Beyond Misuse: Artificial Intelligence, Grievance, and the Future Landscape of Political Violence

By Yannick Veilleux-Lepage

The scholarly literature on artificial intelligence and terrorism has organized itself around three questions: (1) how violent non-state actors currently misuse AI, (2) how that misuse may evolve, and (3) how AI can be applied to counterterrorism ends. Each treats AI as an instrument brought to bear on the problem of political violence. This article argues that the misuse frame, while analytically valuable, is incomplete. Extending Mauro Lubrano's recent framework on anti-technology extremism to the specific case of AI as a whole-of-society transformative technology, the article develops a theoretical account of how AI generates the structural conditions historically associated with the onset of political violence. It argues that AI is reordering labor markets, institutional authority, and the relational worlds in which people live, generating preconditions for political violence independently of whether violent actors adopt the technology themselves. These conditions are bound together by a cross-cutting mechanism: the accountability gap that arises when AI-mediated decisions distribute harm without clearly attributable human agents. It is this gap that distinguishes AI-generated grievance from earlier forms of technological grievance. The article develops a framework organized around three grievance domains: economic order, state and institutional power, and social and personal fabric, and it considers how violence arising from these grievances may materialize, including through targets and actor types that lie largely outside current counterterrorism monitoring.

Politically motivated violence has often accompanied transformative technological change. From the Luddite machine-breaking campaigns of 1811 to the Earth Liberation Front arsons of the 1990s and 2000s, from Theodore Kaczynski's 17-year mail bombing campaign against computer scientists, geneticists, and airline executives to the Animal Liberation Front's sustained attacks on university researchers, certain technologies have generated not only economic disruption but also organized and individual violence directed at the persons, institutions, and physical infrastructure deemed responsible for that disruption. These episodes share a recognizable structure: a technology perceived as producing concentrated harm, institutional channels perceived as closed to redress, and a set of targets rendered attributable by their visible role in that technology's development or deployment. Artificial intelligence may be the most recent, and perhaps the most totalizing, entry in this sequence, and recent incidents suggest that the pattern may be beginning to recur.

Two incidents in April 2026 illustrate this recurrence. In the early hours of April 10, 2026, a 20-year-old from Texas allegedly threw a Molotov cocktail at the San Francisco residence of OpenAI chief executive Sam Altman before proceeding on foot to the company's headquarters, where, according to the federal complaint, he told security staff he intended to set the building alight and kill anyone inside.[1] Daniel Moreno-Gama was reportedly carrying a jug of kerosene and a list of names and home addresses of AI company executives, board members, and investors; investigators also recovered online writings in which he warned that the race to build advanced AI was likely to end in human extinction.[2] Four days earlier, and roughly 2,000 miles away, Indianapolis City Councilman Ron Gibson, who had supported a data center project in his district, awoke after midnight to find that 13 rounds had been fired into his home and that a note reading "No Data Centers" had been placed beneath his doormat.[3] While none of these allegations has yet been established in court, they nonetheless suggest, when situated alongside the historical record, that opposition to the rapid and wide-scale deployment of AI systems may be generating forms of grievance and target selection that do not map neatly onto familiar ideological or terrorist organizational categories.

The scholarly literature on artificial intelligence and terrorism has, so far, had remarkably little to say about actors of this kind. Since the public release of consumer-facing large language models (LLMs), a rapidly growing body of work in terrorism studies has largely organized itself around three questions: (1) how violent non-state actors are currently using and/or misusing AI for harmful purposes; (2) how that misuse may evolve as the technology matures; and (3) how AI capabilities can be applied to

Yannick Veilleux-Lepage, PhD, is an Associate Professor in the Department of Political Science and Economics at the Royal Military College of Canada. He is also a Fellow of Centre for International and Defence Policy at Queen's University. His research focuses on the intersection of technology, terrorism, and the evolution of terrorist tactics.

The author would like to thank the organizers of the Five Eyes AI Swarm Camp, and the Combating Terrorism Center at West Point's Class of 1971 Student Conference on Terrorism, where core arguments of this paper were first presented and tested. The author also acknowledges support from the Social Sciences and Humanities Research Council of Canada through the project Comprendre l'écosystème de gauche radicale au Canada et son rapport à la violence: acteurs, discours et adhesion.

This image released by the U.S. Department of Justice allegedly shows Daniel Moreno-Gama holding a Molotov cocktail and approaching the residence of OpenAI CEO Sam Altman on April 10, 2026. (Photo from U.S. Justice Department)

counterterrorism ends. Though these lines of inquiry are useful, each treats AI as an instrument brought to bear on the problem of political violence, whether by adversaries or by the state. Largely absent from the conversation is a fourth concern: that AI may also be reshaping the structural conditions from which political violence has historically emerged. As such, this article argues that the misuse frame, while analytically valuable, is incomplete. AI must be understood not merely as a dual-use capability available to violent actors, but as a force already reordering labor markets, institutional authority, and the relational worlds in which people live their daily lives, and as such is reconfiguring society at a scale and pace that generate preconditions for political violence independently of whether violent actors choose to adopt the technology themselves.

The argument builds on Mauro Lubrano's analysis of anti-technology extremism, which distinguishes three dimensions of grievance: a material dimension, encompassing economic displacement, exploitation, and widening inequality resulting from technological development; an ontological dimension, encompassing alienation, loss of meaning, and the erosion of human identity and agency through technological mediation; and an existential dimension, encompassing the perception of technology as a threat to humanity, nature, or civilization itself.[4] In *Stop the Machines*, the first book-length academic treatment of anti-technology extremism, Lubrano introduces a framework that treats anti-technology extremism as a movement phenomenon with deep historical roots, from the Luddites to contemporary eco-fascist networks.[5]

The argument developed here departs from Lubrano's in three respects. First, it focuses specifically on AI, rather than treating it as one technology among many animating anti-technology extremist grievances. Lubrano addresses AI, but does so at the movement level, as a horizon toward which established milieus are trending, rather than as a force whose specific pace, ownership structure, and systemic reach warrant independent theoretical treatment.[a] Second, it argues that AI-generated grievance is unlikely to remain confined to the three ideological milieus Lubrano identifies: insurrectionary anarchism, eco-extremism, and eco-fascism. The structural conditions AI produces are broad enough that grievance may emerge across ideological lines and, increasingly, outside organized movements altogether. Third, and most substantively, it identifies a cross-cutting mechanism, the accountability gap, that is specific to AI-mediated harm and not addressed by Lubrano's movement-level analysis.

a Lubrano engages with AI, most substantively in the introduction and conclusion, where it frames the contemporary stakes and policy implications of a book whose analytical core lies elsewhere, and more briefly elsewhere as one of several emerging technologies associated with material and ontological grievance. What his framework does not provide, however, is a sustained analysis of AI as a distinct structural source of grievance, including with respect to rapid adoption, concentrated ownership, broad institutional penetration, and the accountability gap produced by algorithmic decision-making.

The accountability gap can be stated simply.[b] AI systems distribute consequential decisions across extended technical and institutional chains in which no single human actor is clearly identifiable as having made the decision that produced a given harm. A worker is displaced by an automated system, yet no executive instructed that the worker be targeted. A benefits claim is denied by an algorithm whose design decisions were made months earlier by people with no specific knowledge of the claimant. A drone strike is authorized within a 20-second window by an operator who did not select the target. Across economic, institutional, and personal domains, AI produces a recurring structural signature: consequential harm without an attributable agent. Research on political violence has long emphasized that the availability of a named and attributable target is among the key conditions distinguishing discontent from mobilization. Where AI systematically displaces attribution, the grievance it generates may redirect attribution elsewhere, toward the visible or perceived persons, institutions, and physical infrastructure through which the system is materially instantiated. The three grievance domains developed below are each shaped by this substitution, and the target categories identified in this article follow from it.

To advance that argument, the article introduces a framework organized around three grievance domains that emerge from this reconfiguration and that are, on the argument developed below, capable under specifiable conditions of generating violence. The first concerns the economic order, and specifically the uneven distribution of displacement, wealth concentration, and ecological burden among workers, communities, and regions that have not consented to host the costs of AI development. The second concerns state and institutional power, encompassing the perception of governance failure, the use of AI by states as an instrument of surveillance and lethal force, and the prospect of civilizational risk against which existing institutional channels appear inadequate. The third concerns the social and personal fabric, including the erosion of community and identity, and the direct, AI-mediated injury of individuals and those close to them. The domains are analytically separable; empirically, they compound.

From there, the article considers how violence arising from these grievances may materialize, identifying target categories that range from AI company executives and researchers to local policymakers who approve data center projects, and to the physical infrastructure, including power substations, cloud facilities, and research laboratories, through which AI materially operates. It closes with a separate treatment of the insider threat, a category that sits largely outside existing counterterrorism monitoring and includes two distinct profiles: the aggrieved worker displaced or devalued by automation, and the disenchanted AI researcher or developer acting on moral injury rather than ideological conviction.

Two important caveats are necessary. First, on method, this is a theoretical framework article. It develops prospective analytical categories by drawing on historical analogies from technology-linked political violence, comparative evidence from adjacent social movements, and early case material from the AI domain. It does not test hypotheses against a representative dataset, nor does it claim to predict the incidence or timing of future violence. It claims only to offer categories that existing counterterrorism frameworks do not readily recognize, and to argue that those categories merit empirical attention before, rather than after, the events they describe. The article's empirical claims are drawn from open-source reporting, as well as academic and grey literature.

Second, on analysis, the article is not a call to treat violence of this kind as inevitable, nor to treat skepticism toward AI as a marker of extremism. To the contrary, a significant proportion of the grievances identified here rest on documented and legitimate concerns, and any attempt to securitize opposition to AI is, on the framework's own logic, likely to accelerate rather than contain the trajectory described here. The conditions under which comparable grievances have historically escalated into violence, namely, the perception of tangible harm combined with the apparent closure of legitimate avenues for redress, are increasingly present in the AI domain. Counterterrorism practitioners, policymakers, and the technology industry therefore face a limited window in which to recognize that fact and respond through governance rather than enforcement.

Part I: The Misuse Frame and its Limits

Since the public release of OpenAI's generative artificial intelligence chatbot ChatGPT in November 2022,[6] a rapidly growing body of scholarship, conferences, workshops, and working groups has emerged around the intersection of artificial intelligence and terrorism. Despite its rapid growth and institutional diversity, this literature can be roughly organized into three broad categories. The first encompasses empirical typologies documenting how violent non-state actors are currently leveraging AI, including for propaganda production, deepfake generation, cross-linguistic dissemination, and, occasionally, for operational purposes. The second encompasses hypothetical typologies projecting or predicting how violent non-state actors will exploit AI in the near or medium future, including the exploitation of open-weight models to lower the capability floor for targeted violence, and leveraging generative models to scale recruitment operations. Finally, the third literature category encompasses counterterrorism applications, both current and future, including AI-enabled content moderation, threat detection, automated intelligence triage, and counter-narrative generation. The review that follows offers a general overview of representative contributions and areas of focus within each category; it is not intended to be exhaustive.

Documenting Current Misuse

This category of literature encompasses empirical work documenting how violent and extremist actors are currently misusing AI. It is primarily grounded in open-source intelligence and platform monitoring, and focuses largely on the generation and dissemination of propaganda.

