CSIS BRIEFS

Artificial Intelligence and the Future of Terrorism

By Daniel Byman and V.S. Subrahmanian

JULY 2026

THE ISSUE

This brief examines how AI may affect terrorism in the coming years. It argues that AI is likely to reinforce existing trends in decentralization, lone-actor violence, digital propaganda, and information warfare rather than create entirely new forms of terrorism. It also contends that the same technologies that empower extremists will simultaneously strengthen state counterterrorism capabilities, though often at substantial risk to privacy and civil liberties. The brief concludes by arguing that democratic societies must develop institutional safeguards, appropriate regulatory oversight, and public-private coordination mechanisms. At the same time, AI's rapidly expanding capabilities require continued monitoring and regular reappraisals.

INTRODUCTION

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.

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.

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.

Although AI improves the efficiency of propaganda production and may even aid in refinement of messaging, it does not guarantee persuasive success.

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.

AI may reinforce existing grievances and ideological trajectories, but it will still operate within broader social, political, and psychological contexts.

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.

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 intel-

ligence 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.

The same technologies that empower extremists will also strengthen counterterrorism capabilities.

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.

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.*

This report is made possible by general support to CSIS. No direct sponsorship contributed to this report.

CSIS BRIEFS are produced by the Center for Strategic and International Studies (CSIS), a private, tax-exempt institution focusing on international public policy issues. Its research is nonpartisan and nonproprietary. CSIS does not take specific policy positions. Accordingly, all views, positions, and conclusions expressed in this publication should be understood to be solely those of the author(s). *© 2026 by the Center for Strategic and International Studies. All rights reserved.*

