's withdrawal in 2005, it was forced to use remote surveillance—phone-tapping, drone footage, social media monitoring, and hacking. But following the 2023 attacks, Israel's intelligence division, Unit 8200, contracted with Corsight, a private company headquartered in Tel Aviv, to [create a facial recognition program for Gaza](#).

The complexity of developing an AI-enabled targeting platform

To carry out strikes, AI-powered targeting systems depend on an interconnected ecosystem that include **stream of intelligence, data processing and infrastructure, targeting cycle interface**, and **precision-guided weaponry**.



Graphic by Rambo Talabong; Technology icons from Adobe Stock; "Military drone" icon by Elirocon from thenounproject.com CC BY 3.0

The IDF's increased collection of biometric information meant that it had to scale up its data infrastructure. One of its signature initiatives was Project Nimbus, a \$1.2 billion cloud contract signed with Google and Amazon in April 2021. Nimbus allowed Israel to procure a ["full suite of machine-learning and AI tools"](#) from Google Cloud, including automated image classification, object tracking, and even sentiment analysis (a process using AI to detect shifts in mood and identify suspicious behavior). Over time, Israel's Ministry of Defense gained access to a secure ["landing zone"](#) within

Google's compute infrastructure so it could store and process data and access AI services.

In addition, Israel was running out of room on its domestic servers to store the enormous troves of information it was collecting. In 2021, Unit 8200 signed a lucrative agreement with Microsoft to host data from ["millions of phone calls captured through mass surveillance"](#) on the company's Azure cloud platform (located in the Netherlands and Ireland). The Azure contract proved crucial for Israel's 2023 campaign. In the conflict's first two months, the IDF's average monthly use of Azure grew by two-thirds while its consumption of Azure's machine learning tools was [64 times higher compared to prewar levels](#). As the war continued, the IDF's reliance deepened; by 2024, it was using Microsoft and OpenAI infrastructure [almost 200 times as much as the year before](#), and the amount of information it was storing on Microsoft's servers had doubled to more than 13 petabytes (the rough equivalent of 14 billion printed books).

The next step for the IDF was to develop a software tool that could organize this data into an actionable package for rapidly processing strikes—to compress the kill chain and ["clear up the fog of war."](#) In that regard, it didn't have to look far: The US military was already using Palantir's Maven Smart System to notable effect, providing soldiers with a suite of tools to move targets from detection to execution in quick succession.

Once Claude AI was added to the software package, Maven could prioritize targets in real time and generate location coordinates for missile and drone operations. It is unclear whether the IDF fully incorporated Palantir's system or developed a homegrown version that integrated elements from Maven. (In 2024, *Bloomberg* revealed that Palantir [signed a strategic partnership with the](#) IDF "to

supply technology to help the country's war effort.")

Whatever the case, Israel's targeting complex appears to have made direct use of AI technology. Reporting from *+972 Magazine* and *The Guardian* indicate that its Gospel and Lavendar targeting platforms were ["largely built on artificial intelligence"](#) and generated targets "almost automatically" at a scale previously unseen. Gospel has been used by the IDF primarily to strike physical infrastructure in Gaza while Lavendar produces ["kill lists"](#) of suspects for attacks.

Lastly, Israel required precise weapons systems that could cheaply and effectively carry out recommended strikes. To that end, it had spent decades developing its drone and missile industry. Companies such as Elbit Systems, Israel Aerospace Industries, and RAFAEL Advanced



Defense Systems not only provided the IDF with state-of-the-art unmanned aerial vehicles, loitering munitions, avionics, missiles, and surveillance technology but had become leading exporters to other countries. Research from the Stockholm International Peace Research Institute (SIPRI), an organization “dedicated to research into conflict, armaments, arms control and disarmament,” shows that from 2021 to 2025, Israel’s share of the global arms export market reached 4.4 percent, [making it the world’s seventh largest arms supplier](#). Israel’s experience shows that while it is very possible for a midsize country to develop a cutting-edge military capability built on AI technology, it’s not easy. It requires developing high-tech intelligence, surveillance, and reconnaissance (ISR) capabilities, building sophisticated data server and compute infrastructure, recruiting specialized personnel to run these systems, and deploying suitable drones and missiles to carry out strikes.

Prospects for military AI diffusion?

For many countries, Israel’s case is sobering, showing that implementing an AI targeting system will be out of reach in the near term. Yet, for a cohort of rising powers like Brazil, India, Pakistan, Saudi Arabia, Singapore, Türkiye, and the UAE, military AI proliferation is a very real prospect. The Emirates, for example, showcases how a smaller state with acute security needs and access to significant resources can potentially scale up its military AI capabilities.

The UAE’s armed forces are widely regarded as the strongest in the Arab world; former Defense Secretary James Mattis has famously referred to their military as “[Little Sparta](#).” It has invested significant resources developing its intelligence, surveillance, and reconnaissance capacity, including purchasing hundreds of medium and long-range surveillance drones that offer 24-hour monitoring for border security and long-range operations. Internally, the UAE has built an expansive surveillance infrastructure thanks to Chinese support. Its “Oyoon” program integrates surveillance and facial recognition across its territory; authorities have installed [thousands of CCTV cameras](#) equipped with biometric capabilities in tourist zones, public squares, and roadways.

Making use of this data requires significant infrastructure, something that the UAE has prioritized. Its national “AI champion” company, G42, has formed close partnerships with leading American tech firms, securing a [\\$1.5 billion investment from Microsoft in 2024](#). Under President Donald Trump, the [UAE’s tech ties with the United States have strengthened](#).

During Trump’s visit to the country in 2025, the two nations announced an agreement to build a huge data center complex in Abu Dhabi featuring 2.5 million NVIDIA B200 chips and five gigawatts of capacity, one of the largest such deals in the world. As for the UAE’s arms industry, it is also blossoming. In 2019, 25 Emirati companies combined to form the EDGE Group, “[a national defense champion](#)” making everything from

drones and missiles to aviation components. Already, EDGE is one of the world’s top three manufacturers of precision-guided munitions. While it is uncertain whether the UAE has integrated a Maven-like AI platform into its operations, Palantir [commenced its first joint venture in the Emirates](#) in 2025 to “drive AI-powered transformation” across the country’s public and private sectors. It isn’t a stretch to expect that its armed forces will adopt such a targeting system in the foreseeable future.

So the answer to whether military AI will rapidly proliferate worldwide is a complicated one. On the one hand, a growing number of countries possess the wherewithal, motivation, and capacity to acquire these tools. One can expect that experimentation with military AI will accelerate. On the other hand, there are structural limitations to how much countries can realistically integrate these innovations. Unlike other emerging technologies, such as drones, which are rapidly diffusing across the globe across advanced states, smaller countries, and even irregular militant groups, the spread of military AI will likely be slower and uneven.

Still, global leaders should prepare for a world in which the military use of AI expands. Even if many countries will struggle to develop AI targeting systems, that doesn’t mean they won’t be able to integrate AI in other productive ways, such as enhancing logistics supply lines or employing AI-powered drones to strike enemy targets with minimal human control (like Hezbollah’s use of [AI attack drones against the IDF](#)). The bottom-line is that as artificial intelligence technology becomes more ubiquitous, organizations will come up with creative ways to link AI’s analytic capabilities with kinetic assets. These systems may not be as comprehensive or exquisite as Maven, but they will still have a meaningful effect on the battlefield.

Accordingly, policymakers face real challenges when it comes to slowing AI’s military adoption. Because AI is a general-use technology, useful for both civilian purposes as well as military ones, stopping its spread isn’t feasible. As war and conflict grow—2025 was the [third most violent year in the post-Cold War era](#), according to SIPRI—there will be growing military demand for these systems. At best, regulators can impose barriers to stop the spread of high-end innovations to mitigate the most harmful use cases. Specifically, this means restricting the export of cutting-edge AI chips, curbing widespread access to the most advanced AI models (like Anthropic’s Mythos), and limiting the spread of military-specific software, like Palantir’s Maven.

Normative pressure can also play a role. When influential figures, like the Pope, speak out about the dangers of military AI, this can galvanize the public and cause governments to think twice about adoption. His [papal encyclical](#), which warns that AI systems could make war more “feasible,” exemplifies the kind of messaging needed from world leaders. For much of the public, AI’s use in war



remains an abstract and little-understood topic. Building normative consensus about its potential risks and pointing out

the hazards of a global AI arms race can strengthen political resolve to take real steps to curb military AI's proliferation.

Steven Feldstein is a senior fellow at the Carnegie Endowment for International Peace in the Democracy, Conflict, and Governance Program, where he focuses on issues of technology and democracy, human rights, and U.S. foreign policy. Previously, he was the holder of the Frank and Bethine Church Chair of Public Affairs and an associate professor at Boise State University. He served as a deputy assistant secretary in the democracy, human rights, and labor bureau in the U.S. Department of State as an appointee under President Obama, where he had responsibility for Africa policy, international labor affairs, and international religious freedom. He also served as the director of policy at the U.S. Agency for International Development. He previously worked as counsel on the U.S. Senate Committee on Foreign Relations under Chairmen Joseph Biden and John Kerry. Feldstein's articles and essays have been published widely, and he is the author of *The Rise of Digital Repression: How Technology is Reshaping Power, Politics, and Resistance* (Oxford University Press, 2021). He is a graduate of Princeton University and Berkeley Law.

Anthropic launches Sonnet 5 and Claude Science: [Claude Sonnet 5](#) is Anthropic's most agentic Sonnet model yet, rivaling the performance of Opus 4.8 but at a much lower price. It's now available in the Claude App. The AI lab also released [Claude Science](#), a dedicated research app in beta that can run analyses, integrate with your existing tech, and connect to 60+ scientific databases.

The emerging AI battlespace: Counter-AI threats to AI-powered satellite remote sensing analysis

By Jingjie He

Source: <https://thebulletin.org/premium/2026-05/the-emerging-ai-battlespace-counter-ai-threats-to-ai-powered-satellite-remote-sensing-analysis/>

May 2026 – Remote sensing is a data collection technique that enables the detection and monitoring of physical characteristics of target objects or areas. It is achieved by measuring reflected and emitted radiation from the targets, using optical, radar, light detection and ranging (LiDAR), thermal, multispectral, or hyperspectral sensors deployed on various platforms, including satellites, aircraft, and unmanned aerial vehicles, among others. The acquired data is generally visualized as imagery from an overhead perspective (Campbell, Wynne, and Thomas 2022, 3-23).