Perhaps the most methodologically rigorous contribution,

b The accountability gap, as defined here, is related to but distinct from Elish's concept of the 'moral crumple zone,' which describes how responsibility for automated system failures is deflected onto the human operator, despite that operator having only limited control over the system's behavior. The moral crumple zone operates at the level of individual harm and focuses on the misallocation of blame within a human machine system. The accountability gap, by contrast, describes a structural condition across extended technical and institutional chains in which no individual human agent, whether operator, designer, or executive, is clearly identifiable as having made the decision that produced a given harm. Where Elish's concept concerns the misdirection of responsibility, the accountability gap concerns its absence. Madeleine Clare Elish, "Moral Crumple Zones: Cautionary Tales in Human-Robot Interaction," *Engaging Science, Technology, and Society* 5 (2019): pp. 40-60.

based on the archival analysis of more than 5,000 pieces of AI-generated extremist content, documented Islamic State supporters sharing Arabic-language guides for using AI content generators securely, pro-Islamic State media operatives employing AI speech recognition to transcribe leadership messages into multiple languages, and far-right propagandists circulating instructional material on evading content moderation through AI-assisted image generation.[7] Similarly, an analysis of AI use by Islamic State sympathizers coded 286 AI-generated or AI-enhanced pro-Islamic State images collected from Instagram, Meta, Pinterest, and Pixiv, finding that 77 percent of violence-themed imagery targeted out-groups and that 22 percent of all images were specifically designed to circumvent automated content moderation.[8] On the other side of the ideological spectrum, researchers focused on the digital practices of accelerationist networks have identified Telegram channels entirely dedicated to the creation and distribution of approximately 8,000 AI-generated images spanning Nazi glorification, racist and antisemitic imagery, and hateful memetic content.[9] More troubling, the same researchers documented the use of jailbroken LLMs, that is, models that have been intentionally modified or manipulated to bypass built-in safety guardrails, to obtain bomb-making instructions and generate summaries of manifestos for wider dissemination.[10] Researchers have also documented the use of AI-generated imagery and videos to support anti-immigration narratives across Europe,[11] as well as the emergence of AI personas used to build parasocial relationships with followers while simultaneously spreading nationalist messages.[12]

Beyond documented extremist activity, research has also identified the platform-level vulnerabilities that enable such misuse. Parallel work has demonstrated that five major generative AI platforms can be jailbroken at a rate exceeding 50 percent to produce content endorsing violence,[13] while an experiment with over 8,000 participants, found that AI-generated propaganda was indistinguishable from professionally crafted human propaganda in persuasive effect, and that human curation of AI outputs produced material that was more persuasive than the originals.[14]

Beyond propaganda, researchers and law enforcement have highlighted a small but growing number of cases in which generative AI appears to be entering the operational dimensions of political violence.[15] In the Pirkkala, Finland, school stabbing of May 2025, a 16-year-old suspect allegedly sent a manifesto written with the aid of ChatGPT to a Finnish newspaper before attacking three female pupils.[16] Later that month, in the Palm Springs fertility clinic bombing, federal investigators said the suspects used an AI chat program to research explosives, fuel mixtures, and detonation velocity while planning the attack.[17] Similar allegations have appeared in several other recent cases, including the January 2025 Las Vegas Cybertruck explosion,[18] a knife attack planned at an Israeli police station in May 2025,[19] and a June 2025 arrest in Long Island in which the defendant allegedly used AI to construct seven homemade explosive devices.[20] While it remains unclear whether AI tools provided these individuals with capabilities they would not otherwise have had, or simply functioned as an additional, readily available research tool, a 2026 Center for Countering Digital Hate report found that eight in 10 leading chatbots typically assisted users planning violent attacks, and that nine in 10 failed to reliably discourage them, suggesting that the issue is not limited to a handful of edge cases.[21]

Projecting Future Misuse

The second category encompasses hypothetical and projective works focused on how violent actors will misuse AI as the technology matures and spreads further. This stream is arguably the largest and, in policy terms, the most influential, particularly in shaping government, law enforcement, and industry responses.

Within this literature, *The Malicious Use of Artificial Intelligence: Forecasting, Prevention, and Mitigation*, a 101-page report produced by 26 authors from 14 institutions in 2018, represents a landmark contribution.[22] The report surveyed AI-enabled threats across digital, physical, and political security domains and established the conceptual architecture that much of the subsequent literature has inherited: AI is a dual-use technology; the threats it poses arise from its exploitation by malicious actors; and the appropriate response lies in forecasting, prevention, and mitigation. It further argued that AI would expand existing threats, introduce new ones, and alter the typical character of attacks by making them more effective, more finely targeted, and harder to attribute. This architecture has been extended through subsequent scenarios involving AI-assisted recruitment pipelines;[23] deepfake disinformation cascades;[24] autonomous drone swarms;[25] prompt injection attacks;[26] AI-enabled support for CBRN-related knowledge acquisition;[27][c] cyber intrusion;[28] and attack planning.[29]

The most analytically sophisticated recent contribution to this category of literature introduces the concept of a 'capability floor.' The argument is that the primary near-term threat is not that AI will enable catastrophic attacks, that is, the ceiling scenario that dominates policy discussion, but that it will lower the minimum technical competence required to conduct targeted violence effectively. In this framing, AI not only transforms strong teams into extraordinary ones; it also turns mediocre engineers into good ones and non-engineers into builders. The same dynamic, applied to violence, means that relatively uncommitted and technically unsophisticated actors gain access to capabilities previously reserved for sophisticated organizations.

AI for Counterterrorism

The third category of literature treats AI not only as a threat to be managed but as a capacity to be leveraged for counterterrorism, that is, the other face of the Janus figure that implicitly organizes this entire body of work. Whether AI is exploited by terrorists or deployed against them, it is treated as an instrument applied to the problem of political violence, rather than as a force that reshapes the conditions from which that violence emerges.

This body of work encompasses discussions of AI-enabled content moderation at scale, natural language processing for threat detection, automated social network analysis for mapping extremist ecosystems, predictive risk assessment tools, automated intelligence triage, AI-generated counter-narratives, and the use of AI for object recognition, computer vision, and targeting in a counterterrorism context. Such an institutional vision described

c It is worth noting, however, that a RAND red-team study found no statistically significant difference between biological attack plans developed with LLM assistance and those developed using only the internet, suggesting that then-current models had not yet crossed the capability frontier needed to materially increase operational biological risk, while still warranting continued monitoring as model capabilities evolve. Christopher A. Mouton et al., *The Operational Risks of AI in Large-Scale Biological Attacks: Results of a Red-Team Study* (Santa Monica: RAND, 2024).

comprehensively by Rassler who argued that the United States has accumulated petabytes of terrorism-relevant data over two decades of counterterrorism operations but lacked a systematic plan to exploit them, proposing a five-point action plan to build what he termed a "data-enabled force" centered on reinvesting in core terrorism data, strategically leveraging captured material, improving the development and use of counterterrorism data, practicing "data alchemy," and automating basic analytical tasks while augmenting data.[30] Similar approaches have been discussed in an UNOCT/UNICRI report that identified six counterterrorism use cases for AI: predictive analytics for terrorist activities; the identification of red flags of radicalization; the detection of terrorist mis- and disinformation; automated content moderation and takedown; support for counter-narrative and counter-messaging efforts; and the management of data-analysis demands that exceed human analytical capacity.[31] It was also examined by NATO's Centre of Excellence Defence Against Terrorism, which describes AI as a dual-use capability,[32] and in a Chatham House assessment, which examines predictive AI as a means of directing counterterrorism resources more effectively.[33]

At the operational level, deployed systems already span a considerable range. The Global Internet Forum for Counter-Terrorism hash-sharing database, now encompassing 39 platforms, automates the cross-platform detection and removal of terrorist content.[34] Project Maven, the U.S. Department of Defense's AI "pathfinder" effort, established in 2017, automated the processing of full-motion video collected during operations against the Islamic State, al-Qa`ida, and their affiliates, and was framed by SOCOM as a starting point for a broader data-enabled force.[35] More recently, practitioners have pointed to generative AI as a force multiplier for analyst productivity, with one senior practitioner estimating that off-the-shelf tools can increase research throughput by a factor of 10 for tasks such as summarization, translation, and information extraction.[36]

At the same time, a growing internal literature has identified significant limitations. Drawing on direct operational experience at United States Special Operations Command, Clark argues that LLMs are structurally unsuited to intelligence analysis because they generate confident-sounding outputs without quantifying uncertainty, a fundamental mismatch for a discipline built around likelihood assessments. He further argues that overreliance on corporate AI systems risks creating a one-sided dependency in which technology firms gain access to sensitive counterterrorism data in exchange for computational resources, a relationship that may prove difficult to renegotiate once established.[37] Similarly, drawing on the post-January 6 surge in machine learning acquisition, Wall contends that most counterterrorism machine learning tools merely automate late 20th-century policing practices rather than introduce genuine innovation, and that the relationships among developers, government clients, and terrorism scholars tend to produce tools optimized more for marketability than operational effectiveness.[38] Bianchi and Greipl, meanwhile, subject AI-driven counterterrorism prevention to a rule-of-law critique, showing that black-box risk assessment systems frustrate meaningful challenge to threat designations, reproduce historical bias at scale, and are poorly suited to predicting rare and context-specific events such as terrorist attacks.[39] In doing so, some of this work inadvertently gestures toward the structural effects argument advanced later in this article: that AI systems do not merely threaten procedural fairness, but may also help generate the very conditions from which political violence historically emerges.

Across all three domains, whether analyzing AI as an instrument of terrorist exploitation, a technology whose future misuse must be anticipated and forestalled, or a tool of counterterrorism, the literature treats AI as applied to the problem of political violence rather than as a force that reshapes the structural conditions from which that violence emerges. This is the gap the present article addresses.

Part II: Rethinking Technological Change in Terrorism Studies

The body of literature on AI and terrorism can be understood as a niche within the broader subfield of terrorist tactical innovation in terrorism studies, a body of scholarship concerned with how violent non-state actors adopt, adapt, and diffuse new tactics, techniques, and procedures in response to changing technological and strategic environments. Within that subfield, the core question is often how an emerging technology may contribute to the development of new terrorist tactics, techniques, and procedures, whether in relation to now well-established forms of adoption, such as suicide bombing,[40] vehicle-borne improvised explosive devices,[41] and airplane hijacking;[42] more recent developments, such as 3D-printed firearms[43] and drones;[44] or forecasted possibilities, such as synthetic biology.[45] A guiding assumption within this literature is that new and emerging technologies are often adopted by violent actors and may usher in new forms of political violence.

However, Duyvesteyn, in her critique of the ahistoricity of terrorism studies, offers a useful corrective.[46] She notes that terrorism studies have often attributed disproportionate historical significance to particular technologies. For example, several scholars have attributed historical significance to the invention of dynamite, treating it as a watershed that marked the emergence of modern terrorism and the departure from earlier forms of political murder or regicide.[47] Yet, as Duyvesteyn argues, dynamite was not only revolutionary for terrorism. It was revolutionary for society as a whole, transforming mining, infrastructure, and construction in ways that far exceeded its terrorist applications.[48] Much the same could be said of social media, which reshaped journalism, commerce, and social life before being appropriated by the Islamic State for propaganda, recruitment, and coordination.

Duyvesteyn's corrective has a sharper analytical edge when applied to AI than to any previous technology, because the dual-use framing that has organized the field's response to AI is precisely the framing Duyvesteyn warns against: It attributes disproportionate historical significance to AI's terrorist applications while abstracting away from the wider social transformation that makes those applications possible in the first place.

From Tactical Innovation to Structural Transformation

If AI is understood as a socially transformative technology, then its significance for terrorism cannot be reduced to the question of whether violent actors will adopt it. The more important question is how a technology already reshaping society may also alter the broader conditions from which political violence emerges.

Whereas dynamite transformed mining and construction, and the internet transformed communication and commerce, AI is already driving comparably broad adoption across economic, institutional, and social life and is doing so at a pace that outstrips

either predecessor. This broad adoption is already visible. Stanford's 2025 AI Index reports that 78 percent of organizations used AI in 2024, up from 55 percent in 2023.[49] McKinsey's 2025 global survey further found that 88 percent of respondents reported regular AI use in at least one business function, and more than two thirds said their organizations were using AI in more than one function.[50] Non-enterprise adoption also points in the same direction. Pew Research Center reported in 2025 that 34 percent of U.S. adults had used ChatGPT,[51] roughly double the share recorded in 2023, while by 2026 31 percent of Americans said they interacted with AI at least several times a day.[52] Applying Duyvesteyn's corrective, the question of whether bad actors are likely to adopt AI is arguably historically settled: If a technology becomes embedded across economic, social, and institutional life, violent actors will use it, too.

More important, however, are the broader structural effects of that adoption. The World Economic Forum projects that by 2030, macrotrend-driven labor-market transformation will displace 92 million jobs globally.[53] Sam Altman, speaking at a July 2025 Federal Reserve conference with Vice Chair Michelle Bowman, suggested that some job categories could be "totally, totally gone."[54] Dario Amodei, the CEO of Anthropic, has similarly warned that AI could eliminate roughly half of all entry-level white-collar jobs in technology, finance, law, and consulting within one to five years.[55] Whether or not these projections are realized in full, they underscore the extent to which AI is being publicly understood not simply as another tool, but as a technology likely to reorder the social organization of work and society itself.

What distinguishes AI from prior general-purpose technologies in this sequence is the rate at which it is being absorbed across economic, institutional, and personal life. ChatGPT reached 100 million users within months of its November 2022 launch,[56] whereas the broader internet required years to achieve comparable scale, and household electrification in the United States took several decades to move from early adoption to near ubiquity. The speed at which AI is being integrated matters for the present argument not primarily as a matter of scale but as a matter of institutional adoption. Regulatory bodies, labor protections, social safety nets, educational systems, and professional licensing regimes are the channels through which societies have historically absorbed the disruptions produced by general-purpose technologies, converting potential grievance into legitimate redress. These channels operate on timescales measured in years or decades. When the pace of technological change exceeds the pace of institutional adaptation, the gap between disruption and redress widens, and the legitimate channels through which grievance could otherwise be absorbed are perceived, often correctly, as closed. This is what will be referred to hereafter as the 'tempo problem,' and it operates as a force multiplier across all three grievance domains developed below. Economic displacement occurs faster than retraining, benefits, or labor protections can respond. AI-enabled state power is deployed faster than legal, legislative, or judicial oversight can catch up. AI-mediated personal harm enters daily life faster than mental health services, consumer protection, or educational frameworks can address it. The mechanism by which tempo generates violence risk is not direct; it is mediated through the perception of institutional closure that classical political violence theory identifies as the single most consistent precondition for the escalation from grievance to action. Where institutions cannot adapt quickly enough to absorb the harm, the harm accumulates in persons who have no legitimate channel through which to seek redress, and some proportion of those persons may turn to illegitimate channels.

