Will Generative AI Fundamentally Change Terrorist Threats?

By Andrew Glazzard, David McIlhatton, and Paul Martin

This article examines whether generative artificial intelligence (Gen AI) represents a genuine transformation in terrorist capability or whether its risks are being overstated within contemporary security discourse. Drawing on affordance theory, it explores how terrorists exploit technologies in ways not intended by their designers and assesses the practical implications of Gen AI across propaganda, recruitment and radicalization, reconnaissance, and attack planning. The article argues that while Gen AI tools, particularly large language models, can improve the efficiency, accessibility, and scale of certain terrorist activities, there remains limited evidence that they fundamentally alter the nature of terrorism or significantly enhance operational capability. Instead, many anticipated threats remain hypothetical and are often shaped by speculative worst-case scenarios rather than demonstrated use. The article further argues that terrorists historically favor technologies that are cheap, accessible, reliable, and easy to adapt, rather than highly sophisticated innovations.

Technology is, for obvious reasons, a major focus for counterterrorism analysts. Modern terrorism originated in the second half of the 19th century partly from the convergence of developments in weaponry, such as the invention of dynamite (1867), and in communication technologies, such as electric telegraphy (patented 1837; first deployed 1844), steam-powered rotary printing presses (1843-47) and Linotype (1866).[1] The vast expansion of commercial aviation in the 1960s gave terrorists something new and spectacular to attack, while the arrival of television in people's homes at around the same time provided them with the means to provide the spectacle.[2] Digital technologies from the 1990s enabled terrorist networks to supersede closely knit organizations and achieve global reach; social media and smartphones allowed terrorist brands to exist virtually as well as on the ground.[3] As various forms of artificial intelligence are now making real and, in some cases, profound changes to our lives, the public realm is becoming filled with discussion of their potential benefits and risks. It is therefore only to be expected that terrorism analysts have turned toward the potential downsides of these remarkable technologies, and of generative AI (Gen AI) in particular. This article considers the terrorist risk presented by widespread, freely available Gen AI tools, and specifically whether Gen AI offers, as some fear, a step-change in terrorist capability.

Andrew Glazzard is Professor of National Security Policy and Practice in the Protective Security Lab at Coventry University.

David McIlhatton is Associate Pro Vice Chancellor for Defence and National Security, and Director of the Protective Security Lab at Coventry University

Paul Martin is Professor of Practice at the Protective Security Lab at Coventry University and a Distinguished Fellow of RUSI.

The Affordances of Gen AI for Terrorist Groups and Movements

Terrorism scholars have productively deployed the theory of affordance to explain how terrorists exploit technology. Affordance (a concept developed by the psychologist James J. Gibson) is what an environment provides that an animal or individual can use: A tree, for example, may be useful for escaping from predators, as a source of food, or as a vantage point for surveying a landscape.[4] Applied to technology, affordances are what a tool enables an individual or group to do; affordance theorists make a subtle but important distinction between a tool's features, which are the elements of the tool itself, and affordances, which are the actions that can be performed with the tool.[5] In this context, there is a subtle difference between a tool's purpose and how it is used: Affordances are not necessarily intended by the technology's designers.[6] This apparently simple observation is at the heart of Relational Affordance Theory, which proposes that a tool's affordance does not derive exclusively from its inherent properties but emerges from a dynamic and iterative interaction between the tool and the humans (individuals, groups, societies) that use it.[7] This is important in terrorism studies as most technologies are not designed for the purposes of terrorism: Terrorists discover the properties within technologies they find useful, irrespective of their intended purpose.[8] One example is the Memopark timer, a Swiss-made 60-minute mechanical timing device intended to remind motorists of metered parking expiry times. The Provisional IRA (PIRA) discovered in the 1970s that the Memopark was an extremely cheap and reliable method of controlling the initiation of improvised explosive devices (IEDs), and the group used them frequently and to deadly effect in their time and power units.[9] The inventors of the timer presumably had no idea that their tool could be used for such malevolent purposes.

One important and useful distinction, developed by Donald Norman, is between the *perceived* affordance and the *actual* properties of a technology, giving rise to the concept of hidden affordance (where the affordance is not easily perceived).[10] Another distinction is between functional affordance (what a tool can do in the physical world) and cognitive affordance (what it can do mentally or informationally).[11] And these distinctions raise another

important consideration: A tool may appear to offer functional or cognitive affordance to the user, but this perception may be mistaken, either because the user has misperceived its function, or because the tool is ineffective or has a misleading design feature. This is known as false affordance.[12]

What are the affordances of Gen AI from the perspective of 21st century terrorists? More simply, what can terrorists do with Gen AI that they could not do yesterday? Most of the research and commentary on this has focused, unsurprisingly, on large language models (LLMs), the form of Gen AI that has captured the most attention and is making the most progress into our daily lives. A survey of recent papers suggests at least four areas of concern:

