

TECHNICAL REPORT | MARCH 2026

AI and Bomb Plots: Distinguishing Potential Effects from Language Models

Connor A. Stewart Hunter and Luca Righetti

Abstract

Future AI systems could impact the risk of terrorist bomb plots. This technical report analyses two channels by which this risk could increase: (1) future AI could increase the technical skill of bomb plotters by providing higher quality assistance than what they could otherwise access, and (2) future AI could make plotters harder to detect. We note that detection could decline if future AI systems provide higher quality operational advice on how to avoid getting caught, or if they reduce the need for terrorists to contact others for support. A preliminary review of historical cases studies of terrorism finds precedent for the relevance of these channels. We show that a large fraction of bomb plots are foiled by law enforcement, making the detection mechanism important and potentially understudied in current AI safety evaluations. We also show parallels between mentorship that some plotters currently receive from other terrorists via the internet and several aspects of emerging AI capabilities, such as multimodality, troubleshooting, and tailored guidance. We conclude that assessments of AI risk should account not just for potentially dual-use technical knowledge, but also for consequences on the ability for threat actors to evade current law enforcement detection strategies.

This work represents the views of its authors, rather than the views of the organisation, and does not constitute legal advice. GovAI technical reports have received extensive feedback but have not gone through formal peer review.

AI and Bomb Plots: Distinguishing Potential Effects from Language Models

Connor A. Stewart Hunter¹ and Luca Righetti²

¹Work completed while a GovAI Research Assistant

²GovAI

Introduction

Law enforcement and intelligence officials have warned that AI systems could be used by terrorists in bomb plots.¹ In 2025, AI chatbots were consulted in at least three explosives incidents in the US, including two incidents in which devices detonated.² For example, in the Palm Springs bombing in May 2025, the FBI recovered the perpetrator's technical discussions with an AI chatbot about explosive mixtures to create a more powerful blast.³ British security officials have warned that AI tools may circumvent counter-terrorism efforts.⁴ The terror group ISIS has reportedly published instructions to its followers on how to use AI in ways that are difficult to trace.⁵

How concerned should society be? Terrorist bombings in developed countries result in few fatalities compared to other causes of premature death.⁶ In Western Europe, there has been about one lethal incident annually since 2000, and even fewer in the US – but *attempted* bombings are more frequent, with dozens of documented terrorist bomb plots each year

¹ Federal Bureau of Investigations, "[Director Wray's Remarks at the U.S. Military Academy at West Point](#)," March 4, 2024; Dan Sabbagh, "[Terrorists could try to exploit artificial intelligence. MI5 and FBI chiefs warn](#)," *The Guardian*, October 18, 2023.

² Las Vegas Metropolitan Police, "[LIVE: Sheriff Kevin McMahon provides new info in the explosion investigation](#)," YouTube, January 7, 2025; Tom Winter and Jonathan Dienst, "[Helped by AI, man built bombs he planned to detonate in Manhattan, officials say](#)," *NBC News*, July 23, 2025; [United States v. Michael Gann](#), U.S. District Court for the Southern District of New York, Criminal Complaint, 1:25-cr-00331-DEH, June 6, 2025; [United States v. Daniel Jongyon Park](#), U.S. District Court for the Central District of California, Criminal Complaint, 5:25-mj-00400-DUTY, June 4, 2025. See also Clara Broekert and Lucas Webber, "[AI Use in Terrorist Plots and Attacks Surges in 2025](#)," *Militant Wire* (blog), December 23, 2025, for summary of these and other incidents.

³ [United States v. Daniel Jongyon Park](#).

⁴ Mason Boycott-Owen, "[Prepare for the 'coming wave' of AI terrorism, top adviser warns UK government](#)," *Politico*, July 15, 2025.

⁵ Mason Boycott-Owen, "[ISIS teaching recruits how to use AI 'responsibly'](#)," *Politico*, February 23, 2026.

⁶ Bastian Herre, Veronika Samborska, Hannah Ritchie, and Max Roser, "[Terrorism](#)," *Our World in Data*, 2023.

across the US and Western Europe.⁷ Terrorist bombings can also spike in some years due to major geopolitical events or one terrorist attack inspiring additional plots.⁸

Fortunately, most documented bomb plots are foiled by law enforcement.⁹ The extent to which *current* AI systems have enabled explosives misuse by terrorists beyond existing tools remains unclear, but is likely limited.¹⁰ However, several measures of AI are improving rapidly¹¹ and *future* AI systems might improve execution and reduce the number of plots that are foiled. Further, by studying AI's impact on bomb plots, we may learn about how AI could affect other, rarer, and harder-to-study forms of Chemical, Biological, Radiological, Nuclear, and Explosives misuse (CBRNE).¹²

As such, this piece asks: how might AI affect terrorist bomb plots?

To answer this question we review case studies and data on terrorism. Doing so, we distinguish two ways which future AI systems could affect bomb plots:

(1) Future AI could *improve plotters' skill*. Multiple AI security evaluations have looked at how much help AI provides beyond the open internet¹³, with a particular focus on technical uplift. AI systems may be thought of as playing a similar role to how bomb plotters often currently get assistance in the form of mentorship from more experienced terrorists, either in-person

⁷ We reviewed these numbers against several datasets on terrorism covering the US and Western Europe. For example, one dataset we used was the Global Terrorism Database (GTD). We filtered for incidents in the US and Western European countries since the year 2000 in which the primary weapon was explosives. GTD data includes the number of deaths due to an incident, as well as the number of deaths among perpetrators. We deducted the latter from the former and summed all of those incidents, giving 5 lethal incidents in the US and 22 in Western Europe.

⁸ Reinier Bergema and Olivia Kearney, [Rise O Muwahhid, Wherever You May Be: An Analysis of the Democratization of the Terrorist Threat in the West](#), International Centre for Counter-Terrorism (ICCT), May 12, 2020. For the US, the Global Terrorism Database shows a much higher frequency of terrorist bombings in the 1970s. For Western Europe, as another example, there was a wave of terrorist bombings in 2016 and 2017, according to data from the Jihadi Plots in Europe Dataset.

⁹ Martha Crenshaw, Erik Dahl, and Margaret Wilson, [Comparing Foiled, Failed, Completed, and Successful Terrorist Attacks: Year 5 Final Report](#), National Consortium for the Study of Terrorism and Responses to Terrorism (START), December 2017; Erik J. Dahl, ["The Plots that Failed: Intelligence Lessons Learned from Unsuccessful Terrorist Attacks Against the United States,"](#) *Studies in Conflict & Terrorism* 34, no. 8 (2011): 621–648.

¹⁰ Luca Righetti, ["Lessons from the Palm Springs Bombing,"](#) *Previous Instructions* (blog), June 9, 2025; Christopher A. Mouton, Caleb Lucas, and Ella Guest, [The Operational Risks of AI in Large-Scale Biological Attacks: A Red-Team Approach](#), RR-A2977-1, RAND Corporation, October 16, 2023.

¹¹ METR, ["Task-Completion Time Horizons of Frontier AI Models,"](#) last modified February 20, 2026.

¹² Rebecca Hersman and Cassidy Nelson, ["The Weapons of Mass Destruction AI Security Gap,"](#) *Time*, February 12, 2026.

¹³ UK AI Security Institute, [Frontier AI Trends Report](#), Department for Science, Innovation and Technology, December 2025; OpenAI, ["Building an early warning system for LLM-aided biological threat creation,"](#) January 31, 2024; Anthropic, ["Why do we take LLMs seriously as a potential source of biorisk?"](#) September 5, 2025.

or online.¹⁴ Assistance can be technical (e.g. bomb-making chemistry) or operational (e.g. attack planning). Using datasets on terrorism, we find a significant number of terrorist bombs do not detonate. This indicates there is room for technical uplift.

(2) Future AI could make plotters *harder to detect*. This could happen in two ways: first, AI could provide better quality advice on how to avoid getting caught (that is, advice on *operational security*).¹⁵ Second, AI could provide assistance that is *less visible* to law enforcement by substituting for communication with other terrorists, reducing the plotter’s footprint. Historically, law enforcement has surveilled and infiltrated channels of communication between plotters and their mentors, helping foil bomb plots. If, instead of seeking mentorship from other terrorists in group chats or online forums, a bomb plotter in the future consults AI systems, then detection may be harder. This operational security aspect appears less discussed but may also prove an important consideration for counter-terrorism efforts. We find qualitative evidence for the importance of law enforcement detecting mentorship connections in historical cases of foiled and successful bomb plots.

We summarise our risk model in Figure 1.