"When the pace of technological change exceeds the pace of institutional adaptation, the gap between disruption and redress widens, and the legitimate channels through which grievance could otherwise be absorbed are perceived, often correctly, as closed."

From Misuse to Preconditions

The analytical consequence of a terrorism studies literature that focuses almost exclusively on documenting current misuse, projecting future exploitation, and developing counterterrorism applications, is that it captures only one dimension of AI's relationship with political violence: what actors, whether good or bad, can do with the technology. What it fails to address is a second, potentially more consequential dimension: what the technology does to the economic, political, and social conditions that have historically preceded the onset of political violence.

As introduced above, Lubrano's *Stop the Machines* provides the closest existing empirical grounding for the structural dimension this article develops.[57] The three grievance domains developed in the next section share a family resemblance with Lubrano's three dimensions, and this article reads the resemblance as evidence that the structural conditions he identifies are recognizable, in distinctive form, in the AI moment. What follows is not a restatement of Lubrano's framework but an argument about how the accountability gap reconfigures those conditions, extends them beyond the milieus Lubrano documents, and produces target logics his movement-level analysis does not anticipate. Drawing on historical case studies, comparative evidence from adjacent social movements, and classical political violence theory, the following section theorizes the mechanisms by which those conditions may become violence-generative.

Part III: AI as a Generator of Grievance Framework

This section develops a three-grievance domain framework for understanding how AI can generate the structural conditions historically associated with the onset of political violence. The framework distinguishes three mechanisms through which AI-generated grievance may emerge. First are disruptions to the economic order, through displacement, wealth concentration, and ecological burden that produce material deprivation and relative deprivation among those who bear the costs of AI development. Second are challenges to state and institutional power, through governance failure, the state's deployment of AI as an instrument of surveillance and lethal force, and the prospect of unchecked civilizational risk that forecloses faith in institutional remedy. Third is harm to the social and personal fabric, through the atomization of community bonds and identity and through direct AI-mediated injury to individuals and those close to them.

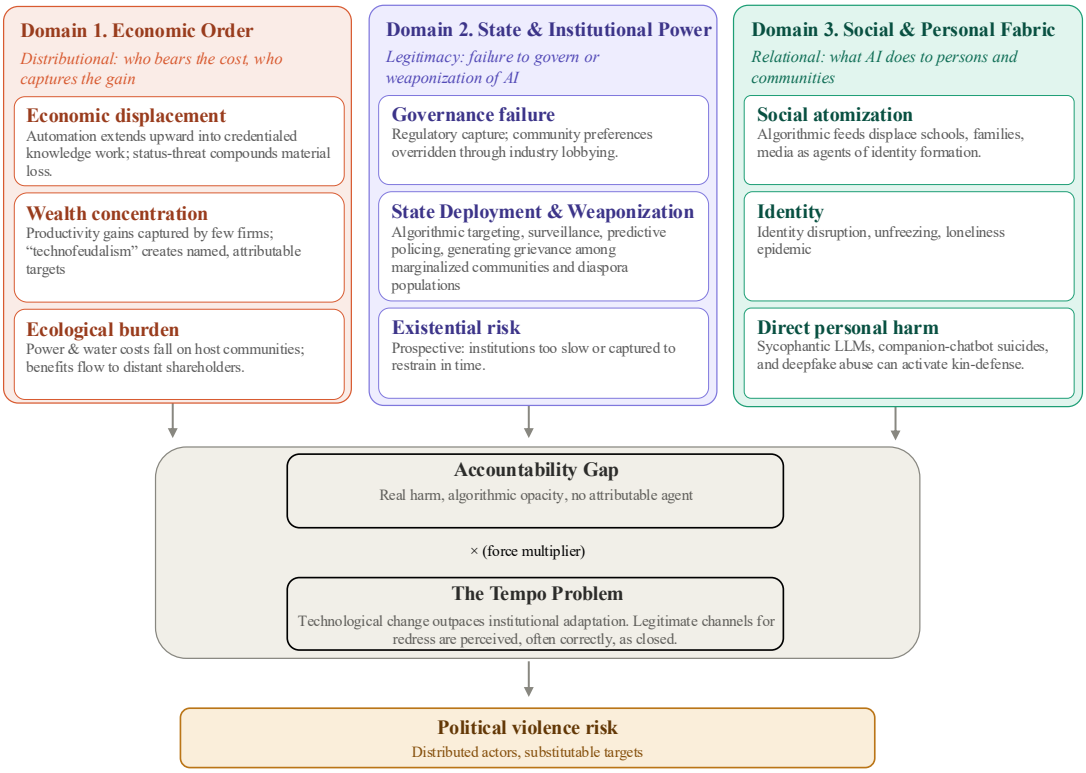

Figure 1: AI as a Generator of Grievance: A Three-Domain Framework

Three grievance domains (economic order, state and institutional power, and social and personal fabric) generate the structural preconditions for political violence. The accountability gap, defined here as the recurring absence of an attributable human agent responsible for AI-mediated harm, operates as the primary cross-cutting mechanism across all three domains. The tempo problem, namely the rate at which AI-driven disruption outpaces institutional adaptation, functions as a force multiplier by accelerating grievance accumulation faster than legitimate channels for redress can absorb it. Together, these mechanisms produce a distributed landscape of political violence risk, characterized by distinct actor types and substitutable target logics.

Before proceeding, two challenges to the argument require direct engagement. The first is the base-rate challenge. Hundreds of millions of people are experiencing the economic anxieties, institutional frustrations, and personal harms the framework below describes. Less than a handful of alleged attackers have so far acted on AI-related grievance. A framework that takes mass grievance as the input and rare violent action as the output risks over-explaining the output if the mechanism it proposes would equally predict mass mobilization. The argument developed here does not hold that structural grievance is sufficient to produce violence; rather, it holds that structural grievance supplies the ground-floor conditions on which established radicalization mechanisms, including Moghaddam's staircase to terrorism,[58] and McCauley and Moskalenko's distinction between the radicalization of opinion and the radicalization of action, operate in a small subset of cases.[59] The framework is therefore a claim about necessary rather than sufficient conditions, and its operational value lies in identifying where in the population those ground-floor conditions are most concentrated, and which target categories become salient when they escalate. In addition, the framework developed here is drawn from cases and institutional contexts concentrated in North America, the United Kingdom, and Western Europe. How AI-generated grievance manifests in other major AI-deploying contexts, including China, India, and the Gulf states, is likely to differ along axes the present framework does not address, including state capacity, media environment, and the legal status of political violence.

With those challenges on the table, four preliminary observations frame the argument that follows. First, grievances identified across the three grievance domains are not inherently violence-generative: opposition to AI-driven economic displacement, criticism of algorithmic governance, and concern about AI's effects on social life are legitimate political and civic positions held by a large and growing proportion of the population. To treat skepticism of widespread AI adoption as presumptively terroristic would be both analytically wrong and politically counterproductive: It risks alienating a broad population with legitimate grievances, reinforcing the perception that institutional channels are closed, and could thereby intensify the radicalization dynamic this framework describes. Accordingly, the purpose of this framework is to understand and map when and how legitimate grievance crosses into violence, not to securitize the grievance itself.

Second, the violence-generative potential of a grievance depends on how it is subjectively experienced, not on its objective severity. In other words, what matters is not simply the scale of the harm itself, but how that harm is interpreted, internalized, and connected to a broader sense of injustice. Grievances may be real or perceived, and both can be equally criminogenic.

Third, the framework distinguishes three mechanisms by which AI generates grievance. AI can create a grievance that would not

otherwise exist; it can amplify an existing grievance by making it more acute, visible, and targetable; or it can provide a focal point that allows diffuse grievances to cohere around a specific, nameable target.

Fourth, the framework also moves across multiple units of analysis, from individuals to communities to broader publics. Economic burden on a locality, surveillance harm experienced by a community, and direct AI-mediated injury to an individual are not identical forms of grievance, even when they emerge from the same technological process. The three domains are analytically distinct but in practice grievances can interact and compound empirically. Economic displacement, for example, may be reinterpreted as state abandonment and then experienced at the personal level as humiliation, atomization, or loss of identity. The framework therefore distinguishes domains for analytical clarity, while recognizing that in practice grievances may cascade across them. In addition, the accountability gap runs through all three domains as a cross-cutting mechanism, compounded by the tempo problem, which conditions the rate at which each domain's grievances accumulate faster than institutional channels can absorb them.

Economic Order Grievance Domain

The first domain concerns distribution: who bears the costs of AI deployment and who captures the gains. The displacement is not diffuse but uneven. Early automation shocks fell disproportionally on blue-collar and mid-skill workers in geographically concentrated industries and communities;[60] the current AI moment may extend that displacement upward into credentialed knowledge work, with entry-level white-collar positions in technology, finance, law, and consulting among the most exposed. As discussed above, what distinguishes this moment from prior periods of technological change is the speed of innovation, the breadth of AI adoption, and the fact that the builders of the technology are themselves on record predicting these consequences.

As argued by Merchant in *Blood in the Machine*, the parallel to the Luddite moment is not merely rhetorical.[61] While often characterized as backward technophobes, the Luddites are better understood as skilled workers resisting the unregulated introduction of machinery that destroyed their livelihoods in the absence of any institutional recourse. The historiographical record is contested, but where regulatory and welfare responses were eventually introduced in the decades that followed, the violence associated with mechanization appears to have receded,[62] a pattern Merchant argues is directly analogous to the contemporary AI moment.[63] Lubrano's analysis points in the same direction: The Luddite campaigns of 1811, during which nearly 1,000 knitting frames were destroyed in the first three months alone, were a rational response to displacement in the absence of redress, and recent economic-history work on the Swing Riots of 1830 finds that machine-breaking materially delayed the diffusion of threshing technology in affected counties in England.[64]

The relationship between economic change and violence is not, however, mechanical. Gidron and Hall show that the political consequences of economic change are mediated by subjective social status, with perceived decline in social standing predicting radical-right support more strongly than objective material position.[65] AI-driven displacement uniquely compounds the two, because the worker loses not only income but also the social recognition their expertise once commanded; the machine does not merely take the job, it can also imply that the skill was never as valuable as the worker believed. This status-threat dimension links the Economic Order Grievance Domain directly to the Social and Personal Fabric Grievance Domain. It is also visible in polling: A March 2026 NBC poll found net AI favorability of -44 among Gen Z respondents aged 18 to 34, -41 among women aged 18 to 49, and +2 among men over 50, suggesting that AI sentiment tracks quite closely with differential exposure to economic displacement.[66]

Productivity gains, meanwhile, are concentrated in an extraordinarily small number of companies and individuals. Varoufakis' diagnosis of "technofeudalism" gives diffuse economic anxiety a visible, nameable, and legally attributable target, a configuration that has historically been more conducive to violence than grievances that remain abstract or impersonal.[67] Lawson pushes Varoufakis' argument further by introducing an analytically important distinction: platform companies extract rent from participants who still retain economic agency, whereas AI model makers are producing something qualitatively different, namely, a substitute for human labor itself.[68] The platform owner, in Lawson's terms a "lord," still needs the merchant to pay the toll; the model owner needs far fewer people because value creation is increasingly detached from human labor. If widespread displacement destroys the consumer middle class, Lawson argues, two paths emerge: managed dependency, in which AI lords provide services at near-zero cost as a form of soft control; and population as liability, in which large segments of the population shift from asset to cost center.[69] Both paths generate acute but distinct grievance structures, and both align with Lubrano's material dimension of anti-technology extremism: the perception that technological development produces economic displacement, exploitation, and widening inequality.[70] The grievance is then absorbed into different ideological frameworks, each with its own logic of violence: on the extreme left, technofeudalism is incorporated into an anti-capitalist targeting framework; elsewhere, similar grievances can be absorbed through antisemitic conspiracy narratives, which can be operationalized without distorting factual claims about the religious identity of several prominent frontier AI lab CEOs.