- **Propaganda:** Gen AI can enable terrorist/extremist groups to produce propaganda at scale and speed, to produce new kinds of content (such as highly credible but inauthentic synthetic content or 'deepfakes,' or immersive propaganda using extended reality applications), to target content more efficiently, and to evade existing countermeasures to moderate or remove terrorist content—all with the potential to accelerate and vastly expand their ability to recruit and to radicalize.[13]
- **Recruitment and radicalization:** Gen AI chatbots potentially increase the capacity of terrorists and extremists to engage with potential recruits, or even to guide attacks, without the need for a human in the loop, while self-starting, undirected 'lone actor' terrorists can potentially create their own radicalizers. Counterterrorism analysts (and the United Kingdom's Reviewer of Counter-Terrorism Legislation) have drawn attention to the case of Jaswant Singh Chail, convicted in the United Kingdom for treason (rather than terrorism) after he was arrested at Windsor Castle with a crossbow on Christmas Day 2021; Chail had discussed his plan to assassinate Queen Elizabeth with his Gen AI 'girlfriend,' a chatbot named Sarai.[14]
- **Research and reconnaissance:** Gen AI enables terrorist/extremist groups to conduct research on potential targets and attack methods at a level previously unattainable using, for example, conventional search engines, while 'hostile reconnaissance' (obtaining information on a potential target), which can already be conducted remotely to an extent, could be made more efficient by Gen AI-enabled digital twins (digital representations of real-world artifacts or environments).[15] Indeed, the United Kingdom's National Protective Security Authority has brought forward guidance for government and businesses for mitigating the risks of this.[16]
- **Attack methodology:** Gen AI could theoretically be used by terrorist groups to develop innovative attack methods, overcoming a model's guardrails (designed to prevent malicious misuse) and developing, for example, novel explosives or designing toxins or biological agents for use in weapons,[17] or innovating with delivery methods such as connected and autonomous vehicles. Some forms of Gen AI could be used in the delivery of attacks (e.g., by creating and delivering malicious code into the computers and networks of their targets).[18]

Taken together, this suggests there is cause for worry. As Hauser and Dong note in their systematic review of the literature on terrorism and AI, "it is a matter of *when*, not *if*, terrorist organizations will begin to weaponize a broad range of AI applications and services for malicious purposes."[19] If Gen AI can enhance terrorist capability in at least four ways, we should prepare ourselves for more (and more effective) attacks. And there should be no doubt that terrorists, like anyone with an internet connection, are using Gen AI and increasingly so. However, in assessing whether Gen AI will be a game-changer for the terrorist threat, one needs to look closer at its affordances from a terrorist perspective, and not merely construct hypothetical vignettes of imaginable possibilities. The following section analyzes the affordances of Gen AI for terrorist use.

Capability, Availability, and Effect

Terrorists do not use any and every tool that has an affordance from their perspective. There have been several attempts to identify what makes a technology appealing or useful to terrorists. Baele and Brace, drawing on affordance theory in their study of the value of AI to extremist movements, identify two factors in the actualization of technology affordance: consistency with social norms and values, and strategic considerations.[20] In other words, terrorists or extremists will consider not only the practical utility and costs of a technology in achieving their desired objectives, but also the extent to which the technology aligns with the movement's morality and culture. A somewhat more elaborate framework, using Situational Crime Prevention theory to evaluate the attractiveness of attack methods, was developed by Clarke and Newman in 2007. The components of their MURDEROUS framework are: Multi-purpose, Undetectable, Removable, Destructive, Enjoyable, Reliable, Obtainable, Uncomplicated, and Safe.[21] More recently, Cronin has proposed a similar (and overlapping) set of characteristics for the technologies with the greatest application for violent non-state users: They will be accessible, cheap, simple to use, transportable, concealable, effective, multi-use, mature, commercially available, and customizable.[22] Neither framework has been empirically tested and validated, and several of the components proposed could be challenged. Despite these limitations, the two frameworks are at least a starting point for identifying what might be the more specific affordances of a technology such as Gen AI for a terrorist. Seen through the prism of affordance theory, the 19 components of these frameworks can be reduced to three categories:[a] i) capability, by which we mean the actual properties and perceived affordances of the technology (including but not restricted to the capability in conducting an attack); ii) availability, meaning the ease (or perceived ease) with which the terrorist can exploit those affordances; and iii) effect, i.e., what the technology achieves both functionally (in the physical world) and cognitively—an important distinction as terrorists are in the business of creating or changing perceptions in their target populations.[23]