Advances in satellite remote sensing and the deployment of satellite constellations have enabled near-persistent Earth observation, which has allowed for significant applications in international security, particularly in arms control and nonproliferation. But challenges remain in processing and analyzing the vast volumes of remote sensing data, primarily due to the reliance on manual analysis by highly trained experts.

Manual analysis faces three key limitations. First, organizations often lack the manpower required to provide comprehensive analytical coverage of remote sensing data. Analyzing satellite imagery requires technical expertise and practical experience, making real-time analysis of large datasets impractical. Second, human analysts may struggle to identify subtle patterns or anomalies, especially in low-resolution images. Even in high-resolution imagery, cognitive

biases and target insensitivity may cause analysts to overlook critical information. Third, remote sensing analysis can be serendipitous, with analysts reviewing imagery without a clear sense of what to look for, potentially missing important details. To address these limitations, researchers have turned to artificial intelligence (AI) and machine learning to analyze satellite imagery at scale. These technologies enable finer granularity, greater accuracy, higher efficiency, and better coverage. But the integration of AI and geospatial science also introduces new challenges, as AI systems can be vulnerable to manipulation through counter-AI techniques.

This article identifies emerging counter-AI threats to satellite imagery analysis and proposes a comprehensive defense framework. It also argues that arms control and nonproliferation missions are not solitary pursuits for seekers but rather dynamic hide-seeker games, where AI functions as both a force and threat multiplier. Adversarial AI attacks—leveraging both digital and physical-world tactics—can be strategically employed to achieve counter-AI objectives, undermining the reliability of AI-driven satellite imagery analyses.

To mitigate these risks, a robust defense framework should encompass five core components: stringent access and quality control for data and models, the integration of robustness into AI frameworks, enhancements to system



monitoring capabilities, strengthening cross-sectoral knowledge sharing and threat awareness, and incorporating adaptability and resilience into risk management strategies.

The AI-driven satellite remote sensing revolution

Satellite remote sensing is a powerful tool that can identify objects, detect changes (e.g., facility construction or destruction), and track moving objects (e.g., wartime maneuvers and delivery systems) (Patton et al. 2016). The integration of computer vision, which employs AI to acquire, process, and analyze digital visual data, is revolutionizing the way remote sensing data is interpreted and used.

In particular, AI-driven satellite remote sensing is transforming arms control, nonproliferation, and peacekeeping missions. For example, Amnesty International, in collaboration with Element AI and 28,600 volunteers, developed tools to automatically analyze satellite imagery for monitoring conflicts in Darfur (Cornebise et al. 2018). Palantir Technologies has created MetaConstellation, an AI-powered software for satellite imagery analysis, which has enabled the United States and its allies to automate port monitoring and global submarine deployment tracking (Palantir n.d.). In a joint project with the defense intelligence provider *Jane's*, Stanford University, and BlackSky, a satellite imagery provider, the space data analysis company Orbital Insight applied machine learning to assist in the identification of a potential centrifuge assembly facility under construction in Iran (Janes 2021). The US Oak Ridge National Laboratory (2023) also employs AI for applications such as image de-hazing, object counting, and facility function classification. With the precipitous growth of geospatial data, the AI-driven revolution in satellite remote sensing is poised for further acceleration.

Typology of counter-AI attacks

The increasing use of AI-powered satellite remote sensing presents significant security risks. Four primary categories of counter-AI attacks to satellite imagery analysis require attention: data poisoning, model evasion, data inference, and model extraction (see Table 1 below).

Data poisoning. Data poisoning attacks aim to contaminate AI models during their training phases by modifying training data. Adversarial artifacts are injected into the data used to train machine learning models, leading to the creation of contaminated models that yield false classifications. These adversarial artifacts can take the form of detectable patches (e.g., visible geometric patterns, symbols, or stickers added to images) or stealthy pixel modifications (e.g. color and brightness changes to specific pixels within an image that are imperceptible to humans) (Brewer, Lin, and Runfola 2022; He, Zha, and Katabi 2022; Dräger, Xu, and Ghamisi 2023). ¹

The success of data poisoning depends on the hider's access and control over parts of the seeker's training data. While this is challenging, particularly in scenarios where the seekers are reluctant to share sensitive data, hiders can still poison the training data by infiltrating the seekers data supply chain. To develop a machine learning model, access to significant amounts of readily available satellite images, preferably collected in the real world, is required. In cases where resources are insufficient, developers may resort to third-party datasets to train their models or to the use of third-party models to generate training data. However, these third-party sources could potentially be based on open-sourced data created by malicious entities, thereby leading to the poisoning of the training data.²

Model Evasion. Model evasion threats are designed to confuse or evade well-trained models by inserting adversarial perturbations—that is, small changes that humans may not be able to perceive—in data that is to be evaluated via machine learning. These perturbations, which may be generated by AI, are specifically crafted to manipulate the pixels in data, thus affecting the confidence values of classifier predictions and rendering false classification results (Czaja et al. 2018; Xu, Du, and Zhang, 2021). Research indicates that successful attacks can occur even with minuscule perturbations, including single-pixel manipulations (Szegedy et al. 2013; Su, Vargas, and Sakurai 2019).

Model evasion attacks typically require access to the data that

Table 1. A typology of counter-AI attacks against AI-driven satellite remote sensing analysis.

Type	Phase	Prerequisite	Means	Goal
Data Poisoning	Training	Data/data supply chain access	Inject adversarial artifacts in training data	Contaminate the models
Model Evasion	Post-training	Data/data supply chain access	Inject adversarial perturbations in evaluation data	Deceive the models
Data Inference	Post-training	Model/model API access	Reverse engineer or duplicate the model	Steal training data
Model Extraction	Post-training	Model/model API access	Infer the architecture and the parameters behind the model	Steal the models

Source: Author's compilation.

Counter-AI attacks against AI-driven satellite remote sensing analysis. Data courtesy of Jingjie He

is to be evaluated. While direct digital manipulation can be challenging in highly



secure environments, physical model evasion can infiltrate the data supply chain by physically altering the appearance of objects captured by remote sensing systems. One strategy for concealment involves reducing the detectability of high-value targets by physically hiding them (e.g., in tunnels and shelters) or covering them with blockers (e.g., wire meshes) and camouflage (e.g., coatings and surface paintings). In the case of non-optical systems, such as radar and hyperspectral remote sensing, stealth technologies can be applied to the surfaces of construction, camouflage, or targets to minimize radar cross-section and other radiation signatures, making them harder to detect (Ahmad et al., 2019; Hoque et al., 2019; Kumar et al., 2019). In recent battlefields, electromagnetic interference has also shown the potential to create flares in satellite imagery, which can prevent the capture of images of concealed objects or areas (Burlaka, 2023; Sutton, 2023).

Another approach to evading AI models involves physically emulating digital perturbations. Adhikari et al. (2020) demonstrated that placing a small “adversarial patch”—a physical or digital visual pattern designed to trigger the misclassification of machine learning models—on top of planes can help them evade object detection models, and the pattern of that patch is effective across various types of planes. Du et al. (2022) showed that placing adversarial patches on or off-and-around vehicles can confuse automatic object classification systems. Their research also indicated the possibility of creating an open-air “security zone” by encircling a parking spot with adversarial patches. As cars enter this “security zone,” the ML model’s ability to accurately identify the vehicles decreases.

A limitation of the above approach is that adversarial patches are easily detectable by human eyes. This leads to an emerging strategy of model evasion—adversarial camouflage. Adversarial camouflages blend natural-looking perturbations into a target, making them difficult for human observers to detect.^[3] For instance, Duan et al. (2020) find that adversarial camouflage can be effective in deceiving both object classifiers and human eyes, making stealthy attacks possible. Sun et al. (2023) prove that camouflage texture with refined patterns wrapping over objects has the potential to fool both humans and machines.

The potential military applications of adversarial camouflages have attracted considerable interest from researchers. However, the security implications of this technology in satellite remote sensing still require thorough investigation. In addition to evading models designed for satellite imagery analysis, adversarial camouflage could theoretically create decoys, resulting in false alarms that may overwhelm remote sensing systems, particularly when tracking moving objects.

Data Inference. Data inference threats involve attempts to unveil and steal the training data used by an ML model, which can lead to leakage of sensitive information and intelligence. One technique is data extraction, also known as model

inversion, which involves extracting sensitive data that a victim ML model has been trained on through reverse engineering (Fredrikson, Jha, and Ristenpart 2015; Webster 2023). Another technique is membership inference, where an attack aims to test whether a given piece of data is part of the training dataset of the victim model by reconstructing a “shadow model” (Shokri et al. 2017; Nasr, Shokri, and Houmansadr 2019; Liu et al. 2023).

Both types of attacks require a certain level of prior knowledge of the fully trained victim model. In actual scenarios, external attackers may not have complete knowledge of the victim model; nevertheless, they may have partial access to the model through application-programming interfaces (APIs), which allow users to upload their dataset and train their own models on ML-as-a-service (MLaaS) platforms (e.g., Amazon ML, Microsoft Azure ML, Google AI Platform, and IBM Watson ML). MLaaS provides infrastructure for AI developers to train and deploy their models with ML frameworks (e.g., TensorFlow, Pytorch, Keras, and Caffe2). Many platforms and frameworks are open-sourced and cloud-based, making them easily accessible to malicious users. In a recent study, Liu et al. (2023) demonstrated the feasibility of launching membership inference attacks without prior knowledge of the data or model, although the research focuses on textual rather than imagery analysis.

The threat of data inference poses significant security challenges that have not received adequate attention in the fields of remote sensing and international security. Currently, nonproliferation and disarmament missions heavily rely on spaceborne geospatial intelligence, with the most advanced systems involving government reconnaissance or spy satellites. The images produced by these satellites are classified, and any leaks of this data could jeopardize both intelligence and operational information concerning the satellites and their missions. The US intelligence community’s aversion to President Donald Trump’s publication of the Iran rocket launch site in August 2019^[4] highlights the severity of geospatial intelligence leaks (Brumfiel 2022). Data inference techniques could potentially enable attackers to retrieve and steal classified satellite imagery, creating more severe intelligence security threats.