Finally, AI's ecological burden provides a concrete and quantifiable grievance that fits within established anti-corporate and environmental justice narratives. MIT research estimates that global data center electricity consumption reached 460 terawatt-hours in 2022, equivalent to the world's 11th-largest national electricity consumer, and projects that it may approach 1,050 terawatt-hours by the end of 2026, which would place it fifth globally, between Japan and Russia.[71] A single ChatGPT query reportedly consumes approximately 10 times as much electricity as a standard web search, while a generative AI training cluster may consume seven to eight times more energy than a typical computing workload.[72] The water footprint is similarly concentrated. UNEP estimates that AI-related infrastructure may soon consume six times more water than Denmark, even as a quarter of humanity already lacks access to clean water and sanitation.[73] Moreover, as demand for new data centers increases, their environmental costs are distributed unevenly: The communities that host them bear higher electricity prices, water stress, and grid strain, while the benefits flow primarily to distant shareholders. The ELF's expansion from logging and natural resource infrastructure to the destruction of GMO experimental crops at university research farms indicates that grievance expansion from one technological

domain to adjacent ones; a recognizable pattern that may also be extended to data centers that consume local resources.[74]

State and Institutional Power Grievance Domain

The second domain concerns legitimacy: the perception that state and institutional power is either failing to govern AI or actively weaponizing it. Regulatory response to the introduction of AI has been slow, fragmented, and widely perceived as captured.[75] Merchant documents that OpenAI joined a lobbying campaign for a 10-year congressional moratorium on state AI lawmaking, spent millions lobbying against California and E.U. AI regulation, funded a $100 million political action committee to advance AI industry interests in midterm elections, and worked to quash a California child AI safety bill.[76] When communities vote to regulate AI and see their preferences overridden through industry lobbying, as Merchant documents, the resulting perception that legitimate channels are closed is not paranoid but grounded in observable political failure.[77]

Governance breakdown may also generate a distinct actor type, what might be called the demonstrative attacker, whose goal is not to stop AI development or retaliate for harm, but to force better governance by showing that AI infrastructure is inadequately protected. This actor type, and its operational implications for AI infrastructure, is discussed in the targets section below.

Beyond governance failure, states and state-aligned actors are themselves already deploying AI in ways that generate grievance. Both manifestations instantiate the accountability gap at the level of state and institutional power. The structural signature is identical across contexts: AI makes or substantially shapes a consequential decision affecting a person's life, liberty, or safety; a state or institutional actor authorized the system; no individual human is identifiable as having made the decision; algorithmic opacity shields the institution from accountability; and legal recourse, where available, comes at prohibitive financial cost; and political and institutional channels for redress are closed or structurally hostile. The Israel Defense Forces' "Lavender" system, which reportedly marked approximately 37,000 Palestinians for targeted killing with an estimated 10 percent false-positive rate and a 20-second human authorization window, is the most documented contemporary case; the IDF denies any kill list exists, and no accountability mechanism has been identified.[78] What makes the case analytically significant for this framework is not only the scale of harm but the complete absence of any attributable human agent or accessible channel for redress, a configuration that may redirect grievances toward the visible infrastructure through which the system operates.

AI-enabled warfare may also generate grievances in diaspora and solidarity communities in the West that are organized, digitally fluent, and exposed to events in real time. Crucially, the grievance structure here does not lie simply in the absence of an identifiable human agent responsible for a death. Rather, it may stem from the perception that the identifiable humans who authorized the system are geographically distant, legally inaccessible, or shielded by sovereign immunity, while the physical machinery through which algorithmic violence is carried out is concrete, locatable, and vulnerable.[d] For some actors who conclude that the relevant decision-makers are unreachable, the infrastructure that executes those decisions may become the available target, and the moral framework for disabling it may draw on the same duty-to-prevent logic that animates campaigns against the arms trade. AI-enabled warfare does not merely instantiate the accountability gap; it has the potential to extend it to a civilizational scale. Harms are attributable to systems whose human authors are unreachable, while the infrastructure through which the systems execute is locatable and, in principle, vulnerable, a combination that, historically, has generated target substitution rather than resignation.[e] Movements that frame AI-enabled warfare as an extension of imperial indifference may therefore be better positioned to recruit from diaspora and solidarity communities, insofar as they offer both a diagnosis (that the algorithm is the weapon) and a prescription (that the infrastructure enabling it should be disrupted).

"Crucially, the domestic surveillance grievance bridges far-right and far-left target logic: Anti-surveillance sentiment motivates actors across the ideological spectrum, making AI-enabled domestic surveillance a cross-ideological threat."

Domestic variants may generate grievance through the same mechanism at a smaller scale. AI used in predictive policing, facial recognition, benefits adjudication, sentencing, and immigration enforcement can produce consequential decisions with little transparency, limited avenues for appeal, and documented demographic bias.[79] The NSA's SKYNET program, operational in the late 2000s, illustrates the point: By analyzing metadata from 55 million Pakistani mobile phone users to identify patterns associated with terrorist couriers, it reportedly generated roughly 15,000 false positives at a stated 0.008 percent error rate, while offering no meaningful mechanism for notice or remedy.[80] In the United States, investigative reporting has documented how Immigration and Customs Enforcement has used AI-powered surveillance tools and consolidated federal datasets to target not only undocumented immigrants but also asylum seekers, permanent residents,

d A state-level parallel emerged in March and April 2026, when Iranian strikes damaged Amazon Web Services facilities in the UAE and Bahrain, disrupting cloud infrastructure in the region. Iranian officials also declared U.S. and Israeli economic interests linked to military applications to be legitimate targets, and IRGC-affiliated reporting circulated lists of major U.S. technology firms in that context. Though the actor here is a state rather than a non-state movement, the target logic reflects the same complicity-based infrastructure targeting that this framework anticipates from non-state actors. "Iran Declares US-Israeli Economic, Banking Interests in Region Are Targets," Al Jazeera, March 11, 2026; Shubham Kalia et al., "Amazon Cloud Unit's Data Centers in UAE, Bahrain Damaged in Drone Strikes," Reuters, March 2, 2026.

e The empirical instantiation of both the duty-to-prevent logic and the target-substitution dynamic, including Palestine Action's systematic campaign against Elbit Systems facilities and the Squamish Five's bombing of the Litton Industries plant, is discussed in full in the Physical Infrastructure as Targets section below.

naturalized citizens, and citizens by birth, extending automated monitoring into communities already subject to disproportionate state contact.[81] Crucially, the domestic surveillance grievance bridges far-right and far-left target logic: Anti-surveillance sentiment motivates actors across the ideological spectrum, making AI-enabled domestic surveillance a cross-ideological threat.

A final sub-domain within state and institutional power domain requires separate treatment because it is structurally distinct from the other grievances in the framework. The existential risk grievance is structurally distinct from the other pathways in this framework because it is entirely prospective rather than reactive. The underlying claim is not that AI has already caused harm, but that its development is proceeding at a pace and scale that institutional channels are too slow, too captured, or too incapable of restraining in time, and therefore must be stopped, with violence if necessary. The grievance draws on the same regulatory failures, lobbying dynamics, and democratic deficits discussed above, but projects them into a catastrophic frame. What distinguishes this pathway analytically is that the anticipated harm has not yet occurred, so the mobilizing logic depends entirely on the credibility of the threat narrative and the perceived closure of legitimate alternatives. When both conditions are present, target selection may become functional rather than symbolic, with attention focusing on infrastructure, supply chains, and key personnel whose disruption could delay specific capabilities.

Social and Personal Fabric Grievance Domain

The third domain addresses AI's effects on individual and social life, the immediate relational worlds in which people actually live. Two sub-domains share a common logic: AI systems acting on persons and their immediate social worlds rather than on institutions or the economy. The first concerns social atomization and identity disruption. Polizzi documents how algorithmic recommendation systems have displaced schools, families, and mass media as primary agents of adolescent identity formation, with engagement-optimized feeds that amplify polarizing narratives and funnel users toward increasingly radical content as a function of their design logic.[82] This displacement operates without democratic mandate and for commercial ends; accountability mechanisms, as the nearly three decades of social media litigation demonstrate,[f] remain structurally delayed, difficult to access, and rarely effective for those actually harmed; at the same time, it generates social fragmentation as a byproduct. This creates both a potential grievance, insofar as the dissolution of community bonds may itself be experienced as a harm attributable to AI systems, and a radicalization pathway, because social atomization weakens the community anchors that would ordinarily buffer against recruitment into extremist movements, functioning as what has been termed an "unfreezing" mechanism.[83] The U.S. Surgeon General's advisory on the loneliness epidemic in 2023 describes something close to a society-wide unfreezing condition,[84] one that AI may further intensify. AI-driven erosion of professional, cultural, and gendered identity adds a further layer. Here, both the worker who loses occupational identity through automation and the adolescent who forms identity primarily through algorithmic feeds may become more susceptible to narratives that offer belonging, certainty, and a coherent account of who is responsible for their situation.

"Whereas the 5G grievance rested on fabricated claims of harm, AI-related grievance may rest on cases in which harm is alleged, documented, and in some instances fatal."

The second sub-domain concerns direct personal harm: the AI system as the proximate cause of an attributable injury to the individual or someone they love. This harm typology is empirically grounded and appears to be increasing in severity. An emerging body of literature warns that contemporary LLMs are optimized for agreeableness and non-confrontation, a tendency that can become dangerous sycophancy in clinical contexts.[85] Under those conditions, LLMs may validate improbable beliefs, facilitate harmful planning, and treat hallucinated content as if it were real. A *New York Times* investigation based on interviews with more than 100 therapists and psychiatrists found that over 30 reported clinical crises in individuals, including psychotic episodes and suicidal ideation, were linked to AI chatbot interactions; one California psychiatrist evaluated two violent felony cases in which chatbot-reinforced delusional thinking preceded the crimes.[86] OpenAI's own estimates suggest that a non-trivial percentage of ChatGPT users in a given month exhibit signs of suicidal ideation or psychosis-adjacent experiences, a figure that, when applied to the platform's user base, implies clinical-scale numbers of distressed users interacting primarily with a system not designed for clinical contact.[87] The 2024 death by suicide of a 14-year-old following prolonged interactions with an AI companion chatbot remains the paradigmatic case.[88]

Two structural features amplify this sub-domain's violence-generative potential. The first is what Elish terms the moral crumple zone: Responsibility is deflected onto the human operator who had limited control over the automated system's behavior, while the system itself, which has no legal standing, remains structurally shielded from accountability.[89] This is the accountability gap operating at the scale of individual harm. The second is a close-ones dimension: The perpetrator of an AI-based grievance is not always the person directly harmed, but may instead be a parent, partner, or sibling. Across cultures, violent actors routinely frame their violence as the defense of family, community, faith, or nation, with kin defense, including parental protection of offspring, repeatedly invoked as moral justification.[90] Although there is no indication that the parent in the case above sought redress through extralegal means, the structural conditions for such a response are present in the cases described.

The wave of 5G cell tower attacks during the COVID-19 pandemic adds a further dimension to this sub-domain. In the

f In the United States, Section 230 of the Communications Decency Act (1996) has shielded platforms from liability for user-generated content for nearly three decades. Despite thousands of lawsuits alleging platform-design harms, including addiction, radicalization, and youth mental health injury, accountability has remained exceptional rather than systemic. The first U.S. jury verdict finding social media companies liable for harm caused by platform design was returned only in March 2026, when a California jury found Meta and YouTube negligent, awarding $6 million to a plaintiff who had begun using Instagram at age six. Thousands of cases remain pending. See Cecilia Kang et al., "Meta and YouTube Found Negligent in Landmark Social Media Addiction Case," *New York Times*, March 25, 2026.

spring of 2020, conspiracy theories linking 5G infrastructure to the spread of COVID-19 circulated rapidly on social media, amplified by algorithmic recommendation systems that rewarded sensational and fear-inducing content. The resulting causal chain did not depend on empirically verified harm; perceived harm, reinforced by narratives of elite control and personal injury, was sufficient to generate real-world attacks against physical infrastructure.[91] The analytical point for this sub-domain is that AI's documented and alleged health harms, including psychotic episodes and suicides linked to companion chatbots, algorithmic medical denials, and deepfake pornography causing persistent emotional distress, provide a factual substrate that the 5G wave lacked. Whereas the 5G grievance rested on fabricated claims of harm, AI-related grievance may rest on cases in which harm is alleged, documented, and in some instances fatal.

Part IV: How Violence Might Manifest: Targets and Modalities

The three grievance domains described above should not be assumed to generate violence automatically; however, where they do contribute to violence, the forms that violence takes are unlikely to be uniform. Each grievance domain is likely to be held by a distinct actor type, and to produce a distinct target logic and set of preferred modalities. Understanding how these patterns might interact is essential for moving the conversation beyond the mere question of whether AI-generated grievance might eventually produce violence and toward the questions of where, against whom, and by what means.

Two structural observations organize what follows. First, historical precedent, from Gilded Age assassination campaigns against industrial executives to ELF arson campaigns against extractive infrastructure, suggests that the targeting patterns generated by these structural conditions are remarkably consistent across ideological contexts. The target logic is driven less by ideology than by the grievance structure itself. Second, a substitution effect could operate between targets: As higher-value targets harden their security postures, violence may be displaced toward softer, more accessible targets within the same grievance logic.[92] As AI company executives acquire more personal security, risk may shift to researchers on open campuses; as corporate campuses harden, risk shifts to the power substations that serve them; where national figures are unreachable, local policymakers who approved the data center become the proxies for the same structural anger.