How does Gen AI measure up against these three dimensions of capability, availability, and effect from the terrorist's perspective? In terms of capability, Gen AI can certainly increase the efficiency of activities, such as propaganda generation or research on targets. For the former, numerous studies have demonstrated that generating extremist text and imagery is straightforward with only basic

a Mapping this taxonomy to Clarke and Newman and Cronin's frameworks, capability comprises Multi-purpose/multi-use, Removable/transportable, Reliable, mature, customizable, concealable; availability comprises Obtainable/commercially available, Undetectable, accessible, cheap, Uncomplicated/simple to use and Safe; and effect comprises Destructive, Enjoyable, and effective.

prompt engineering or fine tuning. As a result of concerns such as these, the developers of LLMs in particular are keen to stress that their models come with 'guardrails,' protections that are aimed to prevent malicious or accidental misuse. However, as with almost any technological protective measure, it is possible to bypass or corrupt the guardrails, using specific prompt engineering (designing tailored inputs into the LLM to create outputs that the designers of the model did not intend or desire) or fine tuning (training the model on specific inputs so it adapts its outputs accordingly). Prompt engineering techniques include 'prompt jailbreak' (where an attacker constructs an input in such a way that it bypasses the restrictions built into the model's algorithms), 'prompt injection' (where the attacker overrides the developer's instructions, either directly using their own instruction or indirectly by injecting malicious code into the model's data sources), and 'prompt leaking' (where an attacker extracts elements of the model's architecture in order to replicate or corrupt it).[24] Experimental studies show that it is indeed possible to generate results useful to terrorists using prompt jailbreaks, though one study intriguingly shows that the success rate was almost as high without having to resort to prompt engineering (i.e., by simply asking the model to complete the task).[25]

It seems clear, then, that Gen AI presents significant capability affordance in propaganda production in terms of efficiency by increasing scale/volume and speed while reducing costs. Set against those efficiency gains are risks identified by Baele and Brace: reduced quality (although they suggest that this risk diminishes as the power of Gen AI increases) and reduced legitimacy and authenticity (which are considered below under Effect).[26] As with the use of any system or service, they may also entail security risks for the terrorist in that their use of Gen AI may provide advance warning or evidence of terrorist activity.

Turning to research and reconnaissance, the efficiency gains of Gen AI for terrorists are equally clear, albeit curiously under-researched.[b] If, for example, one imagines a terrorist group seeking to attack a well-protected site or individual, LLM-powered research has the potential to generate more and better information more quickly than relying on static information sources such as an official website or a Wikipedia entry. Further, one can infer that the benefits of Gen AI for legitimate open-source intelligence (OSINT) research—which have been amply demonstrated in, for example, the rapid processing by investigative journalists using LLMs of 'data dumps' and evidential material—are equally relevant to those with more nefarious objectives.[27] But in this use-case, LLMs are a tool to make better use of the rich information ecosystem already available through the internet and associated capabilities such as remote image capture. Terrorists have been able to conduct virtual reconnaissance from their laptops using existing data and commercial tools for many years, while uncrewed aerial vehicles (drones) offer opportunities to gather new and more granular reconnaissance data of specific locations.[28] Rather than enabling new kinds of planning and research, current and recent case studies show that Gen AI at this time offers efficiency gains rather than novel capabilities, while the seemingly inescapable problem of Gen AI hallucination may also create problems for terrorist users.[29]

For the other use cases identified in the literature, the gains of Gen AI are rather more speculative. Despite a degree of moral panic over internet-enabled radicalization at scale, the process of becoming a terrorist remains what it has always been: a complex, human-centric, socially embedded, variable and contingent phenomenon, in which various technologies may perform enabling roles (e.g., by exposing susceptible individuals to extremist influences) but are rarely causative in and of themselves.[30] At the other end of the scale is the fear that Gen AI tools, especially when combined with algorithmic audience segmentation, can more effectively micro-target radicalizing propaganda at susceptible individuals.[31] However, studies are yet to provide evidence of terrorist groups or individuals using such techniques. 'Radicalizing chatbots' are conceivable, but the literature has yet to show how the messy and contingent process of radicalization can be automated.

Similarly, there is no shortage of scenarios in which Gen AI tools are used to create novel attack methods, but rather less research and evidence showing a viable pathway from an LLM prompt to a previously undiscovered method. Terrorists do experiment, and thereby develop novel approaches, but they rarely do so at the frontiers of knowledge.[32] Rather, they are in the business of discovering hidden affordances in readily available items to people (like the Memopark timer). And for the more concerning applications, such as creating novel chemical, biological, or radiological (CBR) weapons, a great deal of skill and knowledge is still likely to be needed to create viable attack methods. Difficulties in manufacturing and deploying CBR weapons largely explain why we have seen them used so rarely. Even al-Qa`ida, which is known to have invested significant resources in its CBR research programs in the 1990s and early 2000s, and to have developed potentially viable chemical devices, has yet to demonstrate a serious, operational CBR capability.[33]