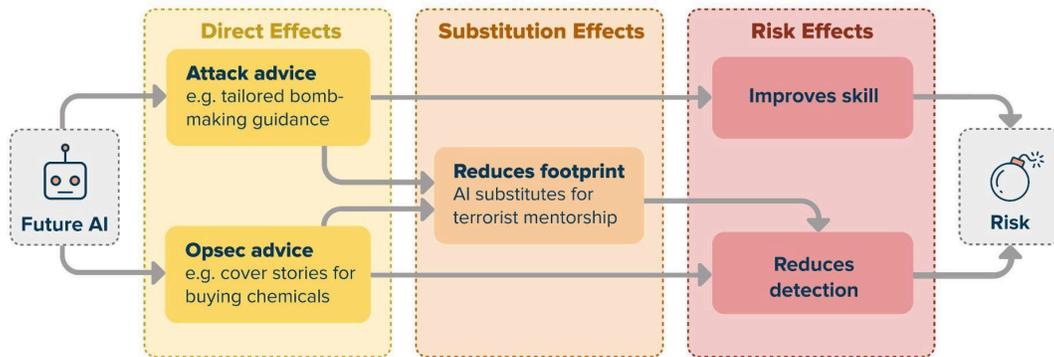


Figure 1 | Our risk model. Future AI could improve skill, but also make attackers harder to detect.

From reviewing four academic datasets on terrorist incidents (Jihadi Plots in Europe Dataset [JPED], Right-Wing Violence and Terrorism dataset [RTV], data provided by Crenshaw, Dahl, & Wilson [2017], and Global Terrorism Database [GTD]) as well as additional case studies, we show that most recorded bomb plots fail either due to lack of skill or by being detected. Mentored bomb plotters are more likely to successfully commit lethal bombings. However, mentorship between terrorists creates opportunities for law enforcement to notice and disrupt bomb plots. In the US and Western Europe, a large fraction of bomb plots are foiled by

¹⁴ AI could provide tailored instructions and troubleshoot problems in ways that are better than static resources and more user-friendly and available than virtual terrorist mentors.

¹⁵ Mouton, Lucas, and Guest, *The Operational Risks of AI in Large-Scale Biological Attacks*.

law enforcement. This happens in many ways, including because of investigative leads from mentorship-driven contact. Thus, it seems possible that future AI systems could have an effect on terrorism involving explosives, either by improving bomb plotters' skill, or by reducing their footprint or providing superior operational security advice, both making them harder to detect.

Whether future AI systems will indeed make bomb plotters harder to detect will depend not just on AI capabilities, but also on adaptation by law enforcement and AI companies. For example, some AI companies say that if a chatbot conversation repeatedly triggers flags for weapons development, they may report such extreme cases to authorities.¹⁶ How to best balance such policies alongside privacy concerns remains an open question.

The analysis presented in this report should be treated as preliminary. The datasets and case studies in this report involve small numbers of incidents: a handful of additional or different cases could meaningfully alter the patterns described. Additionally, we only explored academic datasets; law enforcement may have additional resources to obtain a more detailed picture. Our analysis is intended to help illustrate several important considerations, and we encourage further work to build on it.

Most bomb plots fail, both from lack of skill and detection

Terrorists have historically struggled to cause harm with bombs in the US and Western Europe, according to datasets on terrorism. Many bombs that terrorists build are non-functional and fail to detonate. Many more bomb plots are foiled by law enforcement before being launched. This means both a lack of technical skill and a lack of operational security are important sources of failure for bomb plots. This also means that there is room for future AI systems to meaningfully uplift threat actors on both fronts, should future AI become capable enough.

¹⁶ OpenAI, [GPT-5 System Card](#), August 13, 2025.

We illustrate these trends in Figure 2 by comparing multiple datasets on terrorism.¹⁷ Law enforcement foils a large fraction of bomb plots both in the US and Western Europe carried out by actors across multiple ideologies. Among the smaller fraction of bomb plots that are successfully launched, half or less result in harm (injury or death) to someone other than the perpetrators.

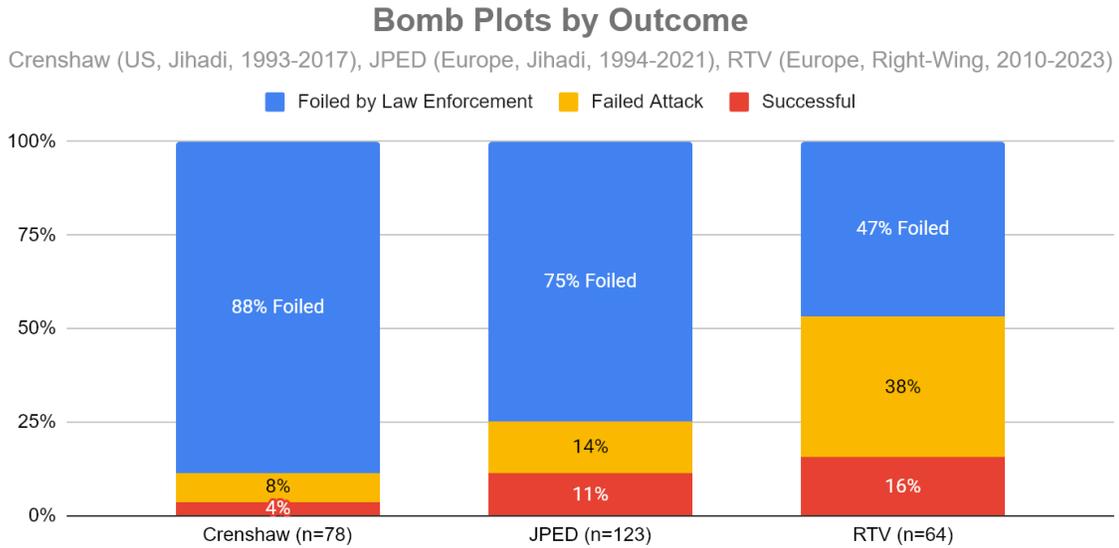


Figure 2 | Bomb Plots by Outcome. Bomb plot outcomes across three key datasets covering Jihadi and right-wing terror groups (“successful” defined as at least one non-perpetrator injury or death)

We can also see this trend of foiled and failed bomb plots when we look more granularly at how far bomb plots have historically progressed before stopping. We show this in Figure 3 using data from the Right-Wing Violence and Terrorism (RTV) dataset, which we expanded

¹⁷ To create Figure 2, we downloaded data from the Jihadi Plots in Europe Dataset (JPED), the Right-Wing Violence and Terrorism dataset (RTV), and data kindly provided by Martha Crenshaw, Dahl, & Wilson’s [Comparing Foiled, Failed, Completed, and Successful Terrorist Attacks](#) work with START. For JPED and RTV, we used variables in these datasets that represented whether a plot was foiled by law enforcement and whether the plot resulted in deaths to someone other than the perpetrator (we exclude vague or semi-vague plots, coded as plot types 4 and 5 in RTV, and plot category C3 in JPED). For Crenshaw, Dahl, & Wilson’s data, we manually recoded the dataset to map to the same information in JPED and RTV. Specifically, we used RTV’s incident type 1 to a launched attack, plot type 3 as a foiled attack, and coded an incident as a success if there was at least one non-perpetrator death or injury, and a failure if it was launched but no non-perpetrator death or injury occurred. We used JPED’s own coding for “launched” vs “foiled” attacks, and similarly counted an incident as a success if there was one non-perpetrator injury or death, and a failed attack if it was launched but did not result in a death or injury to a non-perpetrator.

upon.¹⁸ Among right-wing terrorist bomb plots in Western Europe since 2010, just over half (53%)¹⁹ are foiled by law enforcement before an attacker can plant a bomb. Among launched attacks (i.e. where an attacker has planted a bomb), many (27%) do not result in an explosion. Furthermore, even among attacks that successfully cause an explosion, most do not result in harm (other than to perpetrators), with only a single bombing since 2010 resulting in deaths to non-perpetrators. These trends are also observable in data from the Global Terrorism Database (see footnote).²⁰

¹⁸ To create Figure 3, we downloaded, filtered, and further classified data from RTV. RTV measures both violence as well as terrorism, and so we filtered for premeditated attacks (incident type 1 in RTV) and foiled plots (incident type 3 in RTV) involving explosives as either a primary or secondary weapon where the plot was not vague or semi-vague (that is, we included plot types 1, 2, and 3). We manually reviewed every entry (making a small number of corrections where RTV was found to be incorrect or incomplete), and manually determined in each case at what stage the bomb plot progressed to. We classified a plot as having progressed to a stage if there was credible public information positively demonstrating that stage (or a later stage) had been reached. Since public information does not always comprehensively speak to some aspects of a bomb plot, this coding method may understate the overall progression of plots (for example, some plots may have built a bomb, but news reporting is ambiguous and only notes the presence of chemicals relevant to making a bomb, in which case we mark the incident as “acquired precursors”, and not “built a bomb”). There is also ambiguity in much public information about what constitutes a “bomb”, with news reports sometimes referring to functional bombs as “explosives” and unfinished bombs as “bombs”. Where there was ambiguity, we used our best judgement based on the wording of public information.