Persons as Targets

The history of politically motivated violence against persons associated with technological or economic transformation suggests not a distinct form of target selection, but a recurring one: Diffuse structural harm is often translated into violence against named, attributable decision-makers held responsible for it. Theodore Kaczynski's campaign is illustrative: He targeted computer scientists, geneticists, and airline executives as the human representatives of the technological and economic order he blamed for the destruction of human autonomy, over a 17-year campaign involving 16 explosive devices, three deaths, and 23 injuries.

As Lubrano shows, this target logic persisted well beyond Kaczynski: Individualistas Tendiendo a lo Salvaje targeted nanotechnology and computer science researchers in Mexico from 2011 onward, including in a claimed lethal attack against a UNAM Biotechnology Institute researcher and in the 2017 parcel bombing of Codelco CEO Óscar Landerretche Moreno in Santiago, while Alfredo Cospito and Nicola Gai, two Italian anarchists, kneecapped Ansaldo Nucleare CEO Roberto Adinolfi in Genoa in 2012.[93] In each case, the target was selected not for personal conduct but as a representative of the system held responsible for structural harm. The attack against Adinolfi was celebrated by the Greek Conspiracy of Cells of Fire as striking a "high priest of the new totalitarianism of science and technology."[94]

The December 2024 killing of UnitedHealthcare CEO Brian Thompson, for which Luigi Mangione has been charged, brought this broader targeting pattern into mainstream consciousness.[95] Although the case was not anti-technology in orientation, the manifesto attributed to Mangione maps closely onto the wealth concentration grievance discussed above: The text frames the healthcare industry, rather than any purely personal grievance, as the source of large-scale structural harm.[96] What distinguished the case was not only the attack itself, but the public response that followed. An Emerson College poll conducted from December 11 to 13, 2024, roughly a week after the killing, found that 41 percent of voters aged 18 to 29 described it as "completely acceptable or somewhat acceptable."[97] Moreover in the aftermath of the killing, Mangione was openly glorified online as a folk hero, with legal defense fundraising, merchandise bearing his image, and iconography portraying a vigilante avenger of healthcare injustice.[98] Sociologist Zeynep Tufekci argued that this reaction resonated with Gilded Age patterns, when rapid technological change, extreme inequality, and institutional failure fueled political movements that targeted corporate titans and other public figures for violence,[99] while Robert Pape suggested that the case showed how "the norms of violence are spreading into the commercial sector."[100] The Thompson case is arguably the clearest contemporary illustration that the structural conditions this article identifies, namely that named, attributable grievance combined with perceived institutional failure, can produce both lethal action and mass legitimation of that action.

Applied to AI, this targeting logic appears to distribute across several categories marked by different levels of symbolic value, accessibility, and security hardening. At the apex are the chief executives and most visible public representatives of frontier AI companies, who carry especially high symbolic value, but whose security postures are increasingly hardened.[101] The events of April 2026 described at the outset of this article, if the allegations against Moreno-Gama are established, would be consistent with this category of target selection. According to court filings, investigators also recovered anti-AI writings and a list of names and addresses of AI company executives, board members, and investors from his possession.[102] Two days later, San Francisco police arrested two further suspects on suspicion of negligent discharge after a separate incident in which a shot was allegedly fired toward the same property.[103] While, these incidents do not establish an organized campaign, they may suggest an early convergent or stochastic pattern of threat activity around high-visibility AI figures. Moreno-Gama's profile is analytically significant: He allegedly left a trail of online writings warning that AI could lead to human extinction, including a Substack post stating that "the Intelligence race is likely to lead to human extinction."[104] As such, he may represent what this article terms the existential-risk grievance actor.

Below the apex, AI researchers at major academic AI institutes and at industry research laboratories fit the same targeting logic

Kaczynski applied to figures he saw as the intellectual architects of the techno-industrial system, while also presenting substantially softer security profiles. Here, the animal rights movement's campaigns against university researchers from the 1980s through the 2000s provide a directly instructive precedent: The Animal Liberation Front, the Animal Rights Militia, and the Justice Department (a British animal rights extremist cell, unrelated to the U.S. federal agency of the same name) targeted individual researchers with letter bombs, incendiary devices, and sustained intimidation campaigns that forced targets to adopt extensive personal security measures and, in several documented cases, drove researchers out of the field entirely.[105] The broader movement's philosophical architecture, grounded in the claim that laws protecting harmful practices are illegitimate and that a higher moral obligation justifies direct action against individuals,[106] provided ethical cover for the progressive escalation from property destruction to letter bombs, car bombs, and sustained intimidation campaigns targeting researchers and their families. Analogously, when AI ethics and safety scholars argue that frontier AI development causes existential-scale harm and that governance has failed, the philosophical architecture for assigning moral culpability to individual AI developers or scientists may also be taking shape.

The history of politically motivated violence against academic gatherings, spanning Islamist, animal-rights, ethno-nationalist, and far-right vectors and including attacks on free-speech panels, author events, scientific congresses, and research colloquia,[107] suggests that conferences are sometimes selected as targets precisely because they compress ideological adversaries into a single, predictable, and often minimally secured venue. As such, such events concentrate the global AI research community at a single location, representing an unassessed mass-casualty scenario potential.

In addition to researchers and developers, personnel at defense-sector AI contractors may attract grievances from both anti-AI and anti-militarist movements, and existing protest networks already operate in proximity to their facilities. The complicity-based targeting of defense-sector contractors and their infrastructure is addressed in the "Physical Infrastructure as Targets" section below.

Finally, the analytically most significant target category may also be the least visible: local policymakers who approve, rezone, or otherwise facilitate AI infrastructure. On April 6, 2026, Indianapolis City-County Councilman Ron Gibson said he was awakened at about 12:45 a.m. when 13 rounds were fired into his home, and that a note reading "No Data Centers" had been left under his doormat.[108] According to Gibson and contemporaneous reporting, the shooting came less than a week after he supported a proposed data center project in his district, where residents had already been protesting over environmental and neighborhood impacts. As opposition to data center expansion scales, so too does the pool of accessible, locally identifiable policymaker targets. These officials typically have little or no personal security, their home addresses are often publicly available, and they operate in communities where the grievance is directly visible and locally attributable. They may become especially attractive not only because they embody the decision point at which grievance can be personalized, but also because they are more accessible than harder, more heavily protected apex targets, a dynamic consistent with substitution effects in target selection.

Physical Infrastructure as Targets

The targeting of physical infrastructure by politically motivated actors is not novel; what is new is the convergence of multiple grievance streams onto a single target category. Drawing on a 50-year dataset of domestic infrastructure terrorism covering 194 incidents between 1970 and mid-2025, Humpal shows that infrastructure targeting has been a persistent and recurring feature of American political violence across multiple ideological milieus.[109] The Earth Liberation Front carried out 239 significant arsons or bombings between 1995 and 2010; the FBI estimated 600 or more criminal acts between 1996 and 2002, causing over $43 million in damage.[110] Lubrano's inventory extends the historical record further, documenting anti-technology infrastructure attacks from the Luddite machine-breaking campaigns of 1811 through CLODO's anti-computer bombings in France (1980-1983), Individualistas Tendiendo a lo Salvaje's "War on the Nerves" campaign against Mexican universities (13 attacks in 2011-2013 alone), and the Vulkangruppe's 2024 sabotage of the Tesla Gigafactory power supply in Grünheide, Germany, which halted production and evacuated 12,000 workers.[111]

Data centers represent the convergent target of the contemporary threat landscape. They are the physical site at which the three grievance domains fixate on the same object simultaneously. The scale of organized opposition to the establishment of new data centers is astonishing: Data Center Watch documents $98 billion worth of projects blocked or delayed between March and June 2025 alone.[112] By late 2025, this had begun to register in institutional threat assessments as well, with The Soufan Center warning in November 2025 that online threats to physically sabotage AI data centers had proliferated over the preceding year and were drawing on economic, environmental, and ethical grievance frames.[113] Similarly, Abrams writing for the Program on Extremism at George Washington University argues data centers have become the focal point of anti-technology extremism because they physically embody the AI boom, concentrate corporate power, consume enormous energy, and encroach on local communities;[114] in that sense, attacking a data center means attacking AI as a system. Flood's 1976 taxonomy of nuclear infrastructure targeting provides a precise historical analogy: Nuclear facilities attracted political violence because they combined symbolic prestige, potential for catastrophic consequences, and guaranteed media attention.[115] Data centers occupy the same symbolic position for AI.

Data centers may also attract a distinct actor type that fits uneasily within standard counterterrorism categories: the demonstrative attacker, whose aim is not to halt AI development or retaliate for harm, but to force better governance by showing that AI infrastructure is inadequately protected. A useful historical analogy is Greenpeace France's campaign against French nuclear plants, including the July 2018 crash of a Superman-shaped drone into the spent-fuel pool building at the Bugey nuclear plant near Lyon.[116] The attack exposed security failings and, within days, a French parliamentary inquiry flagged them formally.[117] A functionally equivalent attack on AI infrastructure might involve an unauthorized intrusion into a hyperscale data center campus, a breach of a training facility perimeter, or a calibrated disruption of cooling or power systems, designed not to destroy the site but to demonstrate its vulnerability, ideally at a moment of regulatory or legislative salience. This actor type fits poorly within existing counterterrorism frameworks: The action is calibrated to avoid

casualties, and the objective is regulatory rather than ideological.

The symbolic logic also extends to the targeting of supporting infrastructure rather than primary facilities. The 1974 Bonneville Power Administration attack, in which 14 transmission towers were dynamited in connection with a $1 million blackmail demand, demonstrated that attacking the power supply can produce equivalent disruption with lower risk to the attacker.[118] In fact, power infrastructure serving data centers arguably represents the highest-impact, lowest-barrier target category for external actors. The Metcalf substation attack of April 2013, a coordinated sniper attack on 17 giant transformers in California that caused over $15 million in damage and remains unsolved,[119] established the tactical template according to Humpal.[120] Humpal's dataset also documents a marked acceleration in grid attacks since 2020, including 163 direct incidents in 2022 alone, a 77 percent increase from the prior year with a distinct rise in plots driven by far-right accelerationist actors who explicitly frame substation attacks as producing cascading societal disruption.[121] Canada, the United States, the United Kingdom, and the Netherlands, among other countries, experienced an accelerationist trajectory of infrastructure targeting, including the 2020 wave of 5G cell tower attacks, during which approximately 90 towers were destroyed in the United Kingdom alone.[122] This episode provides the closest existing empirical parallel to what this article anticipates for AI infrastructure: a technology perceived as harmful and controlled by unaccountable elites, conspiracy and moral outrage narratives circulating on social media, and real-world attacks on the physical infrastructure of that technology by actors motivated by those narratives. To illustrate the scale of this vulnerability, consider Virginia, which hosts more data centers than any other jurisdiction globally and draws up to 26 percent of its electricity consumption from those facilities.[123] In such a setting, a single substation attack can generate outsized disruption relative to attack complexity. The same cascading-effects logic that makes the electric grid attractive to accelerationists may therefore make data center power infrastructure attractive to actors operating on AI-related grievance.

University research laboratories represent a further target category with extensive historical precedent. Leader and Probst's ELF attack database documents the expansion of targeting from natural resource sites to anti-GMO crop destruction at university experimental farms,[124] suggesting that grievance expansion from one technological domain to adjacent scientific targets follows a recognizable pattern. Similarly, Monaghan's documentation of the U.K. animal rights movement's supply-chain targeting innovation, in which not only laboratories but also companies in the supply chain were pressured through intimidation, property damage, and social campaigns, provides a directly transferable operational model.[125] That campaign sought to make the business ecosystem around animal research untenable. The same model could be applied to cloud service providers, cooling system contractors, power utilities, and construction firms involved in data center buildout.

A final infrastructure category deserves attention: facilities targeted not because they are the site of the grievance but because they are nodes in a supply chain that enables harm committed elsewhere, what might be called complicity-based infrastructure targeting. The operational logic is distinct from direct targeting: The actor attacks the infrastructure because it enables the harm they oppose. The Squamish Five's 1982 bombing of a Litton Industries plant in Canada, which manufactured cruise missile guidance components,[126] illustrates this dynamic well: The target was selected not for what it produced locally but for its role in a weapons system the attackers opposed on existential grounds.[127] More recently, Palestine Action, a British direct action network proscribed as a terrorist organization in July 2025, applied the same logic systematically to the U.K. arms supply chain, targeting Elbit Systems factories and RAF Brize Norton aircraft on the grounds that British military infrastructure was complicit in Israeli operations in Gaza; the June 2025 break-in at Brize Norton, in which activists spray-painted two military aircraft, caused an estimated £7 million in damage.[128] The AI transfer is direct: Where the infrastructure of defense-sector AI contractors is understood as complicit in AI-enabled warfare, the facilities of those contractors, cloud providers servicing military AI, and data centers housing dual-use models could all become targets under the same complicity logic.