Gen AI technologies are self-evidently available to terrorists irrespective of ideology and type of organization (lone actor, networked social movement, directed organization). Tools such as LLMs are extremely widespread, with models either available free of cost or for a modest fee, easy to use, often undetectable and safe to use (or perceived to be so), thus satisfying the availability criteria in Clarke and Newman and Cronin's frameworks. Digital technology, however, is unevenly distributed and not quite as ubiquitous as might be assumed in the developed world: Internet penetration in 2025 was a mere 13 percent in Chad and below 20 percent in Mozambique, two countries facing severe terrorist threats.[34]

Finally, there is the question of effect. Terrorists are already using LLMs to produce synthetic propaganda, as well as using LLMs to enable, for example, automated translation and transcription of existing material.[35] However, the real-world effectiveness of these activities is likely to be limited at least in the near term. Scenarios in which synthetic propaganda is created and circulated at scale are plausible. But the question remains whether such operations really would succeed in recruiting, radicalizing, and inspiring future terrorists at significantly greater scale and speed than before these

b Hauser and Dong's systematic review of AI and terrorism identified 28 studies meeting their inclusion criteria, only one of which addressed the benefits to terrorists of increased information availability from AI tools, and this only did so in passing. The study was Ștefan Săvulescu and Lucian Ivan, "The Impact of Chat GPT on Cybercrime and on the Activities Carried out by the Law Enforcement Structures," *Romanian Journal of Forensic Science* 136 (2023).

technologies became available.[c] Terrorist efforts to leverage Gen AI for these purposes have, to date, been "predominantly low-stakes, low-impact" and limited in scale, despite the efforts of some propaganda outlets (such as Voice of Khorasan, the English-language magazine of Islamic State Khorasan) to encourage uptake of AI tools.[36] This may, in part, be the manifestation of two of the risks of terrorists investing in Gen AI identified by Baele and Brace: "legitimacy hazard" (the terrorist group or individual losing legitimacy in their supporters' eyes because of their reliance on Gen AI) and "authenticity depletion" (the material or the producers being seen as inauthentic).[37] [d] These risks are likely to be particularly high for groups and movements, such as jihadis, which prize canonical sources of religious or spiritual authority and are therefore unlikely to regard algorithms as capable of representing the divine. There may be more practical reasons for the limited uptake of Gen AI tools, such as the limited utility of Gen AI tools in solving operational problems or security concerns that may arise from using applications that are potentially vulnerable to surveillance.[e]

Another limitation on the persuasive and motivational effects of Gen AI tools derives from the fundamentals of communication theory. Human beings are not blank slates waiting to be persuaded to join the Islamic State, or a far-right extremist movement, by something they have seen or read on the internet, no matter how technologically advanced in its production. The hypodermic or 'magic bullet' theory of communication has been largely discredited as it fails to account for human frailties such as stubbornness and inattentiveness, and for the complex processes of selection, interpretation, and social mediation that occur when we do actually pay attention to something we read or see.[38] More complex and empirically supported theories of communication emphasize that people are often resistant to persuasion, filter information inputs through pre-existing beliefs and through social relations, and are highly selective in their attention, leading to relatively limited effects from exposure to media messages.[39] This is true even of the most precisely targeted communication, which perhaps explains why fears of social media algorithms enabling mass persuasion through micro-targeting have proved to be drastically overblown.[40] Communication theory also shows that information ecosystems shape the selection of messages, messengers, or channels,[41] and here, the impact of Gen AI may serve to weaken rather than strengthen terrorist influence: Studies speculating about the effects of Gen AI terrorist content consistently fail to address the informational context of the dramatic rise of Gen AI tools. The capability to produce vastly more content is now available to everyone looking to influence their target audiences, including those with counterterrorism objectives. But the important point here is not so much that Gen AI tools are actor-agnostic but that the information environment is changing radically with the exponential increase in machine-generated content. As Herbert Simon observed in 1971, when information is abundant, attention is scarce, and attention scarcity is likely to be as much of a problem for terrorists as for anyone seeking to influence and persuade populations.[42] Some critics of big tech point to the increasing prevalence of 'AI slop'—low-quality AI-generated content—in the information ecosystem.[43] Will adding quantities of terrorist AI slop make much difference to the threat?

Balancing Risks

Examining terrorist use of Gen AI with affordance theory suggests that the risks of terrorist misuse of Gen AI are often overstated, even though it is also clear that Gen AI has the potential to lower some barriers to entry and make some terrorist activities more efficient. But given that the stakes are so high—a terrorist bioweapon, after all, implies some deeply alarming scenarios—are terrorism analysts correct to focus on the risks posed by Gen AI, so that counterterrorism practitioners in government and industry can close them off? Focusing on high-impact, low-probability risks is, after all, a necessary part of the counterterrorism business given the rarity, at least in the West, of significant terrorist attacks.