¹⁹ Note that this number is slightly different from the number of foiled plots in Figure 2 due to a different method of constructing this figure, described in detail in footnote 5. In short, we did a manual review and recoding, which resulted in the addition, merging, and removal of some entries.

²⁰ Data from the Global Terrorism Database (GTD) indicates the same pattern (see [this spreadsheet](#)). The GTD tracks launched terrorist attacks across the globe as far back as 1970, and although it does not track foiled plots, it does make a distinction between launched attacks that are a “success” and launched attacks which are not. In the case of attacks involving explosives, the [GTD codebook \(2021\)](#) remarks that “in general if a bomb is planted but fails to detonate... the attack is considered for inclusion in the GTD, and marked success=0”. Using GTD data for US bombings, we find that a substantial number of bombs are planted (launched) but fail to detonate (success). GTD also codes for injuries and deaths so we can see how many detonated bombs result in harm (though note in our spreadsheet this includes perpetrator deaths).

Bomb Plot Progression

RTV (Europe, Right-Wing, 2010-2023)

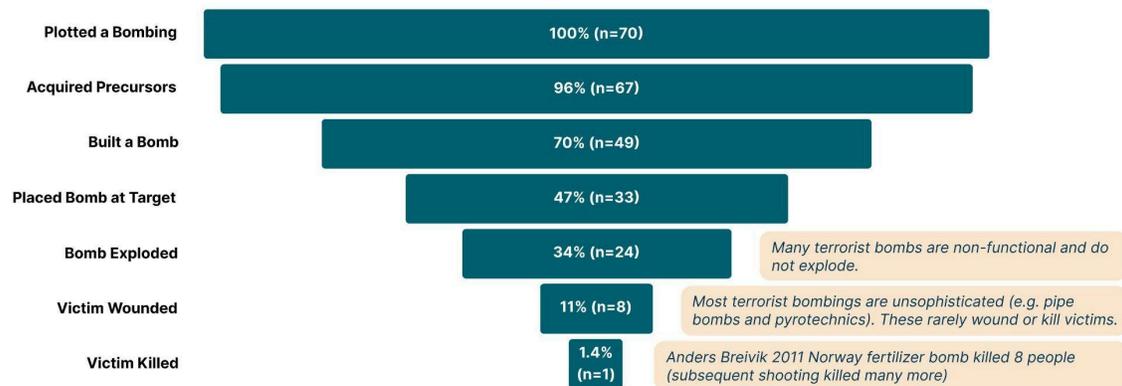


Figure 3 | Bomb Plot Progression. Author’s analysis of RTV dataset describing right-wing bomb plot progression in Western Europe, 2010-2023

These data demonstrate that bomb plots remain difficult and out of reach for many threat actors, due to both technical and operational hurdles.

Many terrorist bomb plotters seek mentorship

These technical and operational obstacles help explain why many terrorists engage mentors when formulating bomb plots. Even those described as “lone wolves” in the literature still often obtain help in procuring weapons (24%) or attempt to form a group or recruit others to their plot (35%).²¹ Although for many terrorists, the primary driver for contact may be social, in the case of bomb plots, another driver is a desire to obtain knowledge from more experienced individuals.

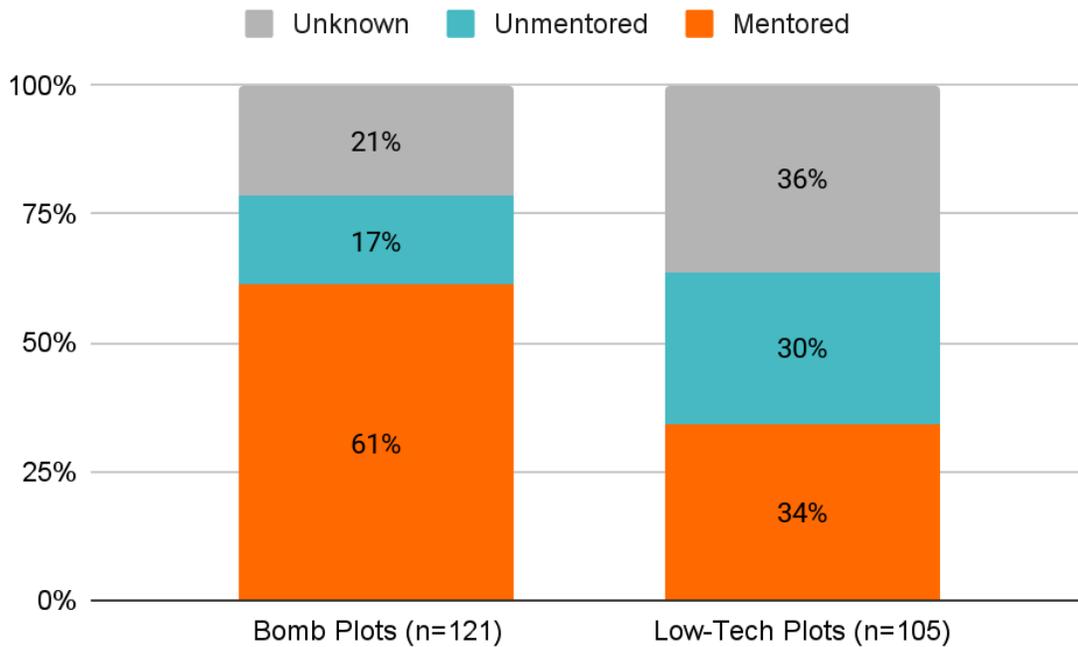
The Jihadi Plots in Europe Dataset (JPED) helpfully breaks down terrorist incidents by the type of external support given to the perpetrator. The dataset demonstrates that bomb plots often involve more external support than other plots: 61% of Jihadi bomb plots between 1994 and

²¹ Paul Gill, John Horgan, and Paige Deckert, “[Bombing Alone: Tracing the Motivations and Antecedent Behaviors of Lone-Actor Terrorists](#),” *Journal of Forensic Sciences* 59, no. 2 (December 2013): 425–435.

2021 in Western Europe involved either external training, external directives²², or both (see Figure 4).²³ By contrast, only 30% plots primarily involving firearms, bladed weapons, or vehicles (“Low-Tech Plots”) feature mentorship. These trends lend credence to the idea that terrorists seek out mentorship in part due to the technical difficulty of bomb-making and need external support to succeed.²⁴

Mentorship in Jihadi Terror Plots

JPED (Europe, Jihadi, 1994-2021)



²² In JPED, training refers to hands-on training, such as that obtained by going overseas to a conflict zone or in terrorist training camps, while directives refer to ordering an attack as well as operational and technical instructions during the preparation and/or implementation of a terror plot (see [JPED's codebook](#), and earlier typology descriptions in Thomas Hegghammer and Petter Nesser, “[Assessing the Islamic State's Commitment to Attacking the West](#),” *Perspectives on Terrorism* 9, no. 4 [August 2015]: 22). It's worth noting that this data represents support that is successfully provided and does not capture failed attempts to obtain or provide support which do not materialize.

²³ To construct Figure 4, we filtered for terrorist plots in JPED involving explosives as a primary weapon, and counted all plots known to involve either training or directives as “Mentored” and plots involving no directives or training (“remote contact but no directives” and “inspiration only”) as “Unmentored”. We compared this to terrorist plots primarily involving bladed weapons, firearms, and vehicles as “Low-Tech Plots”.

²⁴ Another complimentary explanation is that terrorist organisations invest greater resources into bomb plots. It's been argued that al-Qaeda had a preference in the decade after 9/11 for “complex attacks, often featuring improvised explosive devices and firearms” ([Council of Europe 2025](#)), which shifted to an emphasis on “attacks with any means available” in the 2010s ([Gartenstein-Ross & Blackman 2017](#)).

Figure 4 TKTK | Bomb Plot Progression. Author’s analysis of RTV dataset describing right-wing bomb plot progression in Western Europe, 2010-2023

Some terrorists explicitly state that they believe themselves incapable of building bombs unless they obtain mentorship from seasoned terrorists. We illustrate one such case study in Box 1, where a US-based threat actor indicated interest in committing terrorist attacks against the US, attempted to enter a conflict zone specifically to learn how to make “explosives and car bombs”, was unable to receive such training, and ultimately returned to the US to conduct a smaller attack using a firearm instead. If future tools, such as more capable AI systems, can substitute for mentorship, some terrorists may make use of such tools, given the lengths historically terrorists have gone to seek out practical weapons expertise.