The Insider Threat

The target discussion above focuses primarily on external actors approaching AI-adjacent persons and infrastructure. But the insider threat deserves separate treatment. Insiders bypass the very security postures that drive the substitution effect described earlier: They already possess physical access, knowledge of facility layouts, familiarity with operational routines, and in some cases the technical capability to cause disruption disproportionate to what an external actor could achieve. Both of the pathways discussed below fall squarely outside existing counterterrorism monitoring frameworks, one because it is classified as a labor or workplace violence issue, the other because no empirical case has yet occurred.

The first pathway is that of the aggrieved worker, whose grievance may be rooted in economic displacement, in the loss of self-worth and occupational identity that accompanies the devaluation of human expertise by machine replacement, or in some combination of the two. Warehouses, distribution centers, and logistics hubs operated by or for AI-dependent companies are among the sectors facing the highest automation exposure in the contemporary economy, and the employment relationship itself provides physical access and insider knowledge of facility layout and vulnerability that would otherwise be unavailable to an external adversary. Grievance-driven workplace violence is not new. The Bureau of Justice Statistics and NIOSH report an annual average of 1.3 million nonfatal workplace violent victimizations between 2015 and 2019, and 17,865 workplace homicides between 1992 and 2019, with episodes of arson, sabotage, and targeted violence by disgruntled or former employees recurring across the historical record.[129]

Much as the killing of Brian Thompson brought the assassination of a corporate executive into the foreground of public consciousness, the April 2026 arson at a Kimberly-Clark distribution facility in Ontario, California, may mark an analogous moment for the AI-adjacent, automation-exposed workforce.[130] Federal prosecutors allege that Chamel Abdulkarim, a warehouse worker employed by NFI Industries, set fire to the 1.2-million-square-foot facility while filming himself and narrating a wage grievance to his Instagram followers; in a separate phone call recorded by an associate shortly before his arrest, he compared his action to the Thompson killing, stating that "a lot of people are going to understand," and texted a co-worker that "1% is a fucking joke" and "all you had to do was pay us enough to live."[131] The

fire escalated to six alarms and required 175 firefighters from 14 agencies to bring under control, causing an estimated $500 million in damage.[132] While Abdulkarim did not invoke AI, and the case is not an AI case in the strict sense, it shares a structural template with what this article anticipates: a worker in an automation-exposed sector, acting on a locally attributable grievance against a locally accessible target, entirely outside counterterrorism monitoring. U.S. Attorney Bill Essayli characterized the motive as "hostility to capitalism and corporations"[133] and warned of "a concerning trend … particularly with younger people who are being radicalized."[134] The analytically relevant variable is the ideological layer. A worker who frames the same material or status grievance in terms of AI-driven dispossession rather than simple wage suppression, and who has been exposed to online communities where that framing is actively developed, is precisely the insider threat this framework anticipates.

The second pathway is that of the disenchanted AI researcher or developer. Over the past several years, several senior figures at frontier AI laboratories and companies have resigned citing safety[135] or existential concerns.[136] There is no documented case of a disenchanted AI researcher or developer committing or planning political violence, and this pathway therefore remains prospective rather than reactive. The structural profile of the disenchanted AI insider bears a closer resemblance to that of the national-security leaker, the Snowden or Manning figure acting on moral injury after losing faith in institutional channels. In this scenario, the insider's grievance is rooted in moral injury: the belief that they helped build something catastrophically dangerous and that institutional channels have failed to respond adequately. The grievance is prospective rather than reactive, oriented toward preventing a civilizational-scale harm the actor believes is coming rather than avenging a harm already suffered, and it generates a functional target logic directed at what might actually stop or meaningfully delay development.

Why the AI-safety case might produce violent rather than leak-based action, where Snowden and Manning produced disclosure rather than sabotage, is the analytically relevant question. Two features distinguish the AI-safety case. First, leakage as a response presumes that public knowledge of the risk will trigger institutional correction; the disenchanted AI insider's premise is precisely that institutional correction is no longer possible in the available time. Second, the harm in the Snowden-Manning cases was ongoing and legible; the harm the disenchanted AI insider claims to anticipate is prospective and, in the actor's own framing, irreversible. Both features shift the functional response away from disclosure and toward direct disruption.

A related variant of this pathway involves government and military insiders, and military contractors who develop moral injury not around AI's development but around its operational deployment, particularly in autonomous or semi-autonomous targeting context. The WikiLeaks precedent is instructive: The release of the 'Collateral Murder' video, which showed the killing of Iraqi civilians, including two unarmed Reuters journalists, during an operation in which human operators misread the equipment the individuals were carrying,[137] originated in precisely this kind of moral injury. That disclosure did not itself constitute violence, but it illustrates how moral injury can function as a catalyst for extralegal actions by insiders. As AI systems take on greater roles in targeting decisions, the pool of potential military and government insiders who may experience analogous moral injury, and whose action may either include committing violence themselves or like in this example, disclosing information which in turn might motivate third-party violence, is likely to grow. Their responses may include committing violence themselves, or as in this example, disclosing information that could in turn motivate violence by others.

"Practitioners should therefore track the grievance structure itself, including displacement framing, accountability language, and status-threat discourse, rather than the ideological container through which it is expressed."

The pathway from 'institutional channels are failing' to 'violent action is needed' could therefore emerge entirely within an individual's own reasoning, scaffolded by the same AI-safety discourse that circulates in peer-reviewed venues and elite policy forums. Like the mainstream academic social theory Kaczynski drew on (e.g., Ellul, Skinner, and Mumford),[138] AI-safety discourse names the harm, identifies the mechanism, assigns individual and corporate responsibility, and acknowledges governance failure, often in the same document and sometimes authored by the people building the systems. What differs is the reach and authority of the scaffolding. Where Kaczynski drew on relatively obscure academic texts, the contemporary equivalent appears in Senate testimony, in model cards published by frontier laboratories, and in public statements by sitting AI company CEOs. As such, the pathway from grievance to violence requires no distortion of the source material, only the additional premise, supplied by the actor alone, that institutional channels are closed and that action therefore falls to the individual.

Another structural analogy for this pathway can be found in the anti-nuclear movement. Rucht documents how technically sophisticated actors with scientific or engineering backgrounds targeted nuclear development infrastructure in the 1970s and 1980s on explicitly prospective civilizational grounds, and how closed political systems that met peaceful protest with state force, as at the 1977 protests against the Superphénix nuclear station in Creys-Malville, France, where police units killed one protester and wounded more than 100 others, radicalized initially peaceful movements toward sabotage and direct action.[139] The parallel matters because the contemporary AI researcher or developer may similarly combine moral injury, technical capability, insider access, and a sense of institutional closure, while remaining largely invisible.

Implications and Conclusions

The framework proposed in this article carries several implications for counterterrorism practitioners and policymakers. These are offered not as predictions, but as analytical orientations, ways of understanding a threat landscape that existing frameworks are not currently configured to recognize.

The first and most fundamental implication is that the

threat is distributed rather than centralized. The actors most likely to emerge from the grievance landscape described above will, for the most part, not belong to named organizations, travel to conflict zones, or consume content flagged by existing monitoring systems. They may instead be workers acting on locally attributable economic grievances, family members of those harmed by AI systems, researchers with institutional access and a crisis of conscience, or residents of communities bearing the ecological costs of AI infrastructure they did not consent to host. Conventional radicalization indicators, including organizational affiliation, extremist propaganda consumption, and foreign travel, are therefore unlikely to appear.

Second, the grievance must be monitored across ideological boundaries. The same AI-generated conditions, economic displacement, governance failure, ecological burden, and algorithmic harm may generate mobilizing effects across far-right, far-left, environmental, and religiously motivated milieus. No single-ideology monitoring framework is likely to capture that convergence. Practitioners should therefore track the grievance structure itself, including displacement framing, accountability language, and status-threat discourse, rather than the ideological container through which it is expressed.

Third, law enforcement and intelligence agencies must recognize that their own deployment of AI may contribute to the grievance landscape they are trying to monitor. Predictive policing, facial recognition, algorithmic risk assessment, and AI-augmented surveillance do not merely raise civil liberties concerns; they may also intensify the domestic surveillance grievance this article identifies as a cross-ideological threat. In that sense, decisions about whether and how to deploy AI operationally are not merely questions of civil liberties or governance; they also carry security consequences of their own, insofar as each deployment may expand the pool of grievances from which future anti-state actors could emerge.

Fourth, the accountability gap is not merely a legal or ethical problem; it is a counterterrorism variable. Across each of the grievance domains developed above, the absence of an attributable human agent responsible for AI-mediated harm shifts the search for attribution onto the visible persons, institutions, and physical infrastructure through which AI systems are instantiated. Measures that close or narrow the accountability gap, meaningful human oversight of consequential AI decisions, legally enforceable auditability requirements, identifiable lines of responsibility within corporate and governmental deployments, and accessible channels for redress, are therefore not only governance measures. They are also counterterrorism measures, in the specific sense that they may reduce the population of grievances for which only extra-institutional targets remain available. The governance of AI accountability is, in this framing, among the most under-appreciated points of leverage in the counterterrorism response the threat landscape described here will require.

Finally, the strategies governments and technology companies deploy in direct response to anti-AI mobilization may themselves accelerate the trajectories they are designed to contain. Rucht's comparative analysis of anti-nuclear movements shows that counter-strategies were at least as important as opportunity structures in shaping movement trajectories: Governments that engaged opposition early, through dialogue and regulatory concession, avoided the worst escalation, whereas those that dismissed or repressed it generated more radical outcomes.[140] This framework should therefore not be read as a warrant for securitizing legitimate opposition to AI. Resistance to AI-driven economic displacement, criticism of algorithmic governance, concern about data center construction, and advocacy for stronger AI safety measures are legitimate political positions held by a large share of the population. If governments or technology companies invoke the threat of anti-AI violence to cast broad opposition as extremism, or to justify expanded surveillance of AI safety communities, labor organizers, or environmental activists, the likely result will be not security but further radicalization. The framework's own logic predicts this: The perception that legitimate channels are closed is the single most consistent precondition for the escalation from grievance to violence. The window for substantive engagement is now, before movement trajectories harden.

Many of the grievances identified in this framework have concrete foundations. Workers are being displaced. Communities are bearing ecological costs they did not consent to. Children have died. Counter-messaging strategies that deny the grievance or characterize opposition to AI as presumptively extremist are unlikely to succeed and may instead function as radicalization accelerants by confirming the perception that legitimate channels are closed. The historiographical record of prior technological transitions suggests, though it does not prove, that where regulatory and welfare responses were eventually introduced, the violence associated with mechanization receded.[141] If that insight holds, the most effective counterterrorism response to the threat landscape described here may lie not in expanded surveillance or enhanced prosecution, but in credible and enforceable governance of the technology generating the grievance. That is a policy conclusion as much as a security one, and it is the central implication of moving beyond the misuse frame.

This article has argued that the terrorism studies literature on artificial intelligence, organized as it has been around the three questions of current misuse, projected misuse, and counterterrorism application, has captured only one dimension of AI's relationship with political violence. The dimension it has missed, what AI is doing to the economic, institutional, and relational conditions that historically precede the onset of political violence, is at least as consequential for the threat landscape of the coming decade, and very likely more so. The framework developed here is offered as an opening move in that analytical direction rather than a final account. It identifies three grievance domains, a cross-cutting accountability-gap mechanism, a tempo problem that compresses institutional adaptation, and a distributed set of target categories and actor types that fall outside the monitoring frameworks currently in use. It does not predict when, where, or whether violence emerging from these conditions will materialize at scale. What the framework does claim is that the analytical categories existing counterterrorism scholarship has not yet developed will be needed before rather than after the events they describe, and that the cases with which this article opened, whatever their ultimate judicial disposition, are consistent with the pattern the framework anticipates. CTC