Any risk is worth examining. But there are serious downsides to focusing on hypothetical catastrophes. The first is that doing so comes with an opportunity cost. Resources for governments and their security agencies are scarce. And attention—that most scarce resource in the Gen AI-enabled information economy—is an essential resource for national security decision-makers. Every pound, dollar, or hour spent on insuring the public against AI-enabled terrorism is a pound, dollar, or hour not spent on other counterterrorism measures. Furthermore, there are risks from Gen AI in other categories of malevolent misuse—such as those involving cyber-attacks by hostile states[44] or organized criminal networks[45]—that are potentially far greater in terms of impact and likelihood than Gen AI-enabled terrorism, and so deserve greater resource and focus, even though the adoption of Gen AI by cybercriminals has not yet led to a significant disruption of the cybercrime ecosystem.[46]

The second is the risk of creating what Kuran and Sunstein call an "availability cascade," which they define as a "self-reinforcing process of collective belief formation by which an expressed perception triggers a chain reaction that gives the perception [of] increasing plausibility through its rising availability in public discourse."[47] When public concern over the terrorist risks of Gen AI reaches a certain level, it may trigger a series of unproductive or counter-productive actions. Indeed, there are warnings from recent history that highlighting the capacity for terrorist misuse of novel technologies, based on speculation rather than evidence, risks incentivizing counter-productive behaviors in decision-makers. The fear that weapons of mass destruction (WMD) could fall into the hands of groups like al-Qa`ida was not the only issue that prompted the invasion of Iraq in 2003, but it was certainly a major plank of the war's justification. Then-British Prime Minister Tony

c Some research suggests that exposure to online environments accelerates the radicalization process. See, for example, Michael Jensen, Patrick A. James, and Herbert Tinsley, "Profiles of Individual Radicalization in the United States - Foreign Fighters (PIRUS-FF): Infographics," Report to the Office of University Programs, Science and Technology Directorate, U.S. Department of Homeland Security, 2016, and "Overview: Profiles of Individual Radicalization in the United States-Foreign Fighters (PIRUS-FF)," National Consortium for the Study of Terrorism and Responses to Terrorism, last accessed May 22, 2026. However, much of this research links speed of radicalization to online environments in general (and predates recent developments in Gen AI), and in the view of the authors, measuring the velocity of radicalization is fraught with problems of construct validity (the ambiguous status of 'radicalization'), data availability (scarcity of longitudinal data and reliance on retrospective data), and temporal delimitation (the start-points and end-points of a radicalization journey are matters of interpretation).

d Going by the authors' definitions, these two risks are, in the view of this current article's authors, hard to differentiate.

e Thomas Hegghammer has observed that jihadis specifically appear not be using Gen AI at present and has attributed this to a lack of operational need for them to do so. (Hegghammer, unpublished working paper, 2026).

"The transformative potential of Gen AI for terrorist threats can be seen as the latest iteration of a tradition of security rhetoric that elevates the novel over the established and the possible over the actual."

Blair repeatedly expressed his fear that the proliferation of WMD, including to non-state terrorist groups, would be the defining security challenge of the 21st century: In a critical debate in the UK Parliament on the eve of the war, he condemned those who "dispute the link between terrorism and weapons of mass destruction, and dispute, in other words, the whole basis of our assertion that the two together constitute a fundamental assault on our way of life."[48] A more sober and realistic assessment of terrorist capability might have induced greater caution. Once the WMD panic had subsided, terrorist use of cyber-weapons came into view, with President Obama in 2009 warning:

> *Al Qaeda and other terrorist groups have spoken of their desire to unleash a cyber attack on our country -- attacks that are harder to detect and harder to defend against. Indeed, in today's world, acts of terror could come not only from a few extremists in suicide vests but from a few key strokes on the computer -- a weapon of mass disruption.*[49]

Despite some false alarms, and a perception among researchers that the threat of cyberterrorism is actually increasing, there has yet to be a single instance of a destructive cyber-attack that can be confidently attributed to a terrorist organization.[50]

The transformative potential of Gen AI for terrorist threats can be seen as the latest iteration of a tradition of security rhetoric that elevates the novel over the established and the possible over the actual. While presidents and prime ministers were warning of terrorist WMD and cyber 'Pearl Harbors,' actual terrorists were indeed busy—discovering the value in the information age of the most basic and mature technologies such as knives and vehicles. The technological developments of greatest importance were less in planning and executing the attacks and more in broadcasting them. Some counterterrorism researchers were indeed prescient in paying attention to developments such as livestreaming on social media: Their work has aged much better than those that warned of the catastrophic potential of cyber weapons.[51]

Terrorism experts tend toward pessimism in their analysis of trends and possible futures. At a speech given when he was head of the United Kingdom's domestic security agency, Jonathan Evans observed that "when intelligence folk smell roses they look for the funeral." However, he went on to express skepticism about the notion that the problem grows inexorably worse:

> *Those of us who are paid to think about the future from a security perspective tend to conclude that future threats are getting more complex, unpredictable and alarming. After a long career in the Security Service, I have concluded that this is rarely in fact the case. The truth is that the future always looks unpredictable and complex because it hasn't happened yet. We don't feel the force of the uncertainties felt by our predecessors. And the process of natural selection has left us, as a species, with a highly developed capacity to identify threats but a less developed one to see opportunity.*[52]

That capacity is particularly highly developed among the cadre of terrorism analysts who have a record of seeing technological change as an enabler of greater terrorist capability. But what the history of terrorism and technology, going back to the 19th century, really teaches us is that terrorists are early adopters of communication technologies—such as newspapers in the 19th century, television in the 20th century, and social media in the 21st. Terrorism is a form of communication so we should not be surprised that terrorists are quick to seize on developments in communication technology. For attack methods, though, terrorist groups are still more likely to use what is at hand, what is cheap, and what is easy—and they will exploit the affordances of available, cheap, and easy-to-use technologies in ingenious and innovative ways. It takes a certain kind of what is sometimes called malevolent creativity to turn a pressure cooker into a weapon or use an acid-filled condom as a timer for an explosive device.[53]

This article is not saying that terrorists will avoid using Gen AI or that technologies like LLMs lack affordances for terrorist use. It does, however, sound a note of caution about how researchers and practitioners talk about the implications of Gen AI for the terrorist threat and draw attention to the consequences of over-hyping the risks. Much is made of failures of imagination in counterterrorism, and the authors agree that imagination is an important skill for anticipating threats. But researchers and practitioners also need to restrain their imaginations by focusing on evidence, rather than entertaining themselves and others with scenarios of terrorist catastrophe. CTC

Citations

1 Audrey Kurth Cronin, *Power to the People: How Open Technological Innovation Is Arming Tomorrow's Terrorists* (Oxford: Oxford University Press, 2020), pp. 61-93, 171-199; Carola Dietze, *The Invention of Terrorism in Europe, Russia, and the United States* (New York: Verso Books, 2021), pp. 595-596.

2 Jason Burke, *The Revolutionists: The Story of the Extremists Who Hijacked the 1970s* (London: Bodley Head, 2025).

3 Alexander Meleagrou-Hitchens, Audrey Alexander, and Nick Kaderbhai, "The impact of digital communications technology on radicalization and recruitment," *International Affairs (Royal Institute of International Affairs 1944-)* 93:5 (2017): pp. 1,233-1,249; Donald Holbrook, "Social media and terrorism" in Diego Muro and Tim Wilson eds., *Contemporary Terrorism Studies* (Oxford: Oxford University Press, 2022), pp. 347-368.

4 James J. Gibson, "The Theory of Affordances" in *The Ecological Approach to Visual Perception* (Boston: Houghton Mifflin, 1979), pp. 56-60.

5 Jeffrey W. Treem and Paul M. Leonardi, "Social media use in organizations: Exploring the affordances of visibility, editability, persistence, and association," *Annals of the International Communication Association* 36:1 (2013): pp. 143-189.