Box 1. Threat actor seeks bomb mentorship abroad.

Little Rock shooter

Abdulahakim Mujahid Muhammad (born Carlos Bledsoe) was an American who sought explosives mentorship overseas, failed to obtain it, and ultimately carried out a lethal shooting when he returned to the US in 2009.²⁵ Muhammad moved to Yemen in 2007 and lamented while there that he lacked “knowledge and training” to carry out a suicide attack against the US.²⁶ Consequently, he attempted to enter Somalia in late 2008 “to receive training on the making of explosives and car bombs”.²⁷ Muhammed was arrested at the Yemeni border for using a fake passport and having an expired visa, spent a few months in jail, and was deported to the US in early 2009, never obtaining the training he sought. Muhammed reflected after his lethal shooting, “(H)ad I got this training my story would of ended a lot differently than it’s going to end now. My drive-by would of been a drive-in with no one escaping the aftermath!”²⁸

Terrorist mentorship does appear to help plotters

Data and case studies suggest that mentorship is associated with successful bombings. This may be because mentorship can help with technical construction and because mentorship can assist plotters in their operational security that makes them less likely to be caught by law enforcement.

²⁵ United States Secret Service, [Investigating Ideologically Inspired Violent Extremists: Local Partners Are an Asset: A Case Study on Abdulhakim Mujahid Muhammad](#), Department of Homeland Security, December 2015.

²⁶ Ibid.

²⁷ Ibid.

²⁸ Kristina Goetz, [“Muslim who shot soldier in Arkansas says he wanted to cause more death,”](#) Knox News, November 13, 2010.

Using JPED data, Figure 5 compares both bomb plots and low-tech plot outcomes²⁹ against the level of external support the plots involved. It is striking that there are no known cases where unmentored Jihadi bomb plotters caused injuries or deaths in Western Europe between 1994 and 2021. Every successful Jihadi bombing in the dataset involved some form of external support – whether training, directives, or both.³⁰

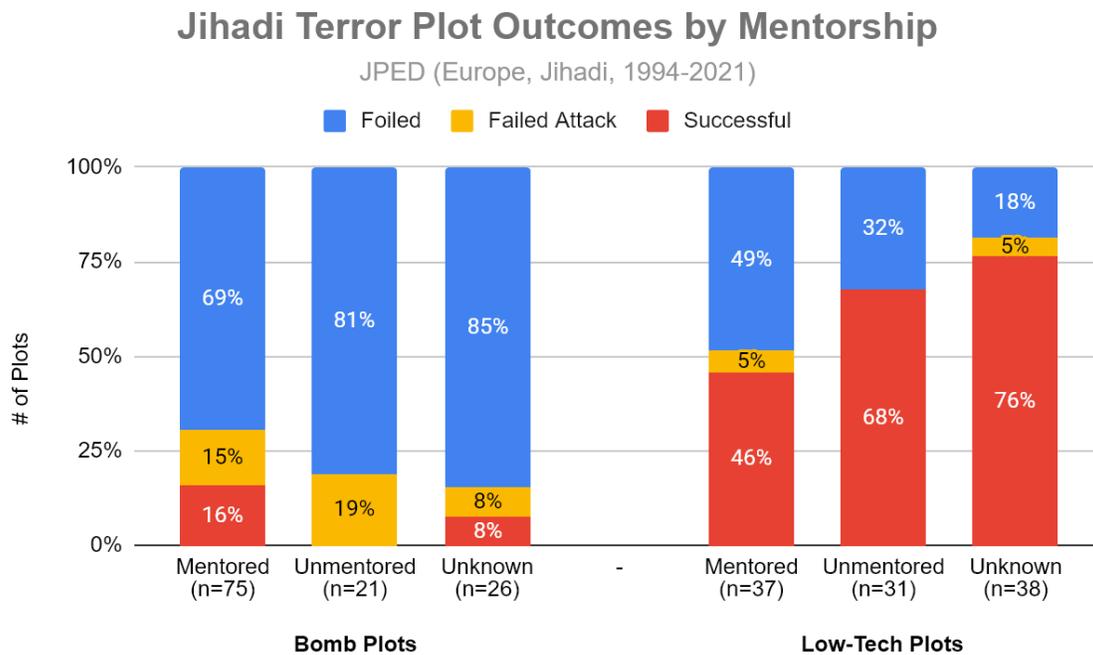


Figure 5 | Mentorship is associated with more success in bomb plots (“successful” defined as at least one non-perpetrator injury or death)

These technical and operational obstacles help explain why many terrorists engage mentors when formulating bomb plots. Even those described as “lone wolves” in the literature still often obtain help in procuring weapons (24%) or attempt to form a group or recruit others to

²⁹ JPED data tracks both the nature of external support and the outcomes of plots, which are either foiled by law enforcement, launched but fail to cause harm, or launched and successfully cause harm (“harm” defined as at least one injury or death to someone who isn’t a perpetrator of the attack).

³⁰ Data from European right-wing terrorist bomb plots offers further complications to the picture worth mentioning. The only lethal bomb plot among right-wing terrorist attacks in RTV since 2010 was carried out by a lone wolf with no known mentorship links (2011 Norway attacks). However, on the other hand, two out of three bombing *campaigns* (where multiple bombs were successfully detonated), and the only ones recorded in RTV as causing injuries, were carried out by groups (Freital Group, and Viktor Melin, Jimmy Jonasson, and Anton Thulin) – the latter group obtaining in-person training abroad. We note that this other data suggests mentorship still features prominently even for right-wing terrorism in Europe and that unmentored bomb plotters can still in rare instances execute fatal bomb plots.

their plot (35%).³¹ Although for many terrorists, the primary driver for contact may be social, in the case of bomb plots, another driver is a desire to obtain knowledge from more experienced individuals.

In JPED, mentorship is also associated with a *lower* foiling rate for bomb plots, but a *higher* foiling rate for low-tech plots. One possible explanation is that bomb plots have significantly higher operational security requirements for success than low-tech plots – so much so that the operational security benefit of mentorship outweighs the operational security penalty of contacting terrorists. Bomb plotters must acquire precursor chemicals without suspicion, synthesise chemicals without creating odours or other visible signs their neighbours might notice, and potentially test a bomb – all things that could be done with better or worse operational security. Operational security advice is known to be a major feature of external support in Jihadi plots. For example, ISIS was reported in 2015 to have set up a 24/7 help desk for “the express purpose of helping would-be jihadists use encryption and other secure communications in order to evade detection”.³² Consequently, the operational security costs of seeking out mentorship (and potentially getting on law enforcement’s radar from these connections) may be outweighed by the operational security *benefits* of mentorship in the case of bomb plots. Per this hypothesis, we see the opposite dynamic for low-tech plots because there is much less operational security advice to give (i.e. the acquisition of a car or a bladed weapon is less likely to require higher operational security skill in the first place), while still creating comparable detection opportunities for law enforcement.

However, we should be cautious in drawing too strong an inference from JPED data alone. The number of incidents is small. Even a small number of successful bomb plots without mentorship would change the picture. Additionally, there are also alternative explanations for the data. For example, terrorist groups that offer mentorship might invest more effort into bomb plots that are more promising, shifting success away from unmentored and towards mentored plotters, but not because mentorship is directly conferring an advantage.

Still, both case studies we investigated suggest that mentorship seems to confer important technical skills to plotters. In Box 2 we describe a pair of cases where separate Jihadi terrorist bomb plots were attempted in London in July 2005 by terror cells which had both received training in Pakistan. Only one of these plots resulted in functional explosive devices. That plot, which resulted in the deaths of 52 people, appears to have received continuous mentorship that “increased in frequency... when the explosive devices were being built”.³³ The less

³¹ Gill, Horgan, & Deckert, “Bombing Alone: Tracing the Motivations and Antecedent Behaviors of Lone-Actor Terrorists.”

³² Josh Meyer, “[ISIS Has Help Desk for Terrorists Staffed Around the Clock](#),” NBC News, November 16, 2015.