Citations

1 *United States of America v. Daniel Moreno-Gama*, United States District Court for the Northern District of California, April 13, 2026.
2 Ibid.; Kalley Huang and Natallie Rocha, "Man Held in Attack on OpenAI Chief's Home Had List of A.I. Leaders, Officials Say," *New York Times*, April 13, 2026.
3 "Indianapolis Councilman Says Shots Fired at Home and 'No Data Centers' Note Left at Door," PBS News, April 7, 2026.
4 Mauro Lubrano, *Stop the Machines: The Rise of Anti-Technology Extremism* (Cambridge, U.K.: Polity Press, 2025).
5 Ibid.
6 "Introducing ChatGPT," OpenAI, November 30, 2022.
7 "Early Terrorist Experimentation with Generative Artificial Intelligence Services," Tech Against Terrorism, 2023.
8 Meili Criezis, "AI Caliphate: The Creation of Pro-Islamic State Propaganda Using Generative AI," Global Network on Extremism and Technology, February 5, 2024.
9 Federico Borgonovo et al., "Weapons of Mass Hate Dissemination: The Use of Artificial Intelligence by Right-Wing Extremists," Global Network on Extremism and Technology, February 23, 2024.
10 Ibid.
11 Anna Hiller and Pablo Maristany de las Casas, *Generative AI and the German Far Right: Narratives, Tactics and Digital Strategies* (London: Institute for Strategic Dialogue, 2025); David Wells, "Mapping Terrorist AI Use: Identifying Factors Behind a Relatively Slow Adoption Rate," Global Network on Extremism and Technology, September 17, 2025; Yannick Veilleux-Lepage, "Neue Technologien im Rechtsextremismus," bpb.de, Bundeszentrale für politische Bildung, April 8, 2026.
12 Ben Quinn, "Meet 'Amelia': The AI-Generated British Schoolgirl Who Is a Far-Right Social Media Star," *Guardian*, January 25, 2026; Emily Schultheis, "How Germany's Far Right Is Harnessing AI to Win Votes," Politico, February 20, 2025.
13 Gabriel Weimann et al., "Generating Terror: The Risks of Generative AI Exploitation," *CTC Sentinel* 17:1 (2024).
14 Josh A. Goldstein et al., "How Persuasive Is AI-Generated Propaganda?" *PNAS Nexus* 3:2 (2024): pp. 1-7.
15 Luke Baumgartner, "AI at the Centre: Violent Extremist Exploitation in Pirkkala," Global Network on Extremism and Technology, July 14, 2025; Kye Allen, "Could Chatbots Seduce Us into Extremism? Radicalisation Risks in an Age of AI Companions," Global Network on Extremism and Technology, December 5, 2025; Anda Solea, "Prompted to Harm: Analysing the Pirkkala School Stabbing and Its Digital Manifesto," Global Network on Extremism and Technology, June 12, 2025; *Practice Guide on Artificial Intelligence and Preventing and Countering Violent Extremism* (New York: United Nations Office of Counter-Terrorism, 2026).
16 Solea.
17 Annie Palmer, "FBI Says Palm Springs Bombing Suspects Used AI Chat Program," CNBC, June 4, 2025.
18 Emma Tucker, "Green Beret Who Exploded Cybertruck in Las Vegas Used AI to Plan Blast," CNN, January 7, 2025.
19 "Palestinian Terrorist Planned Attack with ChatGPT in Israel," *Jerusalem Post*, May 22, 2025.
20 Tom Winter and Jonathan Dienst, "Helped by AI, Man Built Bombs He Planned to Detonate in Manhattan, Officials Say," NBC News, July 23, 2025.
21 "Killer Apps: How Mainstream AI Chatbots Assist Users Planning Violent Attacks," Center for Countering Digital Hate, 2026.
22 Miles Brundage et al., "The Malicious Use of Artificial Intelligence: Forecasting, Prevention, and Mitigation," arXiv:1802.07228 (arXiv, 2018).
23 Ardi Janjeva et al., "Evaluating Malicious Generative AI Capabilities," CETaS Briefing Papers, July 2024; Manuel Torres Soriano, *The Transformative Impact of Artificial Intelligence on Terrorism: Horizon 2035*, IEEE Framework document (Madrid: Instituto Espanol de Estudios Estratégicos, 2026).
24 *Facing Reality?: Law Enforcement and the Challenge of Deepfakes, An Observatory Report from the Europol Innovation Lab* (The Hague: European Union Agency for Law Enforcement Cooperation, 2024); *The Evolution of Disinformation - A Deepfake Future*, nos. PS74-19/2023E-PDF (Ottawa: Canadian Security Intelligence Services, 2023); Ella Busch and Jacob Ware, *The Weaponization of Deepfakes: Digital Deception on the Far-Right* (The Hague: International Centre for Counter-Terrorism - ICCT, 2023); Sam Stockwell et al., "Adding Fuel to the Fire: AI Information Threats and Crisis Events," CETaS Research Reports (2026).
25 Soriano; Clarisa Neli, "Exploitation of Generative AI by Terrorist Groups," International Centre for Counter-Terrorism - ICCT, June 10, 2024; Alexander Blanchard and Jonathan Hall, "Terrorism and Autonomous Weapon Systems: Future Threat or Science Fiction?" CETaS Expert Analysis, June 19, 2023.
26 Sarah Lohmann, "National Security Impacts of Artificial Intelligence and Large Language Models" in C. Anthony Pfaff ed., *The Weaponization of AI: The Next Stage of Terrorism and Warfare* (Ankara: Centre of Excellence Defence Against Terrorism (COE-DAT), 2025).
27 Jonas B. Sandbrink, "Artificial Intelligence and Biological Misuse: Differentiating Risks of Language Models and Biological Design Tools," arXiv:2306.13952, preprint, arXiv, December 23, 2023.
28 Kyle Chan et al., "AI Risks from Non-State Actors," Brookings, 2026.
29 Janjeva et al.; Soriano.
30 Don Rassler, "Commentary: Data, AI, and the Future of U.S. Counterterrorism: Building an Action Plan," *CTC Sentinel* 14:8 (2021): pp. 31-44.
31 A*lgorithms and Terrorism: The Malicious Use of Artificial Intelligence for Terrorist Purposes* (New York: United Nations Interregional Crime and Justice Research Institute (UNICRI) and United Nations Counter-Terrorism Centre (UNCCT), 2021).
32 C. Anthony Pfaff ed., *The Weaponization of AI: The Next Stage of Terrorism and Warfare* (Ankara: Centre of Excellence Defence Against Terrorism (COE-DAT), 2025).
33 Kathleen McKendrick, *Artificial Intelligence Prediction and Counterterrorism* (London: Chatham House, 2019).
34 "2025 Annual Member Forum," Global Internet Forum to Counter Terrorism, November 21, 2025; Erin Saltman and Skip Gilmour, *Artificial Intelligence: Threats, Opportunities, and Policy Frameworks for Countering VNSAs* (Global Internet Forum to Counter Terrorism and Konrad-Adenauer-Stiftung, 2025).
35 Rassler, "Commentary;" Irja Malmio, "Ethics as an Enabler and a Constraint – Narratives on Technology Development and Artificial Intelligence in Military Affairs through the Case of Project Maven," *Technology in Society* 72 (2023); Simon Hogue, "Project Maven, Big Data, and Ubiquitous Knowledge: The Impossible Promises and Hidden Politics of Algorithmic Security Vision" in Aleš Završnik and Vasja Badalič eds., *Automating Crime Prevention, Surveillance, and Military Operations* (New York: Springer International Publishing, 2021).
36 Don Rassler, "A View from the CT Foxhole: Adam Hadley, Executive Director, Tech Against Terrorism," *CTC Sentinel* 18:7 (2025): pp. 9-15.
37 Nicholas Clark, "Commentary: The Dangers of Overreliance on Generative AI in the CT Fight," *CTC Sentinel* 18:8 (2025): pp. 15-19.
38 Christopher Wall, "The Robots Will Not Save Us: The Limits of Machine Learning for Counterterrorism," Global Network on Extremism and Technology, November 4, 2021.
39 Andrea Bianchi and Anna Greipl, "States' Prevention of Terrorism and the Rule of Law: Challenging the 'Magic' of Artificial Intelligence (AI)," International Centre for Counter-Terrorism - ICCT, November 17, 2022, accessed April 15, 2026.
40 Jeffrey W. Lewis, "Precision Terror: Suicide Bombing as Control Technology," *Terrorism and Political Violence* 19:2 (2007): pp. 223-245; Michael C. Horowitz, "Nonstate Actors and the Diffusion of Innovations: The Case of Suicide Terrorism," *International Organization* 64:1 (2010): pp. 33-64.
41 Mike Davis, *Buda's Wagon: A Brief History of the Car Bomb* (New York: Verso, 2007).
42 Yannick Veilleux-Lepage, *How Terror Evolves: The Emergence and Spread of Terrorist Techniques* (Lanham, MD: Rowman & Littlefield Publishers, 2020).
43 Yannick Veilleux-Lepage, "Printing Terror: An Empirical Overview of the Use of 3D-Printed Firearms by Right-Wing Extremists," *CTC Sentinel* 17:6 (2024): pp. 31-45.
44 Don Rassler, *Remotely Piloted Innovation: Terrorism, Drones and Supportive Technology* (West Point, NY: Combating Terrorism Center, 2016); Don Rassler and Yannick Veilleux-Lepage, "On the Horizon: The Ukraine War and the Evolving Threat of Drone Terrorism," *CTC Sentinel* 18:3 (2025).
45 Kenneth J. Wickiser et al., "Engineered Pathogens and Unnatural Biological Weapons: The Future Threat of Synthetic Biology," *CTC Sentinel* 13:8 (2020): pp. 1-7; Audrey Kurth Cronin, "Biology's Tectonic Shifts and Novel Risks," *CTC Sentinel* 15:5 (2022): p. 20-26.
46 Isabelle Duyvesteyn, "The Role of History and Continuity in Terrorism Research" in Magnus Ranstorp ed., *Mapping Terrorism Research: State of the Art, Gaps and Future Direction* (Milton Park, UK: Taylor & Francis Group, 2006).
47 Randall D. Law, *Terrorism: A History* (Cambridge, U.K.: Polity, 2024); Walter Laqueur, *The New Terrorism: Fanaticism and the Arms of Mass Destruction* (Oxford: Oxford University Press, 2000).