Cover Photo: NurPhoto/Getty Images

ENDNOTES

1. For a useful review, see Tyler Houser and Beidi Dong, "The Convergence of Artificial Intelligence and Terrorism: A Systematic Review of the Literature," *Studies in Conflict and Terrorism* 48, no. 7 (July 2025): 1-24, https://doi.org/10.1080/1057610X.2025.2527608.
2. Mahmut Cengiz, "A Comparative Analysis of the Impact and Use of Improved Explosive Devices (IEDs) by Terrorist Groups," Small Wars Journal, April 17, 2025, https://smallwarsjournal.com/2025/04/17/a-comparative-analysis/; and Mahmut Cengiz and Joseph Turgay Karagoz, "Strategic Lethality: Trends in the Use of Explosives by Terrorist Organizations," Homeland Security Today, May 5, 2025, https://www.hstoday.us/featured/strategic-lethality-trends-in-the-use-of-explosives-by-terrorist-organizations/.
3. Andrew Glizzard, David McIlhatton, and Paul Martin, "Will Generative AI Fundamentally Change Terrorist Threats?," *CTC Sentinel* 19, no. 5 (May 2026): 25, https://ctc.westpoint.edu/will-generative-ai-fundamentally-change-terrorist-threats/.
4. Paola Cillo and Gaia Rubera, "Generative AI in Innovation and Marketing Processes: A Roadmap of Research Opportunities," *Journal of the Academy of Marketing Science* 53, no. 3 (2025): 684-701, https://link.springer.com/article/10.1007/s11747-024-01044-7.
5. Chongyang Gao et al., "GEM: Generating Engaging Multimodal Content" in *Proceedings of the Thirty-Third Joint International Conference on Artificial Intelligence* (Jeju: International Joint Conferences on Artificial Intelligence, 2024): 7654-7662, https://www.ijcai.org/proceedings/2024/847.
6. "From Deepfakes to Chatbots, Terrorists Embrace AI to Spread Propaganda," *Africa Defense Forum Magazine*, September 9, 2025, https://adf-magazine.com/2025/09/from-deepfakes-to-chatbots-terrorists-embrace-ai-to-spread-propaganda/.
7. Fabrizio Minniti, "Automated Recruitment: Artificial Intelligence, ISKP, and Extremist Radicalisation," Global Network on Extremism and Technology, April 11, 2025, https://gnet-research.org/2025/04/11/automated-recruitment-artificial-intelligence-iskp-and-extremist-radicalisation/.
8. "Our Approach to Dangerous Organizations and Individuals," Meta Transparency Center, November 11, 2024, https://transparency.meta.com/features/dangerous-orgs-and-individuals/.
9. Chava Gourarie, "A closer look at the man wrongfully accused of being the Boston bomber," *Columbia Journalism Review*, October 9, 2015, https://www.cjr.org/analysis/sunil_tripathi_was_already_dead.php.
10. Aldohn Domingo, "IS Terrorist Group Leverages AI Deepfakes for Propaganda," Tech Times, May 23, 2024, https://www.techtimes.com/articles/304954/20240523/terrorist-group-leverages-ai-deepfakes-propaganda.htm.
11. Souroush Vosoughi, Deb Roy, and Sinan Aral, "The spread of true and false news online," *Science* 359 (2018): 1146-1151, https://www.science.org/doi/10.1126/science.aap9559.
12. Federal Bureau of Investigation, "Criminals Use Generative Artificial Intelligence to Facilitate Financial Fraud," public service announcement, December 3, 2024, https://www.ic3.gov/PSA/2024/PSA241203.
13. Priyank Mathur, Clara Broekaert, and Colin P. Clarke, "The Radicalization (and Counter-radicalization) Potential of Artificial Intelligence," International Centre for Counter-Terrorism, May 1, 2024, https://icct.nl/publication/radicalization-and-counter-radicalization-potential-artificial-intelligence.
14. "A New Industry of AI Companions Is Emerging," *The Economist*, November 6, 2025, https://www.economist.com/international/2025/11/06/a-new-industry-of-ai-companions-is-emerging.
15. Yasir Atalan, "What to Know about Chinese AI Models," CSIS, *Critical Questions*, July 2, 2026, https://www.csis.org/analysis/what-know-about-chinese-ai-models.
16. "Delhi Blast Accused Used ChatGPT, YouTube to Make Bombs: Anti-Terror Agency," Press Trust India, May 24, 2026, https://www.ndtv.com/india-news/delhi-red-fort-blast-news-ai-platform-misused-for-terror-engineering-in-red-fort-blast-anti-terror-agency-nia-11540357.
17. Michael Kenney, "'Dumb' yet deadly: local knowledge and poor tradecraft among Islamist militants in Britain and Spain," *Studies in Conflict & Terrorism* 33, no. 10 (2010): 911-932, https://doi.org/10.1080/1057610X.2010.508508.
18. James Patterson and Lydia Khalil, "The Ungoverned Sky: Drones and the Domestic Extremism Threat," Lowy Institute, *Policy Briefs*, March 22, 2026, https://www.lowyinstitute.org/publications/the-ungoverned-sky-drones-and-the-domestic-extremist-threat; and Yaakov Lapin, "Special Report: Hezbollah FPV Explosive Drone Threat," Alma Research and Education Center, May 11, 2026, https://israel-alma.org/special-report-hezbollahs-fpv-explosive-drone-threat/.
19. "Ukraine using AI drones to strike vital convoys supplying Russian troops," BBC, May 29, 2026, https://www.bbc.com/news/articles/cdjp0n7rn41o.
20. John Jumper et al., "Highly Accurate Protein Structure Prediction with AlphaFold," *Nature* 596 (2021): 583-89, https://www.nature.com/articles/s41586-021-03819-2.
21. Andrea Hunklinger and Noelia Ferruz, "Towards the explainability of protein language models," *Nature Machine Intelligence* 8, no. 5 (May 2026): 649-662, https://www.nature.com/articles/s42256-026-01232-w.
22. See the definition of nihilistic violent extremism in United States' Sentencing Memorandum, United States v. Jack William Rocker, Case No. 8:24-cr-496-VMC-AEP, ECF No. 33, (M.D. Fla. March 12, 2025), https://storage.courtlistener.com/recap/gov.uscourts.flmd.434650/gov.uscourts.flmd.434650.33.0.pdf.
23. U.S. Department of Justice, "Iraqi National Arrested and Charged with Providing Material Support to Iranian-Backed Terrorist Organizations and Directing Attacks Targeting U.S. Citizens and Interests," press release, May 15, 2026, https://www.justice.gov/

opa/pr/iraqi-national-arrested-and-charged-providing-material-support-iranian-backed-terrorist.