³³ Mitchell D. Silber, [The Al Qaeda Factor: Plots Against the West](#) (University of Pennsylvania Press, 2011), 123.

mentored plotters in the other attack attempted to detonate their bombs but failed for technical reasons.³⁴

Box 2. Mentorship a differentiating factor in two highly similar bomb plots

7/7 and 7/21

In 2005, two Islamist terror cells launched two separate terrorist bombings in London on July 7 (“7/7”) and July 21 (“7/21”). The 7/7 attack successfully detonated four bombs and killed 52 civilians, one of the deadliest terror attacks in British history.³⁵ The 7/21 attack tried but failed to detonate four bombs, causing no fatalities.³⁶ Both attacks were executed by four men.³⁷ Both terror cells pursued identical bomb designs: “the bomb-making techniques used by both groups were nearly identical and no other device of its kind had been used in the U.K.”³⁸ These designs were identical likely because both terror cells had members who visited Pakistan during overlapping periods³⁹ where they likely received training in the “production of improvised explosive devices out of hydrogen peroxide”.⁴⁰ In the leadup to their attacks, only the 7/7 attackers received continuous overseas communication. The 7/7 attackers took calls from Pakistan, which “increased in frequency in mid-May through June, when the explosive devices were being built”.⁴¹ One of the 7/7 investigators affirmed that they were “probably connected to some guidance or some means of communicating information concerned with the manufacture of the bombs”.⁴² By contrast, the 7/21 terror cell had “little, if any, ongoing communication between Muktar Ibrahim [the member who received overseas training] and Pakistan in the lead-up to the plot” according to conspirator testimony.⁴³ Although the bomb designs in both of these plots required “no great expertise” and could be learned from “open sources”,⁴⁴ the 7/21 attackers’ devices ultimately did not detonate because the hydrogen peroxide was improperly prepared,⁴⁵ a technical error the 7/7 attackers avoided.

Future AI could potentially replace virtual terrorist mentorship

Bomb plotters seek and obtain mentorship both online and in-person. When conducted online, mentorship mirrors some of the capabilities that we are starting to observe in AI

³⁴ Ibid., 142.

³⁵ Ibid., 107; UK Home Office. *Report of the Official Account of the Bombings in London on 7th July 2005*.

³⁶ Silber, 128.

³⁷ Ibid., 112, 131, 133-135. Note a fifth man was involved in but ultimately abandoned the 7/21 plot.

³⁸ Ibid., 143.

³⁹ Ibid., 140, 143.

⁴⁰ Ibid., 138.

⁴¹ Ibid., 123.

⁴² United Kingdom Coroner’s Inquests into the London Bombings of 7 July 2005, “[Hearing transcripts: 2 February 2011 - Morning Session](#),” National Archives Web Archive.

⁴³ Silber, 143.

⁴⁴ UK Home Office, [Report of the Official Account of the Bombings in London on 7th July 2005](#), May 11, 2006.

⁴⁵ Silber, 142.

chatbots. These include step-by-step and tailored guidance, troubleshooting problems, and the ability to respond multimodally (that is, to pictures as well as text).

In Box 3, we illustrate one case of virtual mentorship, where a Germany-based threat actor was unable to go to a conflict zone and instead received mentorship from a person online. This mentorship was used to plan a domestic bomb plot, including the successful creation and detonation of a test bomb before being arrested by German police.

Box 3. Threat actor mentored remotely

Cologne ricin bomb plot

In 2018, Sief Allah Hammami, a Tunisian immigrant in Germany, attempted to build a bomb laced with ricin (a poison) for a terror attack. Hammami sought and received technical and operational guidance via remote contact with Jihadis alongside assistance from his wife. Hammami had originally wanted to fight in Syria, but in 2017 his two attempts to travel there both only got as far as Turkey.⁴⁶ From December 2017, Hammami began receiving “step by step” instructions via direct messages with an online mentor on making ricin and building a bomb, sending photos of his progress and receiving feedback.⁴⁷ For example, his contact reviewed whether Hammami’s aluminum powder was fine-grained enough,⁴⁸ specified the number of cold packs (which they both mistakenly thought contained ammonium nitrate) necessary for a lethal blast radius of four to five square meters,⁴⁹ and was advised how to incorporate metal balls to make a fragmentation bomb.⁵⁰ Although the ricin part of Hammami’s plot was misinformed, executed poorly, and involved online purchases of castor beans that were flagged by British intelligence monitoring,⁵¹ Hammami was nonetheless able to build an explosive device that he tested in a park and successfully exploded.⁵² Hammami was arrested by German police in June 2018. He had been actively monitored since May 2018 due to a tip-off from British intelligence, but came to German police attention as early as December 2017 due to Hammami reporting his passport missing after his failed attempt to travel to Syria.⁵³

This case study of mentorship has parallels to emerging capabilities in AI chatbots. The Cologne bomb plotter appeared to receive mentorship predominantly over a chat interface, including step-by-step tailored guidance, troubleshooting, and use of knowledge that is difficult to communicate in static texts (i.e. tacit knowledge). We summarise aspects of virtual

⁴⁶ Flade Florian, “[The June 2018 Cologne Ricin Plot: A New Threshold in Jihadi Bio Terror](#),” *CTC Sentinel* 11, no. 7 (August 2018): 1-4.

⁴⁷ Oberlandesgericht Düsseldorf [“OLG”, Higher Regional Court Düsseldorf], [Sief Allah Hammami verdict](#), March 26, 2020, 52-57.

⁴⁸ *Ibid.*, 54.

⁴⁹ *Ibid.*, 56, 157.

⁵⁰ *Ibid.*, 61-62.

⁵¹ Flade Florian, “[The June 2018 Cologne Ricin Plot: A New Threshold in Jihadi Bio Terror](#),” *CTC Sentinel* 11, no. 7 (August 2018): 1-4.

⁵² OLG, Sief Allah Hammami verdict, 64-68.

⁵³ Florian, “[The June 2018 Cologne Ricin Plot: A New Threshold in Jihadi Bio Terror](#).”

mentorship that occurred, as well as aspects not directly assisted in the Cologne case, in Table 1.

Aspect	Cologne ricin bomb plot	
	*	Details
Chat interface	✓	It appears all guidance was text and images shared over Telegram
Tailored guidance	✓	Mentor gave step-by-step guidance through the bomb-making process.
Troubleshooting	✓	Mentor gave advice to ameliorate insufficiently coarse-grained powder.
Tacit knowledge	✓	Mentor confirmed adequacy of equipment and coarseness of a powder.
Multimodal	✓	Plotter sent images, mentor responded to them in technical detail.
Information curation	✓	Mentor sent plotter instructional video; group channels provided more.
Opsec advice	✓	Mentor advised on purchasing behaviour he thought “unsuspicious”.
Precursor acquisition	✗	Plotter travelled to Poland for fireworks, made other purchases online.
Equipment acquisition	✗	Plotter himself bought replacement coffee grinder.
Funding	✗	Plotter appears to have been self-funded.
Synthesis and assembly	✗	Plotter himself physically assembled a test bomb and remote detonator.
Testing	✗	Plotter detonated a test bomb in a meadow on his own.
Recon and rehearsal	✗	Plotter did not engage in target reconnaissance or rehearsal.

Table 1 | Mentorship aspects in the Cologne ricin bomb plot by relevance to AI. * ✓ indicates virtual mentorship assisted in this aspect in the Cologne case, ✗ indicates it did not

Although future AI would likely have to be considerably more capable to substitute for all aspects of *in-person* training, current chatbots can feasibly already provide some aspects of

virtual mentorship.⁵⁴ As a result, several consumer chatbots have built in safeguards to make eliciting dangerous knowledge in this manner difficult. There are also benchmarks for testing whether AI models produce dangerous knowledge, including ones relevant to bomb plots like the chemical security questions in WMDP – as well as additional private and government work.⁵⁵

In extreme cases, we may imagine scenarios where future AI systems may match or even outperform human mentorship. Terrorist mentors can give inaccurate information.⁵⁶ Future AI systems may make fewer of these mistakes.⁵⁷ They could also have structural advantages over human mentors, such as being always available and not becoming tired or distracted. Additionally, consulting an AI does not require reaching out to another person – which, as we show in the next section, matters for circumventing law enforcement’s ability to detect and disrupt the plot.