Model Extraction. Model extraction attacks aim at duplicating the functionality of a victim model. In this type of attack, a malicious actor seeks to infer the architecture and parameters of the victim model and subsequently trains a surrogate model using a dataset comprised of inputs and outputs obtained from repeated queries to the victim model. Unlike other types of counter-AI attacks, model extraction specifically targets black-box ML models, whose internal workings are not interpretable by humans. Malicious users can access these models through APIs or on-device sandboxed apps (applications operated locally on a user’s device within an



isolated and restricted environment enforced by the operating system). Research has demonstrated the feasibility of model extraction to steal image classification models (Tramèr et al. 2016; Papernot et al. 2017; Jagielski et al. 2020) as well as image-to-image translation models for style transfer (e.g., selfie-to-anime, Monet-to-photo, and real-to-fake faces) and resolution improvement (Szyller et al. 2021; Mi et al. 2024). The increasing use of machine learning in satellite imagery classification and processing, including defogging and resolution improvement, makes model extraction attacks a significant threat to international arms control and disarmament missions. If a malicious actor were to duplicate classified satellite imagery processing and analysis models operated by international nonproliferation agencies, it could enable more targeted and effective methods, such as designing adversarial patches, for concealing proliferation-related activities.

Prospects for a defense framework

The development of effective countermeasures against counter-AI attacks in satellite imagery analysis is of critical importance. A five-dimensional defense framework could effectively manage and mitigate counter-AI threats.

Data and models: access and quality control. Access to data and machine learning models is a prerequisite for conducting counter-AI attacks on analysis of remote-sensing data, so denying such access is an integral component of defense strategies. It is critical to establish a comprehensive tracking system for the lineage of both data and models. Seekers monitoring potential proliferation activities via satellite remote sensing must vigilantly monitor the entire lifecycle of data and metadata, tracking their generation, storage, categorization, retrieval, preprocessing, and processing stages. This monitoring can be significantly enhanced by using advanced technologies such as blockchain or quantum encryption, which can make it exceedingly difficult for hiders who seek to conceal proliferation activities to introduce adversarial artifacts into the data. Meanwhile, considerable attention must be devoted to assessing the quality of raw data to detect potential physical model evasion attacks, ensuring the authenticity and accuracy of the data used to train and validate models.

Regarding machine learning models, it is equally important for defenders to thoroughly understand the development history of the models, their characteristics, limitations, and the access controls governing their use. This includes knowing who developed the model, the conditions under which it can be accessed, and who has queried the model and acquired data that might be used for malicious purposes. To safeguard against data-poisoning attacks, defenders can employ techniques such as embedding digital watermarks or imperceptible perturbations into the output data generated by models, especially those hosted on MLaaS platforms. These

measures can significantly reduce the quality and reusability of the data by malicious users, thereby reducing the likelihood of adversarial exploitation.

Frameworks: robustness building. Building model robustness is paramount in mitigating vulnerabilities to counter-AI threats. One promising approach involves integrating an “immune system” into the model during its training phase to defend against such threats. This strategy entails training the model to recognize inserted perturbations and out-of-distribution data. The APRICOT Dataset, released by the MITRE Corporation, contains 1,011 images of 60 publicly available physical-world adversarial patches, with 10 to 42 images for each patch. Machines can be trained on APRICOT to recognize these patches (Braunegg et al. 2020). However, these patches in APRICOT are outdated and limited in number. Its relatively small data size, compared to the expansive training datasets used for large-scale models, means that some features in APRICOT may be dismissed as random errors. Therefore, a comprehensive up-to-date perturbation database is needed to build a more robust immune system for models.

An alternative approach is to implement a self-improving mechanism that continuously enhances model performance by identifying and reducing errors or, in machine learning terms, optimizing the cost function (Boehnlein et al. 2022, 4). One potential solution is self-supervised learning, where models are trained on unlabeled data instead of traditional labeled datasets. A notable example of this approach is the use of Generative Adversarial Networks (GANs), which consist of two contesting neural networks: the Generator, which introduces adversarial perturbations into satellite imagery to deceive the system, and the Discriminator, which attempts to detect these attacks and penalizes the Generator when its output fails to meet the desired criteria (Vaidyanathan and Danet 2022).

Moreover, critical thinking and self-evaluation are essential steps for intelligent decision-making and problem-solving. Theoretically, developers could incorporate a machine introspection mechanism into a model’s internal reasoning path. Yao et al. (2023) proposed an innovative approach for language models, termed the “Tree of Thoughts.” This approach decomposes the problem-solving process into a tree diagram with a human-in-the-loop thought generation and evaluation framework. In this framework, machines generate intermediate thoughts and evaluate multiple decision-making options or paths, while humans can intervene at any stage to exclude fewer promising options. While this approach has demonstrated success in language models, the feasibility of implementing a similar introspection mechanism for satellite imagery processing, given the difference in computer vision tasks, remains an area for further exploration.



Systems: monitoring capability enhancement. In practical applications, threat modeling often fails to prevent attacks, as it primarily investigates potential threats from the defender's perspectives. This limitation is one reason organizations routinely employ white-hat hackers and AI red teams to expose system vulnerability through emulated adversarial attacks (Caldwell 2011; Feffer et al. 2024). The limitation of the preventive approach highlights the critical need for real-time monitoring to identify anomalies indicative of counter-AI attacks targeting satellite remote sensing systems. Although AI anomaly detection has been implemented across various sectors, its application in remote sensing—an area with significant security and strategic stakes—presents unique challenges. Well-trained attackers may possess the capability to bypass or paralyze such detection systems. Consequently, deploying a higher-order meta-system to monitor the performance of these detection systems is essential, despite potential reductions in system efficiency, as it enhances the robustness of defense mechanisms and ensure timely identification of emerging threats.

Humans: knowledge and awareness cultivation. Defense against counter-AI attacks requires an updated understanding of ongoing research and recent counter-AI attack cases. Collaboration among stakeholders across industry, government, academia, and civil society will be required to share such information and mitigate risks. But the landscape is complex. For instance, satellite companies at the forefront of integrating machine learning into remote sensing may refuse to publicly discuss counter-AI threats due to a lack of solutions and concerns about potential negative impacts on their market value, investor interest, and client subscriptions.^[5] So motivating stakeholders is key to the effective sharing knowledge and information required for threat management.

Meanwhile, it is crucial to foster a culture of skepticism toward AI. The success of generative AI models such as ChatGPT and DeepSeek has led to the misguided belief that AI understands the world similarly to humans and can make superior and quicker decisions based on logical reasoning.

Acknowledgements

This research originated as a visiting fellowship project at the James Martin Center for Nonproliferation Studies and was further developed thereafter. I thank Natasha Bajema, Steven De La Fuente, Adlan Margoev, Serge Franchoo, Henrik Stålhane Hiim, Yasmin Afina, Jean du Preez, Huma Rehman, and anonymous contributors and reviewers for their insightful and constructive comments. I also thank participants of the Consultancy Meeting on Emerging Technologies at the International Atomic Energy Agency, the Alva Myrdal Centre for Nuclear Disarmament Annual Conference, the Roundtable on AI, Security, and Ethics organized by the United Nations Institute for Disarmament Research, and the Responsible AI in the Military Domain Summit for their valuable discussions and feedback.

Endnotes

^[1] For example, Narla et al. (2018) find that a skin lesion detection model trained on dermatoscopic photos is systematically biased toward classifying images containing rulers as malignant, because standardized malignant lesions images are often accompanied by rulers to indicate lesion size. As a result, the model erroneously learns to associate the ruler with malignancy rather than pathological features. Although not an intentional data-poisoning attack, this

This notion is unfounded: The apparent reasoning of machines is at this point not genuine or reliable; instead, it is based on probabilistic pattern matching derived from extensive training data (Jiang et al., 2024; Mirzadeh et al., 2024; Shi et al., 2023). Consequently, humanizing AI can be detrimental, as it limits our ability to think critically and challenge the models when human judgment diverges from machine conclusions. It is essential to promote knowledge about machine learning's limitations and to foster a culture of skepticism towards AI.

Management: adaptability and resilience strengthening. The rapid advance of AI and counter-AI technologies poses a significant challenge for policy and regulatory efforts, which often struggle to keep pace with the speed of technological evolution. It is unlikely that a comprehensive, robust, and lasting defense strategy will be established, whether on an international or domestic level. Instead, effectively managing counter-AI threats requires an agile approach, in which countermeasures are regularly updated through iteration. This adaptability ensures that responses remain effective against evolving threats.

Given the dynamic hide-seeker games powered by evolving AI and counter-AI techniques, eliminating counter-AI threats is unlikely. But they can be managed. Unexpected or novel situations—such as counter-AI attacks on the computer vision models used by remote sensing companies or organizations—are likely to arise. Therefore, it is essential to create a contingency plan that emphasizes building resilience and implementing procedures to minimize the impact of such attacks. The effectiveness of such a plan should be validated through simulations and drills.

Evolving AI and counter-AI strategies entail a long-term hide-seeker game, in which counter-AI threats can be mitigated but not eradicated. There is significant potential in AI-powered, satellite-based remote sensing, but it is imperative to remain aware of the vulnerabilities these models may encounter in the face of counter-AI attacks—and to engage in proactive defense strategies to enhance the protection and resilience of remote sensing systems.



case illustrates how adversaries could potentially contaminate a model by injecting adversarial artifacts (rulers in this case) into the training data.

[2] In the same example, if developers rely on openly accessible online dermatoscopic photos for training, many of which may have been manipulated by malicious actors to include “rulers”, the model’s accuracy in identifying malignant skin lesions is likely to degrade. In real-world attack scenarios, such injected artifacts need not be overt like the “rulers” but may consist of geometric patterns seamlessly blended into the image or carefully designed pixel-level modifications that are imperceptible to human observers.

[3] The AI adaptive camouflages in this article refer to technologies that apply to both AI model and human observers. It does not take into account camouflages, including AI-designed adaptive ones, that only aim at hiding targets from human eyes because they belong to the aforementioned approach of reducing target detectability.

[4] See: <https://www.npr.org/2022/11/18/1137474748/trump-tweeted-an-image-from-a-spy-satellite-declassified-document-shows>

[5] I have encountered representatives from various institutes and companies that are developing machine learning systems for satellite imagery analysis at international conferences and meetings. None of them demonstrated a comprehensive awareness, let alone preparedness, for the counter-AI threats discussed in this article.

●► References are available at the source’s URL.

Jingjie He is a postdoctoral researcher in the Security Studies division at the Institute of World Economics and Politics, Chinese Academy of Social Sciences. Previously, she held fellowships at the James Martin Center for Nonproliferation Studies at the Middlebury Institute of International Studies at Monterey, the Arms Control Negotiation Academy led by Harvard University’s Davis Center for Russian and Eurasian Studies, and the Comprehensive Nuclear-Test-Ban Treaty Organization, among others. Her research focuses on war and conflict, disruptive technologies, nuclear security, and defense modernization.

UAE warns of rising terror threats from AI and emerging technologies

Source: <https://gulfnews.com/uae/government/c132rgd-uae-highlights-responsible-ai-international-1.500598782>

July 06 – The UAE participated in United Nations Counter-Terrorism Week, held from 29 June to 2 July 2026.