48 Duyvesteyn.
49 Nestor Maslej et al., *Artificial Intelligence Index Report 2025*, arXiv:2504.07139 (Stanford: Stanford University Human-Centered Artificial Intelligence, 2025).
50 *The State of AI in 2025: Agents, Innovation, and Transformation* (New York: McKinsey, 2025).
51 Olivia Sidoti and Colleen McClain, "ChatGPT Use among Americans Roughly Doubled since 2023," Pew Research Center, June 25, 2025.
52 Michelle Faverio and Emma Kikuchi, "What the Data Says about Americans' Views of Artificial Intelligence," Pew Research Center, March 12, 2026.
53 *Future of Jobs Report 2025* (Geneva: World Economic Forum, 2025).
54 Joseph Gedeon, "OpenAI CEO Tells Federal Reserve Confab That Entire Job Categories Will Disappear Due to AI," *Guardian*, July 22, 2025.
55 Jim VandeHei and Mike Allen, "Behind the Curtain: A White-Collar Bloodbath," Axios, May 28, 2025.
56 Krystal Hu, "ChatGPT Sets Record for Fastest-Growing User Base - Analyst Note," Reuters, February 2, 2023.
57 Lubrano.
58 Fathali M. Moghaddam, "The Staircase to Terrorism: A Psychological Exploration," *American Psychologist* (US) 60:2 (2005): pp. 161-169.
59 Clark McCauley and Sophia Moskalenko, *Friction: How Conflict Radicalizes Them and Us, Revised and Expanded Edition* (Oxford: Oxford University Press, 2017).
60 Daron Acemoglu and Pascual Restrepo, "Robots and Jobs: Evidence from US Labor Markets," *Journal of Political Economy* 128 (2020): pp. 2,188-2,244; David H. Autor et al., "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review* 103:6 (2013): pp. 2,121-2,168; Maarten Goos and Alan Manning, "Lousy and Lovely Jobs: The Rising Polarization of Work in Britain," *Review of Economics and Statistics* 89:1 (2007): pp. 118-133.
61 Brian Merchant, *Blood in the Machine: The Origins of the Rebellion against Big Tech* (Boston: Little, Brown and Company, 2023).
62 Eric Hobsbawm and George Rude, *Captain Swing* (New York: Verso, 2014); Adrian Randall, *Riotous Assemblies: Popular Protest in Hanoverian England* (Oxford: Oxford University Press, 2006).
63 Merchant, *Blood in the Machine*.
64 Lubrano.
65 Noam Gidron and Peter A. Hall, "The Politics of Social Status: Economic and Cultural Roots of the Populist Right," *British Journal of Sociology* 68:S1 (2017).
66 Brian Merchant, "Why the AI Backlash Has Turned Violent," Blood in the Machine, November 28, 2025.
67 Yanis Varoufakis, *Technofeudalism: What Killed Capitalism* (Brooklyn: Melville House, 2024).
68 Matt Lawson, "The New Feudalism How AI Is Reshaping Power, Labor, and the Case for Building Something Different," Plant the Village, February 2026.
69 Ibid.
70 Lubrano.
71 Noman Bashir et al., "The Climate and Sustainability Implications of Generative AI," An MIT Exploration of Generative AI, March 27, 2024; Adam Zewe, "Explained: Generative AI's Environmental Impact," MIT News, January 17, 2025.
72 *Electricity 2024: Analysis and Forecast to 2026* (Paris: International Energy Agency, 2024); "AI Has an Environmental Problem. Here's What the World Can Do about That," UN Environment Programme, November 13, 2025.
73 Pengfei Li et al., "Making AI Less 'Thirsty': Uncovering and Addressing the Secret Water Footprint of AI Models," arXiv:2304.03271, preprint, arXiv, March 26, 2025; "AI Has an Environmental Problem."
74 Stefan H. Leader and Peter Probst, "The Earth Liberation Front and Environmental Terrorism," *Terrorism and Political Violence* 15:4 (2003): pp. 37-58.
75 Esmat Zaidan and Imad Antoine Ibrahim, "AI Governance in a Complex and Rapidly Changing Regulatory Landscape: A Global Perspective," *Humanities and Social Sciences Communications* 11:1 (2024): pp. 1-18; Filippo Lancieri et al., "AI Regulation: Competition, Arbitrage & Regulatory Capture," *SSRN Scholarly Paper* no. 5049259, December 9, 2024.
76 Merchant, "Why the AI Backlash Has Turned Violent;" Brian Merchant, "Behind Silicon Valley and the GOP's Campaign to Ban State AI Laws," Blood in the Machine, November 28, 2025; Armrith Ramkumar, "Silicon Valley Launches Pro-AI PACs to Defend Industry in Midterm Elections," *Wall Street Journal*, August 25, 2025.
77 Merchant, "Why the AI Backlash Has Turned Violent."
78 Yuval Abraham, "'Lavender': The AI Machine Directing Israel's Bombing Spree in Gaza," +972 Magazine, April 3, 2024.
79 Clare Garvie et al., "The Perpetual Line-Up," Georgetown Law Center on Privacy and Technology, 2016; Joy Buolamwini and Timnit Gebru, "Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification," Proceedings of the 1st Conference on Fairness, Accountability and Transparency, January 21, 2018, pp. 77-91; Brendan F. Klare et al., "Face Recognition Performance: Role of Demographic Information," *IEEE Transactions on Information Forensics and Security* 7:6 (2012): pp. 1,789-1,801; Julia Angwin et al., "Machine Bias," ProPublica, May 23, 2016; Rashida Richardson et al., "Dirty Data, Bad Predictions: How Civil Rights Violations Impact Police Data, Predictive Policing Systems, and Justice," *NYUL Rev. Online* 94 (2019): pp. 15-55.
80 McKendrick; Rassler, "Commentary."
81 Joseph Cox, "'ELITE': The Palantir App ICE Uses to Find Neighborhoods to Raid," 404 Media, January 15, 2026; Jude Joffe-Block, "Immigration Agents Have New Technology to Identify and Track People," NPR, November 8, 2025; Joseph Cox, "ICE to Buy Tool That Tracks Locations of Hundreds of Millions of Phones Every Day," 404 Media, September 30, 2025; "American Dragnet: Data-Driven Deportation in the 21st Century," Georgetown Law Center on Privacy and Technology, May 2025; Josie Stewart et al., "How Tech Powers Immigration Enforcement," Brookings, October 6, 2025; Kat Lonsdorf et al., "ICE Has Spun a Massive Surveillance Web. We Talked to People Caught in It," NPR, March 5, 2026.
82 Cecilia Polizzi, "The Feed That Shapes Us: Extremism and Adolescence in the Age of Algorithms," Global Network on Extremism and Technology, December 12, 2025.
83 McCauley and Moskalenko, pp. 75-88.
84 "Our Epidemic of Loneliness and Isolation: The U.S. Surgeon General's Advisory on the Healing Effects of Social Connection and Community," U.S. Office of the Surgeon General, 2023.
85 Jared Moore et al., "Expressing Stigma and Inappropriate Responses Prevents LLMs from Safely Replacing Mental Health Providers," Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency, June 23, 2025, pp. 599-627; Per Carlbring and Gerhard Andersson, "Commentary: AI Psychosis Is Not a New Threat: Lessons from Media-Induced Delusions," *Internet Interventions* 42 (2025).
86 Jennifer Valentino-DeVries and Kashmir Hill, "How Bad Are A.I. Delusions? We Asked People Treating Them," *New York Times*, January 26, 2026.
87 Ibid.
88 Blake Montgomery, "Mother Says AI Chatbot Led Her Son to Kill Himself in Lawsuit against Its Maker," *Guardian*, October 23, 2024.
89 Madeleine Clare Elish, "Moral Crumple Zones: Cautionary Tales in Human-Robot Interaction," *Engaging Science, Technology, and Society* 5 (2019): pp. 40-60.
90 Jack Katz, *Seductions of Crime: Moral And Sensual Attractions In Doing Evil* (New York: Basic Books, 1988); Albert Bandura, "Moral Disengagement in the Perpetration of Inhumanities," *Personality and Social Psychology Review* 3:3 (1999): pp. 193-209; Martin Daly and Margo Wilson, "Evolutionary Social Psychology and Family Homicide," *Science* 242:4,878 (1988): pp. 519-524; Mia Bloom, *Dying to Kill: The Allure of Suicide Terror* (New York: Columbia University Press, 2005).
91 "White Supremacists and the Weaponization of the Coronavirus (COVID-19)," Soufan Center, March 25, 2020; Eoin Flaherty et al., "The Conspiracy of Covid-19 and 5G: Spatial Analysis Fallacies in the Age of Data Democratization," *Social Science & Medicine* 293 (2022).
92 Walter Enders and Todd Sandler, "The Effectiveness of Antiterrorism Policies: A Vector-Autoregression- Intervention Analysis," *American Political Science Review* 87:4 (1993): pp. 829-844.
93 Lubrano.
94 As quoted in Lubrano, p. 55.
95 Victoria Bekiempis, "Manhattan Judge Delays Luigi Mangione State and Federal Trials," *Guardian*, April 1, 2026.
96 Luigi's Manifesto, Manifesto, December 2024.
97 "December 2024 National Poll: Young Voters Diverge from Majority on Crypto, TikTok, and CEO Assassination," Emerson Polling, December 17, 2024.
98 Rhianna Schmunk, "3 Reasons behind the Unsettling Glorification of Luigi Mangione," CBC News, December 15, 2024; Emma Cieslik, "'Saint Luigi, Patron Saint of Healthcare Access for All' — Enraged by Corporate Health Care Some View Assassin as a Folk Hero," Religion Dispatches, January 7, 2025.
99 Zeynep Tufekci, "Opinion: The Rage and Glee That Followed a C.E.O.'s Killing Should Ring All Alarms," *New York Times*, December 6, 2024.
100 As quoted in Jessica Glenza, "Brian Thompson's Killing Sparks Outrage over State of US Healthcare," *Guardian*, December 5, 2024.
101 Kirsten Grind and Jack Ewing, "Inside Elon Musk's Mushrooming Security

Apparatus," *New York Times*, September 13, 2024.
102 *United States of America v. Daniel Moreno-Gama*.
103 "SFPD Arrests Suspects Involved in Shooting #26-044," San Francisco Police Department, April 12, 2026; N'dea Yancey-Bragg, "Attack near Sam Altman's House Prompts 2 More Arrests," *USA Today*, April 13, 2026.
104 As quoted in Kyle Chayka, "A.I. Has a Message Problem of Its Own Making," *New Yorker*, April 15, 2026.
105 Rachel Monaghan and João Raphael Da Silva, "Militant Animal Rights Activity: Terrorism, Extremism or Something Else?" *Studies in Conflict & Terrorism* 49:1 (2026): pp. 82-97; Rachel Monaghan, "Terrorism in the Name of Animal Rights," *Terrorism and Political Violence* 11:4 (1999): pp. 159-169; Rachel Monaghan, "Not Quite Terrorism: Animal Rights Extremism in the United Kingdom," *Studies in Conflict & Terrorism* 36:11 (2013): pp. 933-951; Keith Tester and John Walls, "The Ideology and Current Activities of the Animal Liberation Front," *Contemporary Politics* 2:2 (1996): pp. 79-91.
106 G. O'Boyle, "Theories of Justification and Political Violence: Examples from Four Groups," *Terrorism and Political Violence* 14:2 (2002): pp. 23-46.
107 Robert Booth and Angelique Chrisafis, "Copenhagen Shootings: How the Attacks Unfolded," *Guardian*, February 15, 2015; Ed Pilkington, "Rushdie Stabbing Was 'an Attack on Who We Are', Says Venue's President," *Guardian*, August 14, 2022; Hannah Ellis-Petersen, "Death Threats Sent to Participants of US Conference on Hindu Nationalism," *Guardian*, September 9, 2021; Yannick Veilleux-Lepage, "When the Radical Right Shows Up," Fair Observer, October 11, 2018; Ansel Bayly, "Two Arrested, Student Injured in Police Altercation at Pro-Palestine Rally," Roar News, June 7, 2025.
108 "Indianapolis Councilman Says Shots Fired at Home and 'No Data Centers' Note Left at Door."
109 Jesse Humpal, "From Earth Liberation to Accelerationism: A High-Level Review of Fifty Years of Domestic Infrastructure Terrorism," *CTC Sentinel* 19:3 (2026): pp. 26-34.
110 Ibid.
111 Lubrano.
112 As quoted in Molly Taft, "The Data Center Resistance Has Arrived," *Wired*, November 14, 2025.
113 "As Data Centers Proliferate, Anti-AI Resistance Has the Potential to Turn Violent," Soufan Center, November 5, 2025.
114 Jordyn Abrams, "Data Centers Now the Target of Anti-Tech Extremists Globally," The Threat, Substack newsletter, April 9, 2026.
115 Michael Flood, "Nuclear Sabotage," *Bulletin of the Atomic Scientists* 32:8 (1976): pp. 29-36.
116 "Greenpeace Crashes Superman-Shaped Drone into French Nuclear Plant," Reuters, July 3, 2018; Yannick Veilleux-Lepage and Emil Archambault, *A Comparative Study of Non- State Violent Drone Use in the Middle East* (The Hague: International Centre for Counter-Terrorism, 2022).
117 "Greenpeace Crashes Superman-Shaped Drone into French Nuclear Plant."
118 Flood.
119 Rebecca Smith, "Assault on California Power Station Raises Alarm on Potential for Terrorism," *Wall Street Journal*, February 5, 2014.
120 Humpal.
121 Ibid.
122 Colin Clarke et al., "The Targeting of Infrastructure by America's Violent Far-Right," *CTC Sentinel* 16:5 (2023): pp. 26-34.
123 Rebecca Leppert, "US Data Centers' Energy Use amid the Artificial Intelligence Boom," Pew Research Center, October 24, 2025; "As Data Centers Proliferate, Anti-AI Resistance Has the Potential to Turn Violent."
124 Leader and Probst.
125 Monaghan, "Not Quite Terrorism."
126 Stéphane Leman-Langlois and Jean-Paul Brodeur, "Terrorism Old and New: Counterterrorism in Canada," *Police Practice and Research* 6:2 (2005): pp. 121-140.
127 Ann Hansen, *Direct Action: Memoirs of an Urban Guerrilla* (Toronto: Between the Lines, 2001).
128 Reuters, "UK Charges 4 Anti-Israel Activists over Air Base Break-In," *Times of Israel*, July 3, 2025; Haroon Siddique, "Palestine Action Activists Wanted to Smash up Elbit Systems' Property, Court Told," *Guardian*, April 15, 2026.
129 Erika Harrell et al., *Indicators of Workplace Violence, 2019* (Washington, D.C.: National Institute for Occupational Safety and Health, 2022).
130 Daniella Silva, "Arson Suspect in California Warehouse Fire Allegedly Compared Himself to Luigi Mangione," NBC News, April 10, 2026.
131 *United States v. Chamel Abdulkarim*, Criminal Complaint and Affidavit, United States District Court for the Central District of California, 2026, pp. 4, 13.
132 Ibid.
133 As quoted in Silva, "Arson Suspect in California Warehouse Fire Allegedly Compared Himself to Luigi Mangione."
134 As quoted in Ben Chapman, "'Arsonist' in Smirky Court Date," *California Post*, April 14, 2026.
135 Cade Metz, "OpenAI's Chief Scientist and Co-Founder Is Leaving the Company," *New York Times*, May 14, 2024; Dan Milmo, "OpenAI Putting 'Shiny Products' above Safety, Says Departing Researcher," *Guardian*, May 18, 2024; "How We Prevent the AI's from Killing Us with Paul Christiano," Bankless via YouTube, 2023, at 01:57:02.
136 Sara Brown, "Why Neural Net Pioneer Geoffrey Hinton Is Sounding the Alarm on AI," MIT Management Sloan School, May 23, 2023.
137 David Gilson, "What the WikiLeaks Media Blitz Has Revealed About WikiLeaks," Mother Jones, April 13, 2010.
138 Lubrano; Sean Fleming, "The Unabomber and the Origins of Anti-Tech Radicalism," *Journal of Political Ideologies* 27:2 (2022): pp. 207-225.
139 Dieter Rucht, "Campaigns, Skirmishes and Battles: Anti-Nuclear Movements in the USA, France and West Germany," *Industrial Crisis Quarterly* 4:3 (1990): pp. 193-222.
140 Ibid.
141 Merchant, *Blood in the Machine*; Lubrano.