24. Daniel Byman, *Deadly Connections: States that Sponsor Terrorism* (Cambridge, UK: Cambridge University Press, 2005).

25. V.S. Subrahmanian et al., *A Machine Learning Based Model of Boko Haram* (Princeton, NJ: Springer 2021): 1-116; Laura Mostert et al., *Machine Learning Techniques to Predict Terrorist Attacks Exemplified by Jama'at Nasr al-Islam* (Princeton, NJ: Springer 2025); and United Nations Office of Counter-Terrorism and United Nations Interregional Crime and Justice Research Institute, *Countering Terrorism Online with Artificial Intelligence - An Overview for Law Enforcement and Counter-Terrorism Agencies in South Asia and South-east Asia* (New York: United Nations, 2021), https://unicri.org/Publications/Countering-Terrorism-Online-with-Artificial-Intelligence-%20SouthAsia-South-EastAsia; Raj Bridgelall, "Identifying factors associated with terrorist attack locations by data mining and machine learning," *International Social Science Journal* 73, no. 248 (2023): 539-557, https://doi.org/10.1111/issj.12414; and Lirika Sola, Youdinghuan Chen, and V.S. Subrahmanian, "A Quantitative Geospatial Analysis of the Risk that Boko Haram Will Target a School," *PLOS One* 20, no. 6 (June 2025): e0320939, https://doi.org/10.1371/journal.pone.0320939.

26. Stanley Munamato Mbiva and Fabio Mathias Correa, "Machine Learning to Enhance the Detection of Terrorist Financing and Suspicious Transactions in Migrant Remittances," *Journal of Risk and Financial Management* 17, no. 5 (2024): 181, https://doi.org/10.3 390/jrfm17050181; Prashis Raghuwanshi, "Deep Learning Model for Detecting Terror Financing Patterns in Financial Transactions," *Journal of Knowledge Learning and Science Technology* 3, no. 3 (2024): 288-296, https://doi.org/10.60087/jklst.vol3.n3.p.288-296; Lijun Sun et al., "Routine Pattern Discovery and Anomaly Detection in Individual Travel Behavior," *Networks and Spatial Economics* 23, no. 2 (June 2021): 407-428, https://link.springer.com/article/10.1007/s11067-021-09542-9; and Gideon Menwa et al., "Deciphering terrorist attributions: A global terrorist data-driven comparison of machine learning models," *Social Network Analysis and Mining* 15 (2025): 99, https://link.springer.com/article/10.1007/s13278-025-01498-9.

27. Steve Ressler, "Social Network Analysis as an Approach to Combat Terrorism," *Homeland Security Affairs* 2 (2006): 8, https://www.hsaj.org/articles/171.

28. Mohammad Fraiwan, "Identification of markers and artificial intelligence-based classification of radical Twitter data," *Applied Computing and Informatics* 21, no. 3-4: 304-316, https://doi.org/10.1108/ACI-12-2021-0326.

29. Ahmed Hassan, "AI and Public Safety: Detecting Online Radicalization," 2026, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6642198.

30. "Global Internet Forum to Counter Terrorism," Global Internet Forum to Counter Terrorism, https://gifct.org/.

31. Ariana Genovese et al., "Artificial intelligence in clinical settings: A systematic review of its role in language translation and interpretation," *Annals of Translational Medicine* 12, no. 6 (December 2024): 117, https://atm.amegroups.org/article/view/132363/html; Anand Kumar Singh et al., "Generative AI in Voice Transcription: A Comprehensive Survey" in *13th International Conference on Intelligent Embedded, MicroElectronics, Communication and Optical Surveys* (Jaipur: IEEE 2025): 1-6, https://doi.org/10.1016/j.procs.2025.04.628; and Saydul Akbar Murad et al., "Multilingual Cyber Threat Detection in Tweets/X Using ML, DL, and LLM: A Comparative Analysis," *IEEE Transactions on Computational Social Systems* 13, no. 3 (June 2026): 1758-1772, https://doi.org/10.48550/arXiv.2502.04346.

32. Chenjie Zhao et al., "Multi-Sensor Fusion Driven Surface Vessel Identification and Tracking Using Unmanned Aerial Vehicles for Maritime Surveillance," *IEEE Transactions on Consumer Electronics* 71, no. 2 (May 2025): 6280-6293, https://ieeexplore.ieee.org/document/10833660; and Youseef Alotaibi et al., "Optimal deep learning based vehicle detection and classification using chaotic equilibrium optimization algorithm in remote sensing imagery," *Scientific Reports* 15 (2025): 17921, https://www.nature.com/articles/s41598-025-02491-0.

33. Francesca Spezzano, V.S. Subrahmanian, and Aaron Mannes, "Reshaping Terrorist Networks," *Communications of the ACM* 57, no. 8 (2014): 60-69, https://dl.acm.org/doi/10.1145/2632661.2632664.