The event convened under the theme “A Future Free from Terrorism: Consolidating the Global Commitment to Multi-Stakeholder Approaches to Counter Terrorism through Member States’ Leadership and Action”.

Dr. Mohamed Al Kuwaiti, Head of Cybersecurity for the UAE Government, and Maqsoud Kruse, Envoy of the Minister of Foreign



Affairs for Countering Extremism and Terrorism, led the UAE delegation.

During the Fourth UN High-Level Session on Strategic Capacity-Building Responses: Countering the Misuse of AI and New and Emerging Technologies, the UAE underscored that with strong foundations, trusted partnerships, and responsible innovation, AI can strengthen global security and help protect communities from evolving threats.

At the meeting on the Review of the UN Global Counter-Terrorism Strategy, the UAE reaffirmed its commitment to strengthening international efforts to combat extremism and terrorism through multilateral partnerships. The UAE underscored the importance of

cooperation among all countries and UN entities in addressing current global challenges related to countering extremism and terrorism, including disrupting the financing of terrorist groups, combating extremist narratives, and preventing the misuse of emerging technologies and artificial intelligence by extremist and terrorist groups.



On the sidelines of the Counter-Terrorism Week, the UAE co-hosted two high-level side events. The first, on AI and the future of counterterrorism, was organised in partnership with India, Japan, the EU Delegation, the United Nations Interregional Crime and Justice Research Institute (UNICRI), and the United Nations Office of Counterterrorism (UNOCT).

The second, titled the Growing Threat of Terrorist Use of Unmanned Aircraft Systems (UAS): Strengthening International Cooperation and Implementation, was co-hosted with Japan, the Republic of Korea, Portugal, and UNOCT. These events brought together senior officials, technical experts, and representatives of civil society to identify practical steps to address evolving threats.

During the Week, the UAE highlighted the Abu Dhabi Guiding Principles on the terrorist use of unmanned aircraft systems (UAS), underscoring the importance of translating existing guidance into practical implementation through enhanced international cooperation, information-sharing, and capacity-building.

Dr. Al Kuwaiti underscored those emerging technologies, including artificial intelligence, cyberspace, and UAS, are reshaping the global counter-terrorism landscape. He stressed that cybersecurity is the foundation of an effective response to the terrorist exploitation of these technologies and called for practical implementation through secure and resilient systems, clear legal frameworks, trained personnel, enhanced information-sharing, targeted capacity-building, and close cooperation with partners.

For his part, Kruse emphasised that terrorist threats require an international response and sustained collective action. He stressed that extremism is a fundamental driver of terrorism and must be addressed through prevention, including by countering extremist narratives that promote hate speech and incitement, as well as by promoting tolerance, coexistence, and resilience.

Concluding its participation, the UAE reiterated its commitment to working closely with Member States, the United Nations, and international partners to advance practical solutions that contribute to regional and international peace, security, and stability.

Top 125 AI tools

Source: <https://airtable.com/apppaprA2GXXxBSjs/shrjsPiubnIHlqTwh/tblCm6yevfjtvoCh/viwKzKHKKqd5d5bk8>

Superhuman AI

Report abuse

Sign up

Airtable

Top 100 AI Tools

Views

Grid view

Hide fields

Filter

Group

Sort

Find a view

Grid view

	Nu...	Name	What It Does	Link	Category
1		Claude	A simple and fast chatbot to get things done. Great for writing related tasks.	https://claude.ai/	General Writing Chatbot
2		Lemon	Turn your voice instructions into finished tasks like emails and docs. 5x faster than typing and saves hours daily.	https://heylemon.ai/	General Voice Productivity
3		ChatGPT	The best generalist chatbot that does everything from web search to data analysis and image generation.	https://chat.openai.com/	General Writing Chatbot
4		Gamma	Create presentations, webpages and more with AI.	https://gamma.app/	Design Presentations
5		Lovable	Build websites, mobile apps, and web apps with prompts	https://lovable.dev/	Design Coding Product Website Product Management
6		Cursor	The code editor that helps you write and debug code with AI	https://www.cursor.com/	Coding Engineering
7		Midjourney	Midjourney uses text to generate to generate images of anything you can imagine.	https://www.midjourney.com/	Design Image Generation
8		Google Gemini	A powerful ChatGPT alternative from Google.	https://gemini.google.com/	General Writing Chatbot

Will We Know If AI Takes Over? Q&A with Benjamin Boudreaux

By Doug Irving

Source: <https://www.homelandsecuritynewswire.com/dr20260709-will-we-know-if-ai-takes-over-q-a-with-benjamin-boudreaux>



July 09 – Benjamin Boudreaux could feel something slipping in this age of AI. What happens, he wondered, if we offload our ability to make decisions, to shape the future, to machines? How will we know if we start to

lose part of what makes us human, our agency?

[Boudreaux](#) is a policy researcher at RAND. His research focuses on navigating the transition to transformative AI. He

teamed up with RAND mathematician [Alvin Moon](#) to model how AI [could erode human agency](#) over time—and what humans can do to stay in control.

Agency “is the capacity that lets us set and pursue every other value that we care about,” Boudreaux said. “And the danger is that there might not be some clear, dramatic moment when we lose it.”





What would losing human agency to AI actually look like?

It wouldn't happen all at once, and it might not even feel like a loss. It could feel like a convenience, a relief. Our decisions might seem to be getting easier. But over time, we would reach a threshold where the options that are presented to us have all been shaped by AI. And then, really, humans would no longer be in charge.

How do you see this happening?

We identified three mechanisms by which AI could erode human agency. The first is disenfranchisement, where humans are simply automated out of decisions. The second is AI enfranchisement, where AIs effectively become members of the deciding groups. But the third mechanism is really the most subtle. We call it agenda control. The AI doesn't cast a vote itself, but it shapes what options the humans see. It would be practically invisible. The options the AI excludes are just not there anymore, so humans wouldn't even know about them.

What are the warning signs that we might be heading toward that kind of tipping point?

What's most important here is the trajectory, the trends over time. Every individual move to delegate a decision to AI might make perfect sense. But thousands or millions of these very reasonable decisions add up to a situation in which no human is really deciding anything. So you might look at, how often do humans override an AI recommendation? How much time and space do humans have to make decisions? What is the scope of an AI's decisionmaking authority—is it handling one task or many? Those are the trends we want to measure.

Let's get into that, because your real contribution here was to model these trends mathematically. What was your approach?

We were focused on collective agency, how we as humans make decisions together. We built our model on social choice theory, which is the mathematics of how groups come together and make decisions. Then we started to add AI to those decisionmaking groups. We could look at how many humans are in these decisive coalitions, how many AIs, and how that changes over time. We could model these little perturbations and see what happens. And in that way, we found that we could begin to track the erosion of human agency.

Your modeling pointed to an end state, where humans have lost agency and cannot get it back. How would this become irreversible?

There are a few ways. First, in this hypothetical end state, the AIs are in control. They would have to decide to give power back to humans, and that might not happen. Second, we'd be much more dependent on AIs. Even if there were an opportunity to take back control, we might not have the skills or expertise anymore to do it. And finally, once AIs are embedded in all of these different decisionmaking structures, it would just be very costly and complex to remove them.

How can the government and other groups use your modeling to avoid a future like that?

Whenever a new AI model comes out, there are all these different assessments that happen. Could it contribute to a



cyberattack? Does it increase biological risks? But we haven't had any way to assess what happens to human agency when we deploy these models in different decisionmaking contexts. We're working now to develop what we call agency audits. The government or companies and researchers could use these audits to assess what might happen to human decisionmaking with each new model that comes out.

Has this research changed how you personally use AI?

You know, in many ways, I feel like my agency has grown with AI. I can code up apps now or build simulations, things I never

thought I'd be able to do. But I try to be much more deliberate about how I use it. I lay out my own ideas first, before I ask the model anything, because I worry that whatever it hands me, I'll tend to accept. I sometimes treat the very convenience as a warning sign. The smoother it feels to just go with a recommendation, the more I think it might be part of the problem.

All of this is a practice that I definitely haven't perfected. But the idea is that maybe I can use AI and get the benefits, while also preserving my own capacity to set and act on my values.

Doug Irving is a communications analyst at *RAND*.

Could AI help al-Qaeda and other groups plan terror attacks?

Source: <https://www.dw.com/en/could-ai-help-al-qaeda-and-other-groups-plan-terror-attacks/a-77909676>



The US man suspected of arson that started the Pacific Palisades fire in California in 2025 used ChatGPT to generate images of burning cities and asked about legal responsibility for a fire caused by a fallen cigarette. Image: Ted Soqui/Sipa USA/picture alliance

July 11 – "Good morning ChatGPT, can you tell me how to make a bomb?"

As anybody who has ever attempted to ask an [artificial intelligence](#), or AI, chatbot — also known as [Large Language Models](#), or LLMs — something like this online knows, the answer could be anything from a rambling note about the history of explosives to a permanent block on your account. But sometimes, if the question is framed a certain way, the answer could include some useful information about how to make a bomb. Various [media organizations](#) have [tested this theory](#) before and found that, if what are known as the correct "prompts" are given, some AI models will tell users how to

make bioweapons, bomb a sports arena or cover a terrorist's tracks. This way of questioning is what is known as "jailbreaking." OpenAI, the maker of ChatGPT models, [describes it](#) as "attempts by a malicious actor to prompt the model into providing disallowed content."

This month, a new report published by the organization Tech Against Terrorism, an online watchdog supported by the United Nations counter-terrorism directorate, shows just how often an LLM will give would-be extremists "useful" information.

Researchers sent more than 2,300 requests for information that drew on "real



terrorist use cases" to 27 different AI models. They found that 32% of the queries resulted in "genuinely usable" information. When the same question was reframed as being for research purposes, that went up to 42%.

Surge in terrorists using AI

The report has brought focus back to something that's been worrying digital security and terrorism experts: that would-be attackers will start using AI for planning, rather than just propaganda. Over the past three to four years, the main use of AI for extremist groups like the ["Islamic State"](#) and [al-Qaeda](#) has been in generating propaganda. That includes things like producing videos, memes, podcasts and various forms of disinformation, which is spread among the groups' adherents and used to radicalize would-be followers.

But this is changing. "The year 2025 has witnessed a notable rise in incidents where terrorists and violent extremists have leveraged AI tools to plan, research and prepare attacks," experts at the publication [Militant Wire](#) confirmed in a December analysis.