6 Grainne Conole and Martin Dyke, "What are the affordances of information and communication technologies?" ALT-j 12.2 (2004): pp. 113-124.
7 Paul M. Leonardi, "Theoretical foundations for the study of sociomateriality," *Information and Organization* 23:2 (2013): pp. 59-76.
8 For a discussion of the application of the concept in terrorism studies, see Max Taylor, "Terrorism and affordance: An introduction" in Max Taylor, P. M. Currie, John G. Horgan, and Mark Currie eds., *Terrorism and Affordance* (New York: Continuum, 2012), pp. 1-17.
9 Brian A. Jackson, "Provisional Irish Republican Army," *Aptitude for Destruction* 2 (2005): pp. 93-140.
10 Donald A. Norman, *The Psychology of Everyday Things* (New York: Basic Books, 1988).
11 Rex Hartson, "Cognitive, physical, sensory, and functional affordances in interaction design," *Behaviour & Information Technology* 22:5 (2003): pp. 315-338.
12 William W. Gaver, "Technology affordances," Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, 1991.
13 Stephane J. Baele and Lewys Brace, "AI Extremism: Technologies, Tactics, Actors," *Vox-Pol*, 2024; "First Responder's Toolbox: Emerging Technologies May Heighten Terrorist Threats," National Counterterrorism Center/Joint Counterterrorism Assessment Team, October 14, 2022; "Algorithms and Terrorism: The Malicious Use of Artificial Intelligence for Terrorist Purposes," UNCCT/UNICRI, 2021; David Wells, "The Next Paradigm-Shattering Threat?: Right-Sizing The Potential Impacts Of Generative AI on Terrorism," Middle East Institute, 2024; C. Nelu, "Exploitation of generative AI by terrorist groups," International Centre for Counter-Terrorism, 2024; "Early Terrorist Adoption of Generative AI," Tech Against Terrorism, November 8, 2023; Tyler Houser and Beidi Dong, "The Convergence of Artificial Intelligence and Terrorism: A Systematic Review of the Literature," *Studies in Conflict & Terrorism* (2025): pp. 1-24.
14 Jonathan Hall, "The Terrorism Acts in 2023: Report of the Independent Reviewer of Terrorism Legislation on the operation of the Terrorism Acts 2000 and 2006, and the Terrorism Prevention and Investigation Measures Act 2011," July 15, 2025, last accessed November 17, 2025; Priyank Mathur, Clara Broekaert, and Colin P. Clarke, "The Radicalization Potential of Artificial Intelligence," International Centre for Counter-Terrorism, May 1, 2024; Houser and Dong.
15 Houser and Dong.
16 "Security-Minded Communications: Guidance for Virtual Tours," National Protective Security Authority, April 2023.
17 Ewan Callaway, "AI can design viruses, toxins and other bioweapons. How worried should we be?" Nature News Feature, May 13, 2026.
18 Houser and Dong.
19 Ibid.
20 Baele and Brace.
21 Ronald V. Clarke and Graeme R. Newman, "Situational crime prevention and the control of terrorism," *Nato Security Through Science Series E Human and Societal Dynamics* 21 (2007): pp. 285-297.
22 Cronin, pp. 13-14.
23 See, for example, Jonathan Matusitz, *Terrorism and communication: A critical introduction* (Thousand Oaks, CA: Sage Publications, 2012), pp. 77-101.
24 Baha Rababah, Shang Tommy Wu, Matthew Kwiatkowski, Carson K. Leung, and Cuneyt Gurcan Akcora, "Sok: Prompt hacking of large language models" in *2024 IEEE International Conference on Big Data (BigData)*, 32024, pp. 5,392-5,401.
25 Gabriel Weimann, Alexander T. Pack, Rachel Sulciner, Joelle Scheinin, Gal Rapaport, and David Diaz, "Generating terror: The risks of generative AI exploitation," *CTC Sentinel* (2024): pp. 17-24. See also Stephane J. Baele, Lewys Brace, and Elahe Naserian, "More is More," Perspectives on Terrorism 19:1 (2025): pp. 34-63; Kris McGuffie and Alex Newhouse, "The radicalization risks of GPT-3 and advanced neural language models," arXiv preprint arXiv:2009.06807 (2020).
26 Identified by Baele and Brace in "AI Extremism."
27 See, for example, J. Černý, "Prompt Engineering: Tactics and Techniques in Open-Source Intelligence," *Journal of Information Warfare* 23:3 (2024): pp. 115-135; Besjon Cifliku and Hendrik Heuer, "'This could save us months of work'-Use Cases of AI and Automation Support in Investigative Journalism," Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems, 2025; T.O. Browne, M. Abedin, and M.J.M. Chowdhury, "A systematic review on research utilising artificial intelligence for open source intelligence (OSINT) applications," *Int. J. Inf. Secur.* 23 (2024): pp. 2,911-2,938.
28 Emil Archambault and Yannick Veilleux-Lepage, "Drone imagery in Islamic State propaganda: flying like a state," *International Affairs* 96:4 (2020): pp. 955-973.
29 For the persistence of hallucinations, see, for example, S. M. T. I. Tonmoy et al., "A comprehensive survey of hallucination mitigation techniques in large language models," arXiv preprint arXiv:2401.01313 6 (2024).