Contact with mentors can get terrorists caught

While current human mentorship is valuable for attackers, seeking and providing it is also risky because this contact creates opportunities for law enforcement to discover plots. As noted previously, between 47% and 88% of terrorist bomb plots are foiled according to different datasets. Data from Crenshaw, Dahl, & Wilson (2017) on all plots (bombs and non-bombs) for 1993-2017 shows that the majority of foiled plots are caught through methods that rely on plotters making contact with others. Plots can be foiled by surveillance (36% of foiled plots per Crenshaw, Dahl, & Wilson). For example, law enforcement monitoring social media activity or posting behaviour on known sites where bomb-making is discussed. Plots

⁵⁴ Even though bomb-making knowledge exists on the internet, bomb plotters have historically sought mentorship in-person or online from other terrorists. Why haven’t bomb plotters found internet resources to be sufficient? One reason for this is that although some bomb-making information is available on the open internet, much is accessible either on a darknet (Siqueira & Arce 2020) or through messaging among terrorist networks. For example, online service providers take down some content (e.g. see Europol 2021), and some terror networks exist “primarily” over social media and encrypted messaging apps (see for example the Telegram Collective DoS 2025, Farrell-Molloy 2024). As the Cologne ricin bomb plot indicates, bomb-making content (videos and manuals) is often exchanged among online terror networks. Another reason that bomb plotters seek out mentorship over and above what they can find on the internet is that mentors can provide tacit knowledge, troubleshooting, and highly tailored guidance to the specific project the bomb plotter is working on, as was the case in the Cologne ricin bomb plot.

⁵⁵ On the other hand, designing effective safeguards is complicated by the dual-use nature of explosives knowledge. Many precursors and techniques have legitimate applications in mining, demolition, agriculture, and chemistry education.

⁵⁶ See for example OLG, Sief Allah Hammami verdict, 56, where the mentor did not realize cold packs no longer contain ammonium nitrate.

⁵⁷ Future AI systems could be even more capable if developed by terrorist organisations themselves, fine-tuning AI models on tactical materials these terrorist organisations have historically shared directly with bomb plotters. If this type of future AI system comes about, it calls for separate and different adaptation by law enforcement.

can also be foiled by informants (29% of foiled plots). For example, informants and undercover operatives are known to pose as plotters and mentors. A bomb plotter seeking advice may inadvertently contact an informant or undercover operative and never successfully make contact with a bona fide terrorist. See Figure 6⁵⁸ for reproduction of Crenshaw, Dahl, & Wilson’s (2017) bomb foiling types.⁵⁹

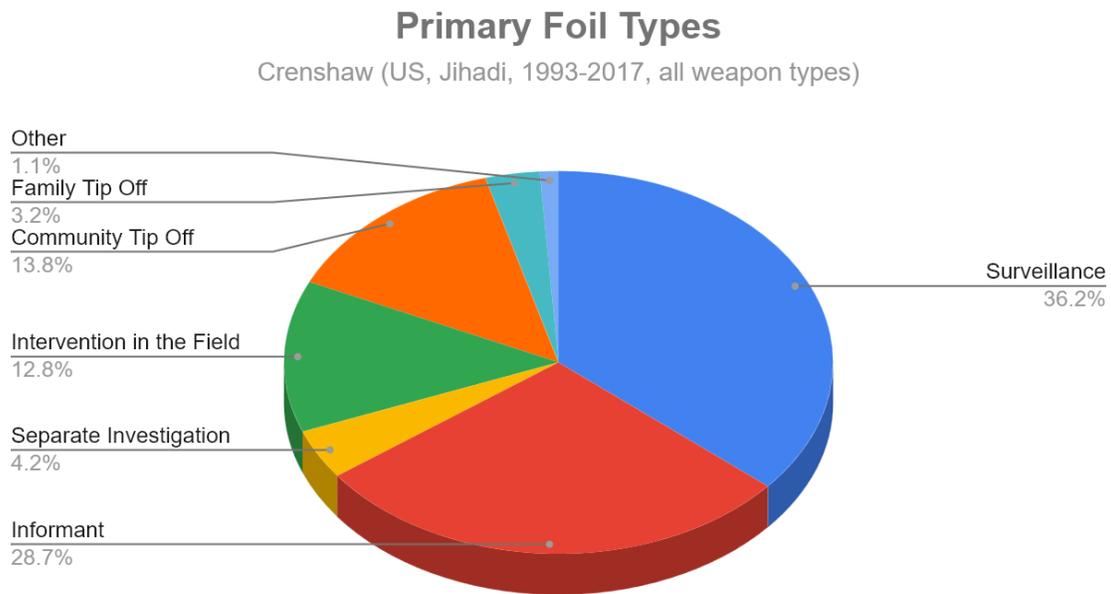


Figure 6 | Reproduced from Crenshaw, Dahl, & Wilson 2017

⁵⁸ Although we have some rough data on these proximate causes of foiling, it is harder to be sure of the underlying causes. In many cases, it appears that the underlying cause of foiling is poor operational security on the part of the plotters. For example, extremists who are the most active online are much less likely to succeed in terrorist plots partially because they are easier to detect by law enforcement. See Michael Jensen, Patrick James, Gary LaFree, Aaron Safer-Lichtenstein, and Elizabeth Yates, [Use of Social Media by U.S. Extremists](#), National Consortium for the Study of Terrorism and Responses to Terrorism (START), University of Maryland, July 2018.

⁵⁹ Crenshaw, Dahl, & Wilson, “Comparing Foiled, Failed, Completed, and Successful Terrorist Attacks”; see also Dahl, “The Plots that Failed: Intelligence Lessons Learned from Unsuccessful Terrorist Attacks Against the United States” for similar results. In the current law enforcement paradigm, some of the mainstay tools of law enforcement can be targeted and less harmful to privacy because the mentorship model of dangerous knowledge transfer structures the flow of information through opt-in terrorist group chats and terror networks. Law enforcement can therefore embed informants into high-concern networks. However, one effect of AI could be to shift the balance of tradeoffs between privacy-security by shifting the site of knowledge transfer to places that require much larger scale monitoring to achieve the same penetration of dangerous knowledge flows and same bomb plot discovery rate.

In some bomb plots, mentorship directly results in foiling. This occurs both for those seeking and providing bomb-making mentorship. In Box 4 we describe an incident in which a bomb plotter sought mentorship and came to the attention of law enforcement. This occurred first through contact with a surveilled affiliate of a terrorist organisation and, subsequently, by forming a relationship with an FBI informant.

Box 4. Threat actor seeks mentorship and is foiled

Omar Al Hardan

In 2013, Omar Al Hardan reached out to Aws Al-Jayab, a member of the al-Qaeda-affiliated al-Nusrah Front, over Facebook Messenger to discuss travel to Syria.⁶⁰ In that chat Al Hardan said “I need to learn from your weapon expertise”.⁶¹ Al-Jayab was a terrorism suspect, and Al Hardan’s contact with Al-Jayab triggered an FBI investigation into Al Hardan. Al-Jayab departed for Syria, but Al Hardan did not, possibly because he believed he needed to first acquire US citizenship.⁶² In 2014, Al Hardan developed a separate relationship with an FBI informant, again requesting training, specifically in constructing remote detonators.⁶³ The FBI informant ultimately provided one hour of training with an AK-47, and Al Hardan was arrested in 2016.⁶⁴

In Box 5 we describe an incident in which law enforcement uncovered a mentor who disseminated bomb-making advice. In this case, a mentor gave guidance on bomb-making on Facebook, triggering law enforcement to gather additional information using an informant.

Box 5. Threat actor provides mentorship and is foiled

Jarrett William Smith

Jarrett Smith was a US infantry soldier, sympathetic to the far-right Azov Battalion in Ukraine, who in 2018 and 2019 “disseminated guidance on how to construct Improvised Explosive Devices”.⁶⁵ The FBI was tipped off to Smith’s online chats and engaged him via an informant and an undercover operative,

⁶⁰ ABC13 Houston, “[Houston terror suspect reportedly wanted to blow up Galleria, Sharpstown Mall](#),” January 13, 2016.

⁶¹ [United States v. Aws Mohammed Younis Al-Jayab](#), U.S. District Court for the Eastern District of California, Criminal Complaint, 2:16-MJ-1-EFB, January 6, 2016.

⁶² Center for Immigration Studies, “[Omar Faraj Saeed Al Hardan](#),” National Security Vetting Failures Database, accessed February 2026.

⁶³ U.S. Department of Justice, [United States of America v. Omar Faraj Saeed Al Hardan](#), Indictment, January 2016.

⁶⁴ Courtney Han, “[Iraqi refugee living in Texas sentenced to 16 years in prison for plotting to make explosives for ISIS](#),” ABC News, December 19, 2017.

⁶⁵ [United States v. Jarrett William Smith](#), U.S. District Court for the District of Kansas, Criminal Complaint, 5:19-cr-40091-DDC-ADM, September 21, 2019.

learning that Smith was interested in conducting attacks within the US.⁶⁶ Smith gave the informant instructions on making an improvised explosive device.⁶⁷ Smith's technical guidance was imperfect, with FBI bomb technicians noting Smith's instructions would in some cases result in viable bombs and in other cases would not.⁶⁸ Smith was arrested in 2019 for disseminating bomb-making knowledge, admitting to officers that he "routinely" gave such advice in order to cause "chaos".⁶⁹

Future AI systems could also make bomb plotters less detectable by reducing their footprint. Both seeking and providing bomb-making mentorship is risky, because law enforcement surveils known suspects (and can pick up leads from suspicious interactions with them), monitors public online behaviour, and has informants who pose as bomb plotters and as mentors.