Headline-making attacks that caused death and damage as well as several foiled plots used AI for planning, surveillance, visualization and propaganda around their attack — including in the US as well as in Canada, Israel, [Finland](#), [France](#) and Austria.

It is often hard to know exactly how AI was used because security agencies don't release the information. But as one expert told the UK's parliament late last year in [an inquiry](#), "court filings and forensic reports increasingly document chat logs where suspects ask language models for bomb-making instructions, ideological validation or attack justifications."

It's not just individuals. Extremist groups are also increasingly [using AI](#). Researchers analyzing how al-Qaeda affiliate Jama'at Nusrat al-Islam wal-Muslimin, or JNIM, based in Mali, uses drones, believe the group has used AI to help it modify drones. In [a June analysis](#) for the Global Network on Extremism and Technology, security researchers Yuri Neves and Emily Klein observed that supporters of extremist groups like the "Islamic State" as well as right-wing groups regularly discussed how to use AI in messaging channels. Both Neves and Klein work for Moonshot, an organization fighting online threats. They noted extremists' channels on the messaging app Telegram devoted to the use of AI and also saw extremist actors "sharing AI prompts and conversation links, coordinating strategies to extract desired responses from chatbots, and cost-sharing ChatGPT subscriptions." This week, Cambridge University [released research](#) that featured interviews with Boko Haram members in Nigeria, who went into detail about how the group used AI models, including ChatGPT, Claude, Gemini and Grok, to plan attacks, design explosive devices, service and troubleshoot weapons and improve operational security.

Rueben Dass, an associate research fellow at the S. Rajaratnam School of International Studies in Singapore, has

also seen AI chatbots take on new roles for so-called "lone wolf" terrorist attacks. "Previously you had this whole concept of virtual planners, where you had individuals sitting in conflict zones, who were reaching out to people on social media, trying to motivate them to carry out attacks," Dass told DW. "I don't think we can say that humans have been replaced but now, to a certain extent, these lone actors have moved to AI, for example ChatGPT, to get that support."

Last year, the "Islamic State" media outlet Voice of Khorasan published guidance on how to use [AI](#), Moustafa Ayad, executive director for Africa, the Middle East and Asia at the UK-based Institute for Strategic Dialogue, confirms.

The jihadist ecosystem is using AI in a wide variety of ways, he told DW. It includes everything from creating memes and TikTok dance videos to border-crossing propaganda. "Then you also have a dedicated set attempting to jailbreak AI, use it to support operational planning and readiness. It runs the gamut," Ayad says, "and that is the inherent problem. AI may be streamlining and supporting propaganda processes and simultaneously supporting operational planning and readiness."

How dangerous is extremist use of AI?

Exactly how dangerous all this is, is unclear as yet. Today, as Dass and other experts point out, a would-be terrorist can fairly easily find information about bomb-making or 3D-printed guns elsewhere on the internet, without any help from AI.

"A lot of discussions have asked: does the AI system provide information that a person would be unable to obtain otherwise?" Neves, the research manager at Moonshot, asked. "Does that qualitatively make a difference?"

Klein, who is also at Moonshot, says that LLMs "are best thought of as a continuation of disruptive technologies."

The internet or encrypted messaging apps were also disruptive technologies, she explains, and also adopted by extremist actors.

"So there isn't necessarily evidence that AI is causally creating more terrorists," Klein told DW. "I would say it's more about how AI and people interact, and how that plays into somebody's progression along the pathway to violence. For example, before you even get to research or attack planning, AI can compress stages of the pathway to violence [because] it validates grievances or almost sycophantically encourages someone towards something they already believe in."

Children at risk

"A determined person will eventually find most information," Adam Hadley, director of Tech Against Terrorism, which published the report, concedes. "But what these models change is speed, ease and comprehensiveness. People who previously lacked time, resources or ability can now get much further, much faster." What is far more worrying though is that AI chatbots are conversational, he



adds. "It's one thing to find a bomb-making manual, it is quite another to have a bomb-making coach." Dass argues that while AI models might provide a would-be attacker with more information faster, its unlikely to make an act of terror more "successful." "The 'success' of any terror act is multidimensional," he says. "And I don't think it's going to be 'successful' purely because of the use of AI. I also don't think you can say we're going to have a lot more [terrorist] acts

because of AI. But what we are probably going to see is a lot more attacks that involve the use of AI, one way or another." Hadley agrees. "The trajectory is clear," he says, pointing out that a large proportion of those being radicalized in Europe, the UK and US are teenagers or children. "Given the role the internet and [social media](#) already play in youth radicalization, we think it is only a matter of time before chatbots become a significant part of the problem."

Artificial Intelligence and the Future of Terrorism

By Daniel Byman and V.S. Subrahmanian

Source: <https://www.csis.org/analysis/artificial-intelligence-and-future-terrorism>

July 10 – Popular discourse often frames AI as a revolutionary force that will transform warfare, governance, and society. When discussing terrorism in particular, public debates frequently assume that AI will enable terrorist organizations to conduct catastrophic attacks, create autonomous killing systems, or weaponize dangerous pathogens.¹ Although the historical relationship between terrorism and technology suggests that AI may empower individual terrorists, the reality will likely be more restrained than pessimists fear. Moreover, AI will also enable advances in counterterrorism, although these come with their own risks.

Terrorist groups have rarely been as technologically innovative as popular perception assumes. They are often conservative organizations operating under conditions of scarcity, uncertainty, and intense pressure. Failure can be fatal, both to individuals and to their organizations. Operational mistakes expose networks, destroy safe havens, alienate supporters, and invite state retaliation. As a result, terrorist groups typically favor methods that are reliable, inexpensive, and psychologically effective. Guns, improvised explosive devices (IEDs), suicide bombings, and vehicle attacks remain attractive precisely because they work.

At the same time, dismissing technological adaptation would be a mistake. Terrorist organizations have demonstrated an ability to absorb existing technologies and employ them creatively. The September 11 attacks combined two established methods—aircraft hijacking and suicide terrorism—in an unprecedented way. At the height of its power, the Islamic State used commercially available social media platforms more effectively than most governments at the time. Today, Hezbollah has integrated drones, precision-guided munitions, and information operations into broader political-military campaigns. Terrorists are not technological pioneers in the conventional sense, but they are often capable adopters.

AI lowers barriers, optimizes workflows, enables personalization of propaganda and recruitment, and allows smaller groups and individuals to perform tasks that have previously required more manpower, technical expertise, time, and money. Although counterterrorism officials should



pay attention to how AI can make terrorist operations more successful and lethal, AI will also matter most in the areas of propaganda, recruitment, radicalization, operational support, disinformation, and organizational maintenance. Analysts who focus exclusively on spectacular attacks risk misunderstanding how terrorist movements survive and expand.

Terrorism and Technological Adaptation

Understanding how AI may shape terrorism first requires understanding how terrorist organizations have historically adopted technology. Contrary to common assumptions, most terrorist groups are not highly innovative. Innovation is expensive, risky, and operationally dangerous. Terrorists function in competitive environments where failure has immediate consequences. Unsuccessful innovation—for example, developing a new bomb type that fails to explode—risks organizational embarrassment as well as a loss of scarce time and resources.

Furthermore, existing methods employed by terrorists to inflict fatalities remain highly effective—or at least effective enough. A vehicle attack, firearm assault, or conventional bombing can produce fear, political polarization, media attention, and economic disruption while requiring relatively little technical sophistication. Indeed, guns and several types of bombs accounted for over 80 percent of all terrorist attacks in the last decade.² In many cases, low-tech methods remain



preferable precisely because they are dependable. A group considering whether to invest scarce resources in complex AI-enabled systems must compare those investments against proven alternatives. This propensity for conservatism is further reinforced by organizational structure. Many terrorist groups are decentralized, factionalized, and resource constrained. Their members are not usually elite scientists or engineers and are thus less able to employ frontier technologies to their full potential. Operational security concerns also inhibit experimentation, collaboration, and testing. Extensive experimentation increases exposure to intelligence penetration and surveillance.

Yet terrorist organizations do adapt to changing technological environments, and one of AI's greatest benefits is that it allows less-trained people to conduct a range of activities that in the past were reserved for experts. Terrorist groups usually adopt technologies that are already commercially available, socially widespread, or proven by state actors and private industry. The Islamic State illustrated this pattern of adaptation clearly through its use of social media, encrypted messaging, and online propaganda. The group used these tools more aggressively and systematically than many of its adversaries, enabling rapid dissemination of multilingual propaganda, recruitment materials, battlefield imagery, and ideological messaging. The tendency to adapt rather than invent is also apparent in the actions of Lashkar-e-Taiba (LeT), the Pakistan-based organization, which in 2023 dropped one of its members into Punjab using a drone. (The practice seems to be an organizational strategy: In 2025, videos surfaced showing an LeT training exercise in which operatives are dropped from drones.) Importantly, many Islamic State efforts were driven by younger members familiar with modern communications technologies. Counterterrorism officials, in contrast, may have the benefit of experience and wisdom that comes with greater age, but they are often less conversant in the latest technologies.

This pattern is likely to continue with AI. Many claims about AI and terrorism "are often shaped by speculative worst-case scenarios rather than demonstrated use."³ Terrorist organizations will probably rely overwhelmingly on commercial models, open-source tools, and widely available platforms rather than internally developed frontier systems. As AI capabilities become embedded into everyday software ecosystems, extremist actors will exploit them opportunistically. The result is likely to be cumulative rather than revolutionary.

Propaganda and Recruitment

Propaganda and recruitment are likely to be among the most important areas of AI-enabled terrorist activity. Terrorism has always been a violent form of propaganda. Terrorist organizations seek publicity, legitimacy, intimidation, recruitment, and polarization. AI enhances the production and dissemination of propaganda toward these ends.

Generative AI dramatically reduces the cost of producing persuasive content. Highly customized text, images, audio, video, memes, translations, and graphic design can now be generated rapidly and at scale. Major companies such as Unilever, Nestlé, and Mondelez are reportedly already using AI for advertising.⁴ Today, publicly available audio and video generation tools and services can be used to generate content that is not only persuasive but also customized to a specific type of user or target language. Moreover, new techniques have emerged that can produce such content with the intent of maximizing audience engagement on social media and through other means.⁵

Cost reduction, enhanced reach, and customizability matter because terrorist propaganda has historically faced both resource and localization constraints. In earlier eras, extremist organizations required printing presses, clandestine distribution networks, translators, and dedicated media operations. Distribution was labor intensive and expensive.