30 Alexander Meleagrou-Hitchens and Nick Kaderbhai, "Research Perspectives on Online Radicalization: A Literature Review, 2006-2016," International Center for the Study of Radicalization (ICSR), *Vox-Pol*, 2017; Chamin Herath and Joe Whittaker, "Online radicalisation: Moving beyond a simple dichotomy," *Terrorism and Political Violence* 35:5 (2023): pp. 1,027-1,048; Joe Whittaker, "Rethinking online radicalization," *Perspectives on Terrorism* 16:4 (2022): pp. 27-40; Sandy Schumann, Jonathan Kenyon, and Jens Binder, "Identifying distinct types of internet use that predict the likelihood of planning or committing a terrorist attack: Findings from an analysis of individuals convicted on terrorism (-related) charges in England and Wales," *Computers in Human Behavior* 168 (2025); Paul Gill et al., "Terrorist use of the Internet by the numbers: Quantifying behaviors, patterns, and processes," *Criminology & Public Policy* 16:1 (2017): 99-117.
31 Nasir Ahmad Ganaie, "The Role of Artificial Intelligence in Radicalization, Recruitment and Terrorist Propaganda: Deconstructing Violent Extremism and Reimagining Counterterrorism in Contemporary Digital Ecosystems," *Frontiers in Political Science* 7 (2025). See also Weimann, Pack, Sulciner, Scheinin, Rapaport, and Diaz; and "Early Terrorist Adoption of Generative AI."
32 Paul Gill et al., "Malevolent creativity in terrorist organizations," *Journal of Creative Behavior* 47:2 (2013): pp. 125-151.
33 René Pita, "Assessing al-Qaeda's chemical threat," International Journal of Intelligence and CounterIntelligence 20:3 (2007): pp. 480-511; Toby Ewin, "Point of View: Inside the Mind of a Bioterrorist" in Filippa Lentzos ed., *Biological Threats in the 21st Century: The Politics, People, Science and Historical Roots* (London: Imperial College Press, 2026), pp. 198-210. 2016; Sammy Salama and Lydia Hansell, "Does Intent equal Capability? Al-Qaeda and Weapons of Mass Destruction," *Nonproliferation Review* 12:3 (2005): pp. 615-653.
34 "Share of internet users in Africa as of February 2025, by country," Statista, last accessed December 10, 2025.
35 "Early terrorist experimentation with generative artificial intelligence services," Tech Against Terrorism, November 2023.
36 Ardi Janjeva, Anna Gausen, Sarah Mercer, and Tvesha Sippy, "Evaluating Malicious Generative AI Capabilities: Understanding inflection points in risk," Centre for Emerging Technology and Security, July 2024. For the Voice of Khorasan advice, see Mason Boycott-Owen, "ISIS teaching recruits how to use AI 'responsibly,'" Politico, February 23, 2026.
37 Baele and Brace.
38 Cristina Archetti, "Terrorism, communication, and the media" in *Understanding Terrorism in the Age of Global Media: A Communication Approach* (London: Palgrave Macmillan UK, 2013): pp. 32-59.
39 Patti M. Valkenburg, Jochen Peter, and Joseph B. Walther, "Media effects: Theory and research," *Annual Review of Psychology* 67 (2016): pp. 315-338.
40 Daniel Kreiss, "Micro-targeting, the quantified persuasion," *Internet Policy Review* 6:4 (2017); Kobi Hackenburg and Helen Margetts, "Evaluating the persuasive influence of political microtargeting with large language models," *Proceedings of the National Academy of Sciences* 121:24 (2024).
41 Desiree Steppat, Laia Castro Herrero, and Frank Esser, "Selective exposure in different political information environments–How media fragmentation and polarization shape congruent news use," *European Journal of Communication* 37:1 (2022): pp. 82-102.
42 Herbert A. Simon, *Designing Organizations for an Information-rich World* (Baltimore: Johns Hopkins University Press, 1971), pp. 37-52.
43 Marina Adami, "AI-generated slop is quietly conquering the internet. Is it a threat to journalism or a problem that will fix itself?" Reuters Institute, November 26, 2024.
44 See, for example, Margaret Davis, "UK targeted by most cyber attacks from hostile states, security chief says," *Independent*, April 22, 2026.
45 See, for example, "ShinyHunters: Cyber Criminal Group Attacks Learning Management System," FBI, Alert Number: I-051526-PSA, May 15, 2026.
46 Blessing Guembe et al., "The emerging threat of AI-driven cyber-attacks: A review," *Applied Artificial Intelligence* 36:1 (2022); Jack Hughes, Ben Collier, and Daniel R. Thomas, "Stand-Alone Complex or Vibercrime? Exploring the adoption and innovation of GenAI tools, coding assistants, and agents within cybercrime ecosystems," arXiv preprint arXiv:2603.29545 (2026).
47 Timur Kuran and Cass R. Sunstein, "Availability cascades and risk regulation," *Stan. L. Rev.* 51 (1998): p. 683.
48 Hansard, March 18, 2003.
49 "Remarks by the President on Securing Our Nation's Cyber Infrastructure," The White House, May 29, 2009.
50 For researcher perceptions, see Stuart Macdonald, Lee Jarvis, and Simon M. Lavis, "Cyberterrorism today? Findings from a follow-on survey of researchers,"

Studies in Conflict & Terrorism 45:8 (2022): pp. 727-752. For possible explanations of the absence of cyber-terrorism, see Jeppe T. Jacobsen, "Cyberterrorism," *Perspectives on Terrorism* 16:5 (2022): pp. 62-72.

51 Maura Conway and Joseph Dillon, "Future trends: Live-Streaming terrorist attacks," *Vox-Pol*, 2016.

52 Jonathan Evans, "The Olympics and beyond: Address at the Lord Mayor's Annual Defence and Security Lecture by the Director General of the Security Service," published June 25, 2012.

53 Paul Gill, John Horgan, Samuel T. Hunter, and Lily D. Cushenbery, "Malevolent creativity in terrorist organizations," *Journal of Creative Behavior* 47:2 (2013): pp. 125-151.