Even if the advice from AI systems is of the same quality as that from human mentors, substituting human contact for advice from an AI could still, in the absence of adaptation, reduce the visibility to law enforcement of information flowing to would-be bomb plotters.

Conclusion

Future AI systems could benefit terrorists aiming to carry out bomb plots in two ways: First by improving their bomb-making skill, and second, by improving their ability to evade detection by law enforcement. We distinguish these two ways in which future AI systems might impact risk by looking at public data and cases of historical terrorist bomb plots. We show that historically the most successful bomb plots involve greater mentorship (online, in-person, or over the phone), but mentorship-seeking and providing also provides an avenue for law enforcement to detect and foil some of these plots.

The analysis in this report should be cautiously considered, as the data on terrorist bomb plots is necessarily limited and our findings remain exploratory. The pathways to elevated harm described in this report – increased threat actor skill and increased operational security – represent only part of the picture. For example, our analysis does not consider radicalisation and recruitment. Our analysis also relies on academic datasets, and we note that governments may be able to draw on more information. Nonetheless, this report provides a framework for decision-makers in government and in technology firms to think about the risk posed by would-be terrorists using AI systems with rapidly advancing capabilities.

⁶⁶ Ibid.

⁶⁷ U.S. Department of Justice, "[Former Fort Riley Soldier Sentenced For Distributing Info on Napalm IEDs](#)," August 19, 2020.

⁶⁸ United States v. Jarrett William Smith.

⁶⁹ Ibid.

AI companies and risk assessments may need to adapt to address these challenges. When assessing the risk of future AI systems, assessors should focus not only on the technical capability that future AI systems may provide over and above what threat actors could get without them, but also account for operational advice and effects on threat actor operational security.

The counter-terrorism community may also need to adapt.⁷⁰ The current counter-terrorism playbook against bomb plots relies significantly on informants and surveillance of extremist networks – methods that assume plotters will reach out to other terrorists. If future AI substitutes for this mentorship, new approaches may be needed.

At the same time, AI chats could in principle be far more monitorable than current information flows, potentially allowing AI companies to detect and report misuse. How to balance public security against user privacy in this context is not obvious, and warrants broader public debate.

⁷⁰ U.S. Department of Homeland Security, [Report on the Use of Artificial Intelligence for Chemical, Biological, Radiological, and Nuclear \(CBRN\) Threats](#), April 26, 2024, 17.

About the Authors



Connor A. Stewart Hunter ✉️ 📧 [in](#)

Former Research Assistant, GovAI

Connor was a winter fellow at GovAI when he began this project and completed it as a research assistant, working on the threat modelling team. Before GovAI, Connor was a professional debate coach and cofounded a debate academy in Vancouver, Canada.



Luca Righetti ✉️ 📧 [in](#) [g](#)

Senior Research Fellow, GovAI

Luca is a Senior Research Fellow at the GovAI, where he leads a team to investigate national security risks from advanced AI systems. He previously worked at Open Philanthropy, the UK Office for AI, and the University of Oxford's Future of Humanity Institute.

Acknowledgements

We would like to thank the following individuals for feedback, review, and contributions to this technical report: Alan Chan, Thomas Clarke, Forrest Crawford, Martha Crenshaw, Erik Dahl, Mario Demetroudi, Ben Garfinkel, John Halstead, Rebecca Hersman, Aidan Homewood, Brian Jackson, Zaheed Kara, John Lidiard, Kamilė Lukošiuūtė, Sam Manning, Amelia Michael, Markus Anderljung, Grant Bradshaw, Roger Brent, Liam Patell, Jacob Aasland Ravndal, Vidar Benjamin Skretting, Bhuvana Sudarshan, Matthew van der Merwe, and Thais Warren.

About GovAI

GovAI is a 501c(3) non-profit organisation. Our mission is to help decision-makers navigate the transition to a world with advanced AI, by producing rigorous research and fostering talent. Researchers at GovAI work on a wide range of topics, with a particular emphasis on the security implications of frontier AI.