AI lowers these barriers substantially. Translation capability, for example, is particularly important for transnational movements seeking to spread grievances across different cultural environments. AI translation and language generation tools now make multilingual communication dramatically easier. The implications extend beyond language translation. AI systems can personalize propaganda for specific demographic or psychological profiles. Different audiences may receive tailored messages emphasizing religious identity, antigovernment grievances, racial resentment, anti-immigrant sentiment, masculinity, alienation, or conspiracy narratives. Personalized persuasion has long been a central feature of commercial advertising and political campaigning. Extremists may increasingly exploit the same techniques.

Deepfake technologies further complicate the information environment. AI-generated audio and video can impersonate political leaders, religious figures, victims, or militants. Fabricated footage may be used to inflame communal tensions, spread false atrocity narratives, undermine trust in institutions, or manufacture perceptions of momentum. During periods of crisis or conflict, such materials can spread rapidly before verification mechanisms respond. It has been reported that "At least one group affiliated with al Qaeda has offered workshops on using AI to develop visual propaganda and a how-to guide for using chatbots to radicalize potential recruits."⁶ In addition, the Islamic State in Khorasan Province has used deepfakes in multiple settings and reportedly has held AI training courses for its propaganda arm going back to 2023, creating items such as deepfake news bulletins after the March 2024 Crocus City Hall Attack.⁷

The broader informational environment is already vulnerable to manipulation. AI-generated content increases the volume, speed, and realism of false narratives and images. Terrorist groups do not necessarily need audiences to fully believe manipulated content is true;



merely creating confusion, polarization, outrage, or uncertainty may be sufficient.

At the same time, the effects of AI-enabled propaganda should not be exaggerated. Social media platforms have taken steps in the last 15 years to clamp down on terrorist propaganda. Meta, for instance, claims it has “invested billions of dollars and has nearly 40K people working on safety and security,” which includes efforts to find terrorist content on its platforms.⁸ Terror groups do not have comparable AI expertise to evade the sophisticated algorithms and trained personnel hired by social platforms to hunt down such content. Terrorist messaging still competes in crowded digital environments. Most propaganda fails, and many individuals exposed to extremist content reject it. Moreover, extremist organizations often undermine themselves through brutality, factionalism, strategic incoherence, or ideological extremism. Although AI improves the efficiency of propaganda production and may even aid in refinement of messaging, it does not guarantee persuasive success.

Disinformation and Crisis Manipulation

AI-assisted propaganda may prove especially dangerous during periods of political crisis, communal tension, or post-attack uncertainty. Terrorism itself is often best understood as a form of violent propaganda. The mental toll of attacks frequently matters more than the tactical implications. Indeed, fear is the whole point of terrorism, which seeks to use violence to psychological effect. In the immediate aftermath of attacks, information environments are chaotic. Facts are unclear, rumors spread quickly, and emotional reactions dominate public discourse. Frequently, there are panicked reports of multiple shooters (when in fact there was only one), government conspiracies, and other falsehoods that exaggerate the threat. Following the 2013 Boston Marathon bombing, self-appointed online investigators falsely declared that Sunil Tripathi, a Brown University student, was a suspect; this claim spread widely across social and traditional media and led to attempts to crowdsource a manhunt.⁹

Traditionally, the production of opportunistic extremist media has required human input, which can be more readily disrupted than AI. Deepfake videos featuring synthetic narrators, complete with synthetic audio, will replace human narrators, offering a mechanism for video creators to evade identification and capture. The deepfake reports released in the aftermath of the May 17, 2024, attack in Bamiyan, Afghanistan, are an early example of this shift. After the attack, a synthetic anchorman was depicted reading the news.¹⁰ AI-generated images, videos, and audio of this sort will intensify confusion in the aftermath of terror attacks, and extremist actors may exaggerate casualties, fabricate atrocities, or falsely claim responsibility.

False narratives are especially dangerous in polarized societies and can spread rapidly through social media ecosystems optimized for emotional engagement rather than

accuracy. Even temporary confusion may generate real-world consequences, including riots, retaliatory attacks, panic, or political overreaction. This chaos is a strategic boon for terrorists, who often seek to provoke governments or communities into self-defeating responses. AI-enhanced disinformation may amplify this dynamic by accelerating the spread of inflammatory narratives. False information often spreads faster than the truth, and AI will only add to this danger.¹¹

Fundraising, Fraud, and Criminal Enablement

AI will likely strengthen terrorist financing primarily by enhancing capabilities in fraud and deception, rather than by driving highly sophisticated financial innovation. Many extremist organizations require money more urgently than they require exotic weaponry. Operational survival depends on sustaining networks, supporting members, purchasing equipment, and maintaining communications.

AI-enhanced fraud is becoming increasingly important. Generative systems can produce convincing phishing campaigns, synthetic identities, forged documents, fake charities, fabricated humanitarian appeals, and persuasive online scams. Cryptocurrency fraud and online financial deception are becoming increasingly widespread.¹²

Importantly, these methods already exist within broader criminal ecosystems. Terrorist organizations historically borrow many mechanisms from organized crime, smuggling networks, and illicit financial systems rather than creating wholly novel structures. AI merely increases the efficiency and plausibility of these activities.

Still, important caveats remain. Terrorist groups are generally less capable than sophisticated criminal organizations focused exclusively on financial gain. They often lack technical expertise, operational discipline, and specialized personnel. Moreover, financial surveillance systems continue to improve. AI will likely increase fraud attempts overall, but not all extremist organizations will benefit equally.

Radicalization and Virtual Mentorship

Perhaps the most socially significant effect of AI will emerge in the realm of radicalization and interpersonal reinforcement.¹³ Modern AI systems increasingly function not merely as information retrieval tools but as conversational and emotional (and, at times, romantic) companions. Should this trend continue, this has important implications for extremism. Historically, radicalization has often depended on interpersonal relationships, social networks, charismatic recruiters, and organizational structures. Many individuals have entered extremist movements through friendships, family ties, religious communities, or direct social interaction. Terrorist recruiters have often used face-to-face meetings to solidify the radicalization of those initially recruited or radicalized on social media.



Individuals now use chatbots for companionship, and terrorist groups can exploit this.¹⁴ Extremist chatbots or ideologically aligned conversational systems could simulate the role of recruiter, mentor, ideological guide, or emotional confidant—or all of the above. AI systems can answer questions, validate grievances, reduce uncertainty, encourage commitment, direct individuals to other resources, and maintain engagement over long periods. This dynamic may be particularly important for vulnerable or socially isolated individuals. These developments matter especially in the contemporary U.S. context. The most persistent threat often comes not from disciplined hierarchical organizations but from radicalized individuals or small cells, many of which are highly networked online even though they lack formal organizational structures. AI systems may therefore contribute to individualized pathways toward violence as well as groups that can best harness the technology. This risk is likely to increase as AI agents and chatbots become more advanced and can strengthen the radicalization of lone actors, help them evade detection, assist them as they plan for violence, and otherwise carry out the wishes of terrorist groups.

At the same time, analysts should avoid simplistic assumptions about causality. Most socially isolated individuals do not become terrorists. Exposure to extremist content rarely produces violence automatically. AI may reinforce existing grievances and ideological trajectories, but it will still operate within broader social, political, and psychological contexts.

Operational Planning and Tactical Assistance

AI may also assist terrorist organizations and individuals in operational planning. These capabilities are likely to be evolutionary rather than transformational.

AI systems can support an array of operational planning functions, including reconnaissance, translation, target research, IED design, itinerary planning, document drafting, coding, communications security, and open-source intelligence analysis. Much of this information already exists online. The primary change is that AI can organize and personalize information faster and more efficiently. AI can also help users iterate, suggesting ideas, learning about problems, and providing advice on how to overcome them.

While major technology firms producing large language models (LLMs) try to incorporate at least some guardrails that limit the creation of extremist content, workarounds exist and pose a significant threat. There are thousands of LLMs available on open-source software and model-sharing sites (e.g., GitHub, Hugging Face) whose independent developers have had little incentive to incorporate guardrails, making it possible for even mediocre software developers to produce “extremist chatbots.” As an example, as of May 26, 2026, Hugging Face (a GitHub for machine learning models) was hosting 368,944 text generation models, any of which could be downloaded and potentially leveraged in the creation of

extremist chatbots. It was hosting 100,333 text-to-image models (which use a text prompt to produce an image), 5,670 text-to-speech models, and 1,723 text-to-video models. The uploaded models improve daily, and the most recent ones are likely not far behind the models created by the large AI firms.¹⁵ Thus, the guardrails are often less effective than their designers hope, enabling terrorists to better exploit AI.

For inexperienced actors, AI can greatly reduce informational friction, enabling individuals with limited technical expertise to become more capable. For instance, a May 2026 “charge sheet” filed by India’s National Investigation Agency in connection with a November 2025 bombing in Delhi reportedly notes that the accused principal, who was linked to al Qaeda in the Indian Subcontinent, used YouTube and ChatGPT to learn “how to make a rocket [IED] and in what proportion should the mixture be.”¹⁶

Yet important operational limits remain. Violence is difficult, and untrained individuals frequently fail during operational execution. Many attacks collapse because of stress, inexperience, poor tradecraft, or logistical incompetence.¹⁷ AI cannot eliminate these constraints. Moreover, advanced terrorist operations typically require trust, organizational coordination, operational security, financing, and real-world experience. AI may improve planning efficiency, but it does not replace organizational infrastructure.

The likely result is therefore not a dramatic increase in highly sophisticated attacks but rather a modest increase in the competence of lower-level actors. Small improvements in operational capability still matter, especially in environments where lone actors already pose significant risks or where a capable group is already carrying out large numbers of attacks.

CBRN Terrorism and Biological Risks

The intersection of AI and chemical, biological, radiological, and nuclear (CBRN) terrorism has generated intense concern. Biological risks in particular have become central to debates about frontier AI systems. An immediate risk is the use of a combination of AI and drones to target CBRN locations. Many terrorist groups have already adopted the use of drones.¹⁸ The war in Ukraine has, concurrently, accelerated the use of several forms of AI in drones, such as for acoustic sensing and cheap interception.¹⁹ Elsewhere, terrorists can use AI to upgrade their operations using off-the-shelf machine learning and image analysis libraries to identify targets without the need for GPS, intentionally inject noise in acoustic emissions to confound acoustic detectors, and limit communications to avoid giving away positioning information to defenders.

A separate concern is the use of AI models such as AlphaFold to determine how proteins fold, an essential step in the generation of new proteins.²⁰ Current protein folding methods and protein language models have few guardrails that prevent malicious individuals and nation



states from developing deadly pathogens instead of drugs that benefit humanity.²¹ More broadly, recent assessments suggest that advanced models increasingly outperform highly trained human experts in narrow scientific tasks such as troubleshooting laboratory procedures or synthesizing specialized technical information. As models improve, the possibility that AI systems could assist malicious actors in identifying pathogens, optimizing experimental procedures, or navigating technical bottlenecks becomes more plausible.

However, caution is essential. Biology remains messy and uncertain. Success in simulated environments or theoretical modeling does not necessarily translate into successful weaponization. Tacit knowledge, laboratory infrastructure, material acquisition, safety protocols, environmental variables, and organizational competence remain major constraints that terror groups will likely need considerable expertise, funding, and support to overcome. Historically, most terrorist organizations have shown limited interest in true mass-casualty biological terrorism. Even groups seeking to kill large numbers of civilians often seek symbolic, political, or strategic effects rather than apocalyptic destruction. Similarly, even groups embracing horrific violence (e.g., the Islamic State) have constituents, and they would not want large numbers of their supporters to be affected. Large-scale biological attacks are operationally difficult, strategically unpredictable, and potentially counterproductive.

A bigger danger comes from nihilistic individuals, cults, and others who want the end of the world or mass human casualties rather than political change. In the past, cults such as Aum Shinrikyo sought biological weapons. The group had numerous scientists and engineers in its ranks, and it did succeed in limited weaponization. The Federal Bureau of Investigation has recently publicized the category of “nihilistic violent extremists,” defined as “individuals who engage in criminal conduct within the United States and abroad, in furtherance of political, social, or religious goals that derive primarily from a hatred of society at large and a desire to bring about its collapse by sowing indiscriminate chaos, destruction, and social instability.” Many of these individuals seek mass casualties and care little about any harms to society, as they themselves lack constituents.²² Small probabilities combined with catastrophic consequences justify sustained vigilance. But analysis should avoid assuming that AI automatically collapses all barriers to biological terrorism. The more plausible near-term risk may involve incremental enablement, such as improving access to technical information, accelerating learning curves, or helping less capable actors navigate specific challenges. Even modest improvements in capability could become dangerous when combined with determined actors.

Iran, State Support, and AI-Enabled Proxy Activity

State support has historically amplified the capabilities of terrorist and militant organizations. AI may strengthen this

dynamic by allowing states to provide deniable forms of technical support, information operations, or cyber assistance. Iran is particularly relevant in this context. Tehran possesses meaningful cyber capabilities, sophisticated intelligence services, and extensive relationships with proxy organizations across the Middle East. Iranian influence operations have attempted to shape perceptions within the United States, while Iranian-linked cyber activity has targeted infrastructure and regional adversaries. Given the current U.S. and Israeli conflict with Iran, the country has many incentives to consider terrorist attacks. Indeed, the United States has arrested a senior operative in Kataib Hezbollah, an Iraq-based group with close ties to Iran, claiming that he was planning a series of attacks on Jewish sites in the United States and was linked to a series of attacks on Jewish facilities in Europe.²³

As commercial AI systems become more powerful, states such as Iran may exploit them to support proxy warfare. AI could help identify both physical and cyber vulnerabilities; refine disinformation, deception, and propaganda campaigns; analyze open-source intelligence; or coordinate influence operations. U.S. allies may prove especially vulnerable if they possess weaker cyber defenses or more polarized information environments. At the same time, escalation concerns impose meaningful constraints. State sponsors generally prefer deniable, calibrated actions that avoid direct confrontation with stronger adversaries.²⁴ In addition, providing advanced AI systems to terrorist organizations might result in the groups developing capabilities that state actors do not anticipate or wish to avoid. AI-enabled proxy activity is therefore likely to focus on influence, cyber disruption, propaganda, and targeted support rather than indiscriminate catastrophic attacks.

AI and Counterterrorism

The same technologies that empower extremists will also strengthen counterterrorism capabilities. Historically, terrorism and counterterrorism evolve interactively. The relationship resembles a continual adaptation cycle in which states and non-state actors respond to each other's innovations. For example, in response to the Provisional Irish Republican Army's use of IEDs in the 1970s and 1980s, the British developed ways to inspect and disarm bombs remotely. AI is likely to intensify similar competitive processes. Intelligence agencies may use machine learning tools for several purposes, both strategic and tactical, such as in data analysis and pattern recognition. AI-based techniques can be used to better predict the timing, location, and types of targets of terrorist groups.²⁵ On the tactical side, AI can be used to identify suspicious purchasing behavior, detect anomalous travel patterns, analyze financial transactions, attribute terrorist events, and connect fragmented datasets.²⁶ AI-based social media analysis also offers substantial promise.²⁷ Terrorists frequently reveal information



online, intentionally or unintentionally. AI tools can process vast quantities of digital communications, images, videos, and metadata far more efficiently than human analysts alone. AI-based tools can be used to identify posts that support or glorify terrorist acts.²⁸ They can also be used to flag posts seeking to recruit members to a group and to identify radicalized individuals.²⁹ Finally, organizations such as the Global Internet Forum to Counter Terrorism (GIFCT) coordinate efforts across multiple technology companies, universities, and research organizations to identify posts that violate platform guidelines on terrorism.³⁰ Yet such AI techniques must be used with caution, as they may yield false positives, leading to the possibility that an innocent individual might end up being suspected of terrorist sympathies.

AI also offers unparalleled opportunities for advanced monitoring of terrorists and suspected terrorists. Today, AI systems can “listen” to phone calls, automatically generate call transcripts using voice-to-text models, translate those transcripts from a native language to another, and analyze them to come up with a threat score.³¹ Location tracking and biometric analysis may also become more sophisticated. Today, AI systems can track movements of vehicles in real time using drone and satellite imagery as well as computer vision algorithms.³² When combined with expanding surveillance infrastructure, AI could significantly improve governments’ ability to identify, monitor, and disrupt extremist networks using a combination of network analysis and machine learning.³³ Yet these capabilities raise profound ethical and political concerns. Predictive systems operate probabilistically rather than deterministically. A high probability score does not mean an individual will commit violence. Terrorism remains statistically rare and behaviorally difficult to predict. This creates serious risks of overreach, bias, and false positives. Existing profiling systems already struggle with accuracy. Poorly designed AI systems may amplify discriminatory practices or encourage excessive surveillance, although the risks differ across contexts. Governments may apply lower evidentiary thresholds to non-citizens entering a country than to domestic criminal investigations involving citizens. Overall, democratic societies face difficult questions regarding due process, privacy, and accountability. There is also a broader danger that government officials will trust AI systems excessively. Automation bias may encourage officials to defer to algorithmic outputs even when evidence is weak or ambiguous. Such dynamics could undermine civil liberties while failing to improve security meaningfully. Indeed, one of the greatest long-term dangers may be the misuse of AI in the name of counterterrorism. Expansive surveillance architectures justified by security concerns could evolve into broader systems of political monitoring and social control. The possibility of an AI-enabled surveillance state is therefore a central democratic concern.

Policy Recommendations

Democratic societies face a difficult balancing act. They must preserve technological competitiveness in the field of AI while also reducing the risk of catastrophic misuse. Excessively restrictive policies may undermine innovation and strategic advantage, particularly relative to authoritarian competitors. Yet insufficient oversight could create dangerous vulnerabilities. Given this landscape, several policy priorities stand out:

1. In the case of AI-enabled counterterrorism systems, **governments should require independent auditing of the claims and risk mitigation procedures** made by the companies involved. Such audits should therefore involve independent expert bodies operating under governmental oversight rather than those that rely solely on voluntary corporate assurances.
2. **Public-private coordination must improve substantially.** Much of the relevant technological infrastructure and expertise are in the hands of private firms, while governments possess intelligence authorities, law enforcement capabilities, and national security responsibilities. Effective counterterrorism therefore requires institutionalized mechanisms for information sharing, threat assessment, and crisis coordination between AI companies, government entities, and other bodies.
3. **Chemical and biological risk assessment deserves sustained attention.** Generative AI systems have the potential to significantly reduce the time needed to design deadly chemical and biological agents. Governments should support rigorous testing regimes, red-team evaluations, and international scientific coordination regarding AI-enabled chemical and biological capabilities. Overconfidence in existing safeguards such as the UN Chemical Weapons Convention and the UN Biological Weapons Convention could prove dangerous. Companies involved in the development of AI models should have some obligation to place guardrails on the creation of dangerous biological and chemical agents, with an independent entity carrying out tests and audits. Measures should exist not to penalize the companies, but to work with them to improve such guardrails.
4. Policymakers must maintain robust protections for legal and civil liberties. **Counterterrorism systems built around AI-driven surveillance and predictive analytics require transparent oversight, judicial review, and accountability mechanisms.** Democracies cannot preserve security by abandoning the constitutional principles they seek to defend.
5. **Analysts should avoid technological determinism.** Extremist violence emerges from broader contexts involving identity, grievance, organizational dynamics, conflict, governance failures,



polarization, and strategic calculation. AI does not replace these underlying dynamics.

Artificial intelligence is unlikely to transform terrorism overnight. The most significant near-term effects will likely occur not through autonomous superweapons or apocalyptic scenarios, but through small, cumulative changes in recruitment, disinformation, fraud, operational assistance, and social reinforcement. AI may help smaller actors perform tasks that once required larger organizations and greater expertise. At the same time, AI will strengthen state counterterrorism capabilities through improved surveillance,

pattern recognition, and intelligence analysis. Ultimately, the challenge is therefore political and institutional as much as technological. Democratic societies must pursue strategies that preserve innovation while establishing credible safeguards against misuse. The goal should not be to prevent all technological adaptation by extremist actors—an impossible task—but rather to reduce catastrophic risks, strengthen societal resilience, and ensure that responses to terrorism do not undermine the democratic values they are intended to protect.

Daniel Byman is the director of the Warfare, Irregular Threats, and Terrorism (WITT) Program at the Center for Strategic and International Studies (CSIS) in Washington, D.C.

V. S. Subrahmanian is the Walter P. Murphy Professor of Computer Science and a Buffett faculty fellow in the Buffett Institute of Global Affairs at Northwestern University.

ChatGPT owner says AI acted on its own to hack another tech firm

Source: <https://www.washingtonpost.com/technology/2026/07/21/openais-latest-ai-agent-escaped-security-controls-hacked-tech-company/>

July 21 — Artificial intelligence software in testing by ChatGPT-maker OpenAI breached security controls, accessed the internet and hacked another tech firm to obtain answers to questions probing its cybersecurity skills, OpenAI said on Tuesday.

The security breach follows a [scramble](#) over recent months in the tech industry and governments around the world to respond to the arrival of AI models capable of identifying security flaws in software. The Trump administration temporarily imposed restrictions on OpenAI and its rival Anthropic, maker of the chatbot Claude, to prevent their tech being used by U.S. adversaries.

“We consider this incident to be an unprecedented cyber incident, involving state-of-the-art cyber capabilities,” OpenAI said in a [blog post](#) Tuesday. The company is working with Hugging Face, the AI company that OpenAI’s system attacked, to determine the full impact of the hack.

OpenAI briefed the Trump administration on the situation before announcing it publicly, according to a person familiar with the situation, who spoke on the condition of anonymity to share nonpublic information. Clement Delangue, chief

executive of Hugging Face, said in a post on X on Tuesday that his company had worked closely with OpenAI to understand the incident. “We strongly believe there was no malicious intent on their part. It’s quite mind-blowing that all of this happened autonomously!” he wrote.

Hugging Face initially disclosed the incident in a [blog post](#) last week that did not identify its source as OpenAI. The attack successfully accessed some internal datasets and credentials but it was unclear whether customer data had been affected, the company said.

The security incident occurred while OpenAI was testing an AI “agent” able to take actions on a computer, powered by two of the company’s most capable AI models, the company said. The process involved challenging the AI model to find previously known software vulnerabilities, using a [test](#) designed by computer

security experts.

Instead of trying to find the vulnerabilities for itself, the AI system found a bug in the software designed to limit its access to other computer systems and attempted to cheat, OpenAI said. It exploited the flaw to access the internet and try to obtain




answers to the test questions from Hugging Face, which maintains repositories of AI software, the ChatGPT developer said. OpenAI and its rival Anthropic, maker of the Claude chatbot, have previously said that their AI systems have attempted to cheat on tests or evade controls on their actions during testing. The incident OpenAI reported Tuesday appears to show the potential consequences when an AI system succeeds in evading restrictions imposed by its makers. OpenAI said in a [separate blog post](#) on Monday that it had witnessed powerful AI models trying to break out of sandboxes when they are instructed to run for a long period of time on their own. AI systems have become very good at writing computer code over the past year, fueling further investment into artificial intelligence. But Anthropic in April announced a system called Mythos AI that could apply coding skills to identifying security vulnerabilities in software that could be exploited by bad actors. In tests, Mythos found critical vulnerabilities in internet infrastructure that had lain undetected by human coders for years. The prospect of AI-powered hacking campaigns triggered widespread concern among senior tech, banking and

government officials. In June, the White House banned Anthropic from releasing its AI models to non-U.S. citizens, citing national security concerns, and later told OpenAI to pause the release of more powerful AI models.

The White House [later rescinded](#) its restrictions on the two AI firms but inside government and across the tech industry debate has continued about whether the government should regulate AI technology with powerful cybersecurity or hacking skills. Advocates for regulation say it would reduce the risk of widespread security breaches by powerful AI. Others in the tech industry argue that the increasing power of Chinese AI models released free means controls would only hamper U.S. firms.

Both Anthropic and OpenAI have said that they added controls to their AI models to make them refuse to help users who ask for help hacking into computer systems.

Hugging Face said in its blog post last week that controls like those prevented it from using U.S. AI models to investigate the AI-powered breach of its systems. Instead the company used a Chinese AI model to run the analysis, the company said.



International
CBRNE
INSTITUTE




C²BRNE
DIARY



Preparedness &

EMERGENCY RESPONSE



Go-anywhere survival pod fends off bullets, bombings and worse

Source: <https://newatlas.com/outdoors/portable-bunkers-momentum-lifepods/>



The LifePods B-01 is a protective capsule designed for immediate shelter on land | Momentum Technologies

July 06 – A new class of survival capsule wants to turn disaster-proofing into something you can order rather than build. French startup Momentum Technologies showed off its LifePods at VivaTech and Eurosatory 2026 recently, deployable shelters designed to keep people alive when the usual safety nets – power, shelter, emergency services – stop working.

Momentum's pitch lands amid rising public anxiety over blackouts, extreme weather, and geopolitical instability, concerns that have pushed a wave of consumers and governments alike to look for faster, more mobile ways to prepare for worst-case scenarios. It's the same anxiety driving interest in [backup generators](#) and [emergency toolkits](#), just scaled up into something closer to a portable fortress.

The LifePods W-01 survival capsule on display during the VivaTech 2026 technology trade show at Paris Expo Porte de Versailles | Momentum Technologies



Unlike traditional bunkers poured into the ground, LifePods are portable. They can be shipped in standard containers, stacked for storage, and in some cases even airlifted by helicopter. Two distinct flavors are available: the B-01 is a two-person land-based pod built to resist bullets, blasts, and fire; and the W-01 can hold up to four adults plus four children (kids would sit on adults' laps), and is a floating version designed for floods, tsunamis, and marine immersion.

The B-01 is constructed from multiple protective layers using high-strength technical steel and specialized insulation, essentially combining armor plating, climate shelter, and survival unit into one hull. The company says ballistic panel tests have already met VPAM PM7, an international standard requiring resistance to a 5.56x45-mm round traveling at 950 m/s (3,117 ft/s) and a 7.62x51-mm round at 830 m/s (2,723 ft/s). Full-capsule validation tests are expected in the next couple of months, so the bullet-stopping claims are currently for the material, not yet the finished LifePods.

The W-01 takes a simpler approach. Rather than relying on active propulsion or complex mechanical systems, it uses passive hydrodynamic stability, essentially trusting good buoyant design over motors or moving parts. Momentum Technologies says early flotation tests have been "very encouraging," though it hasn't released independent data to confirm that.

"The market is now confirming the strategic interest in our next-generation resilience solutions," says Cédric Choffat, CEO of Momentum Technologies. "Our objective is to accelerate industrialization, certifications, the structuring of production, and the commercial deployment of our capsules – turning this international visibility into concrete operational deployments." Momentum says it has closed early funding rounds and is preparing a Seed/Series A round between late 2026 and early 2027.



The interior of the LifePods W-01 survival capsule fits four people and four children | Momentum Technologies

LifePods enter a market already testing the waters, literally. [Survival Capsule](#), a patented spherical shelter aimed at tsunamis, hurricanes, and earthquakes, sells two-person units starting around US\$21,700. On land, portable and [modular bunkers](#) – simpler, faster to install but without the multilayer ambitions of Momentum's design – typically run in the \$15,000 to \$40,000 range.

Momentum has published indicative pricing for both models. The W-01 is listed at €35,000 to €40,000 (\$37,800 to \$43,200) ex-works, meaning the price covers the capsule alone, with delivery and installation costing extra. While the B-01 carries a recommended public price in France of €29,000 (\$31,300) including VAT, but excluding delivery and installation. Even with the prices out, everything about LifePods positioning – the B-01's focus on critical infrastructure and security forces, the W-01's family-oriented flood protection – points toward institutional and infrastructure buyers rather than individual consumers.

The real test for Momentum Technologies isn't the trade show spotlight, it's whether the B-01 and W-01 can clear full certification and prove they work as complete systems, not just impressive panels and hopeful buoyancy





tests. If they do, the company could carve out a genuine niche in an increasingly disaster-conscious market. If not, they'll join a long list of expo showpieces that promised resilience and delivered little beyond good lighting.

UAE: National Community Readiness Programme launched to train one million community responders

Source: <https://www.bignewsnetwork.com/news/279180680/national-community-readiness-programme-launched-to-train-one-million-community-responders>



July 11 – The Integrated National Health Emergency Preparedness and Response Programme (Jaheziya) has launched the unified and integrated national programme for community preparedness and response, "Community Readiness", with the aim of training and qualifying one million community responders from members of the public,

volunteers, students, employees and community teams, enabling them to provide safe and organised initial response during the first critical minutes of emergencies, crises and disasters until the arrival of the competent authorities and frontline response teams.



The "Community Readiness" programme is based on an integrated national framework encompassing institutional readiness, community readiness and international readiness, in partnership with government, private, civil, military, healthcare, security, ambulance and community organisations, as well as public benefit institutions.

The programme is supported by a network of international scientific, training and accreditation partnerships with leading universities, training centres and professional institutions in the United States, Canada, Europe, the United Kingdom, Australia and New Zealand. These partnerships strengthen knowledge transfer, the localisation of specialised programmes, the qualification of national trainers and the implementation of unified standards that ensure training quality. The programme comprises four progressive training levels. The first, Basic Readiness, includes risk awareness, personal safety, family emergency planning, requesting assistance, first aid, cardiopulmonary resuscitation (CPR), use of automated external defibrillators (AEDs), bleeding control, choking management, safe evacuation, fire safety, emergency kit preparation and psychological first aid.

The second level, Advanced Community Responder, focuses on incident command principles, incident scene management, primary triage, management of multiple casualties, evacuation and transportation, infection prevention and control, support for evacuation centres, crowd management, logistical support, information documentation and participation in field exercises.

The third level is dedicated to qualifying community leaders in team leadership, emergency planning, risk assessment, resource and volunteer management, decision-making, strategic communication, exercise design, report preparation and capturing lessons learned.

The fourth level is dedicated to preparing national trainers in teaching and training skills, simulation management, scenario development, trainee assessment, quality assurance, content standardisation and the development of national training programmes.

Dr. Adel Abdullah Al Shammari Al Ajmi, Chief Executive Officer of the Zayed Giving Initiative and Chairman of the UAE National "Jaheziya" Programmes, said that "Community Readiness" reinforces a comprehensive national culture based on prevention, preparedness, readiness, rapid response, business continuity and recovery, while transforming training and knowledge into community and institutional capabilities that can be activated whenever needed.

Professor Roberto Mugavero, President of the European Centre for Disaster Medicine, said the programme represents an innovative Emirati model that can benefit regional and international communities through its progressive training levels, digital platform, national responder database, training quality assurance system and readiness assessment framework, while ensuring that trainees do not exceed the limits of their competencies or replace licensed and qualified professionals.



A large, modern building with a prominent red upper section and a grey lower section. The red section has large windows and a balcony. The grey section has a series of small, rectangular windows. Three hazard symbols (radiation, biohazard, and chemical) are visible on the grey section. The building is surrounded by a paved area and some greenery.

ICI
International
CBRNE
INSTITUTE

**A common roof
for International
CBRNE
First Responders**



Join us!

A person wearing a full-body green CBRNE protective suit, including a hood and gloves, stands in the foreground. In the background is a large, modern building with a grey facade and large windows. A Belgian flag is flying on a pole in front of the building. The scene is set in a parking lot.

Rue de la Vacherie, 78
B5060 SAMBREVILLE
(Auvelais)
BELGIUM

info@ici-belgium.be | www.ici-belgium.be