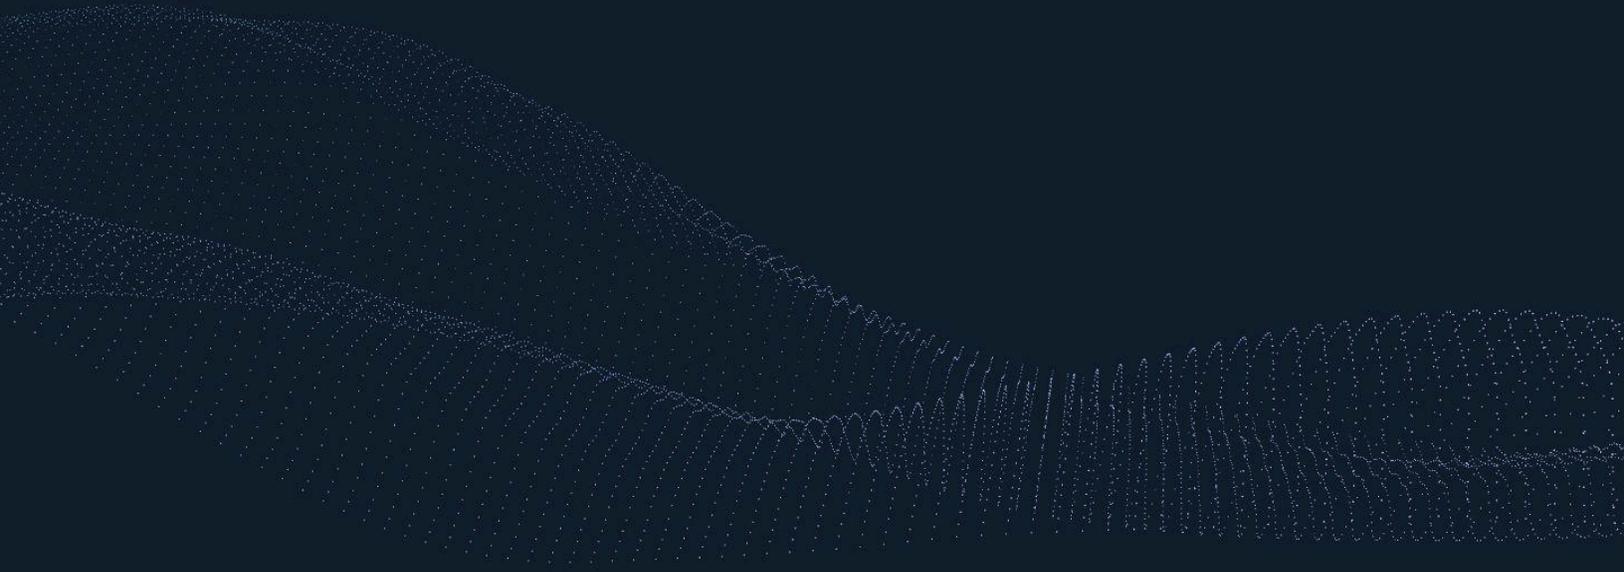
References

- ABC13 Houston. "Houston terror suspect reportedly wanted to blow up Galleria, Sharpstown Mall." January 13, 2016. <https://abc13.com/post/houston-terror-suspect-reportedly-planned-attacks-at-2-houston-malls/1157091/>.
- Anthropic. "Why do we take LLMs seriously as a potential source of biorisk?" September 5, 2025. <https://red.anthropic.com/2025/biorisk/>.
- Bergema, Reinier, and Olivia Kearney. *Rise O Muwahhid, Wherever You May Be: An Analysis of the Democratization of the Terrorist Threat in the West*. International Centre for Counter-Terrorism (ICCT). May 12, 2020. <https://icct.nl/sites/default/files/2023-02/An-Analysis-of-the-Democratisation-of-the-Terrorist-Threat-in-the-West.pdf>.
- Boycott-Owen, Mason. "ISIS teaching recruits how to use AI 'responsibly'" *Politico*, February 23, 2026, <https://www.politico.eu/article/isis-teaching-recruits-how-to-use-ai-responsibly/>.
- Boycott-Owen, Mason, "Prepare for the 'coming wave' of AI terrorism, top adviser warns UK government," *Politico*, July 15, 2025, <https://www.politico.eu/article/ai-terrorism-uk-government-propaganda-jonathan-hall/>.
- Broekert, Clara, and Lucas Webber, "AI Use in Terrorist Plots and Attacks Surges in 2025." *Militant Wire* (blog) December 23, 2025. <https://www.militantwire.com/p/ai-use-in-terrorist-plots-and-attacks>.
- Center for Artificial Intelligence Safety (CAIS). "The WMDP Benchmark: Measuring and Reducing Malicious Use with Unlearning." Accessed February 2026. <https://www.wmdp.ai/>.
- Center for Immigration Studies. "Omar Faraj Saeed Al Hardan." National Security Vetting Failures Database. Accessed February 2026. <https://cis.org/National-Security-Vetting-Failures-Database/Omar-Faraj-Saeed-Al-Hardan>.
- Council of Europe. *Report on the emerging patterns of misuse of technology by terrorist actors*. August 2025. <https://rm.coe.int/report-on-the-emerging-patterns-of-misuse-of-technology-by-terrorist-a/4880281a94>.
- Crenshaw, Martha, Erik Dahl, and Margaret Wilson. *Comparing Foiled, Failed, Completed, and Successful Terrorist Attacks: Year 5 Final Report*. National Consortium for the Study of Terrorism and Responses to Terrorism (START). December 2017. https://fsi-live.s3.us-west-1.amazonaws.com/s3fs-public/start_comparingfailedfoiledcompletedsuccessfulterroristattacks_dec2017.pdf#page=23.
- Dahl, Erik J. "The Plots that Failed: Intelligence Lessons Learned from Unsuccessful Terrorist Attacks Against the United States." *Studies in Conflict & Terrorism* 34, no. 8 (2011): 621–648. <https://doi.org/10.1080/1057610X.2011.582628>.
- Europol. "Terrorist 'how-to' guides - focus of latest Europol Referral Action Day." July 3, 2020. <https://www.europol.europa.eu/media-press/newsroom/news/terrorist-%E2%80%98how-to%E2%80%99-guides-focus-of-latest-europol-referral-action-day>.
- Farrell-Molloy, Joshua. "The proscription of Terrorgram as a terrorist organisation in the UK: Insights from the Independent Reviewer of Terrorism Legislation." *VOX-Pol*. June 26, 2024. <https://voxpoleu/the-proscription-of-terrorgram-as-a-terrorist-organisation-in-the-uk-insights-from-the-independent-reviewer-of-terrorism-legislation/>.

- Federal Bureau of Investigations (FBI). "Director Wray's Remarks at the U.S. Military Academy at West Point." March 4, 2024.
<https://www.fbi.gov/news/speeches-and-testimony/director-wrays-remarks-at-west-point>.
- Florian, Flade. "The June 2018 Cologne Ricin Plot: A New Threshold in Jihadi Bio Terror." *CTC Sentinel* 11, no. 7 (August 2018): 1-4,
<https://ctc.westpoint.edu/wp-content/uploads/2019/01/CTC-SENTINEL-082018-final.pdf#page=4>.
- Gartenstein-Ross, Daveed, and Madeleine Blackman. "ISIL's Virtual Planners: A Critical Terrorist Innovation." *War on the Rocks*, January 4, 2017.
<https://warontherocks.com/2017/01/isils-virtual-planners-a-critical-terrorist-innovation/>.
- Gill, Paul, John Horgan, and Paige Deckert. "Bombing Alone: Tracing the Motivations and Antecedent Behaviors of Lone-Actor Terrorists." *Journal of Forensic Sciences* 59, no. 2 (December 2013): 425-435.
<https://doi.org/10.1111/1556-4029.12312>.
- Goetz, Kristina. "Muslim who shot soldier in Arkansas says he wanted to cause more death." *Knox News*. November 13, 2010.
<https://archive.knoxnews.com/news/state/muslim-who-shot-soldier-in-arkansas-says-he-wanted-to-cause-more-death-ep-407169853-358338211.html/>.
- Han, Courtney. "Iraqi refugee living in Texas sentenced to 16 years in prison for plotting to make explosives for ISIS." *ABC News*. December 19, 2017.
<https://abcnews.go.com/US/iraqi-refugee-living-texas-sentenced-16-years-prison/story?id=51878408>.
- Hegghammer, Thomas, and Petter Nesser. "Assessing the Islamic State's Commitment to Attacking the West." *Perspectives on Terrorism* 9, no. 4 (August 2015): 14-30.
<https://www.jstor.org/stable/26297411>.
- Herre, Bastian, Veronika Samborska, Hannah Ritchie, and Max Roser. "Terrorism." *Our World in Data*. 2023. <https://ourworldindata.org/terrorism>.
- Hersman, Rebecca, and Cassidy Nelson. "The Weapons of Mass Destruction AI Security Gap." *Time*. February 12, 2026. <https://time.com/7373405/weapons-of-mass-destruction-ai-security-gap/>.
- Jensen, Michael, Patrick James, Gary LaFree, Aaron Safer-Lichtenstein, and Elizabeth Yates. *Use of Social Media by U.S. Extremists*. National Consortium for the Study of Terrorism and Responses to Terrorism (START), University of Maryland. July 2018.
https://www.start.umd.edu/pubs/START_PIRUS_UseOfSocialMediaByUSExtremists_ResearchBrief_July2018.pdf.
- Las Vegas Metropolitan Police. "LIVE: Sheriff Kevin McMahon provides new info in the explosion investigation." YouTube. January 7, 2025. <https://youtu.be/LnB6xVjXRmg?t=622>.
- METR. "Task-Completion Time Horizons of Frontier AI Models." Last modified February 20, 2026.
<https://metr.org/time-horizons/>.
- Meyer, Josh. "ISIS Has Help Desk for Terrorists Staffed Around the Clock." *NBC News*. November 16, 2015.
<https://www.nbcnews.com/storyline/paris-terror-attacks/isis-has-help-desk-terrorists-staffed-a-round-clock-n464391>.
- Mouton, Christopher A., Caleb Lucas, and Ella Guest. *The Operational Risks of AI in Large-Scale Biological Attacks: A Red-Team Approach*. RR-A2977-1. RAND Corporation. October 16, 2023.
https://www.rand.org/pubs/research_reports/RRA2977-1.html.

- National Consortium for the Study of Terrorism and Responses to Terrorism (START). *Global Terrorism Database Codebook: Methodology, Inclusion Criteria, and Variables*. University of Maryland. August 2021. <https://www.start.umd.edu/sites/default/files/2024-10/Codebook.pdf>.
- Nesser, Petter. "Introducing the Jihadi Plots in Europe Dataset (JPED)." *Journal of Peace Research* 61, no. 2 (2024): 317-329. <https://doi.org/10.1177/00223433221123360>.
- Oberlandesgericht Düsseldorf ["OLG", Higher Regional Court Düsseldorf]. Sief Allah Hammami verdict. March 26, 2020. https://nrwe.justiz.nrw.de/olgs/duesseldorf/j2020/6_StS_1_19_Urteil_20200326.html.
- OpenAI. "Building an early warning system for LLM-aided biological threat creation." January 31, 2024. <https://openai.com/index/building-an-early-warning-system-for-llm-aided-biological-threat-creation/>.
- OpenAI. GPT-5 System Card. August 13, 2025. <https://cdn.openai.com/gpt-5-system-card.pdf>.
- Righetti, Luca. "Lessons from the Palm Springs Bombing." *Previous Instructions* (blog). June 9, 2025. <https://substack.com/@lucarighetti/p-165530535>.
- Sabbagh, Dan. "Terrorists could try to exploit artificial intelligence, MI5 and FBI chiefs warn." *The Guardian*, October 18, 2023. <https://www.theguardian.com/technology/2023/oct/18/terrorists-exploit-artificial-intelligence-a-i-mi5-fbi-chiefs-warn>.
- Silber, Mitchell D. *The Al Qaeda Factor: Plots Against the West*. University of Pennsylvania Press, 2011. <https://muse.jhu.edu/pub/56/monograph/book/11248>.
- Siqueira, Kevin, and Daniel Arce. "Terrorist Training: Onsite or via the Internet?" *European Journal of Political Economy* 63 (June 2020): 101878. <https://doi.org/10.1016/j.ejpoleco.2020.101878>.
- UK AI Security Institute (AISI). *Frontier AI Trends Report*. Department for Science, Innovation and Technology, December 2025. <https://www.aisi.gov.uk/frontier-ai-trends-report>.
- UK Home Office. *Report of the Official Account of the Bombings in London on 7th July 2005*. May 11, 2006. <https://assets.publishing.service.gov.uk/media/5a7c7bc840f0b626628ac62e/1087.pdf#page=5>.
- United Kingdom Coroner's Inquests into the London Bombings of 7 July 2005. "Hearing transcripts: 2 February 2011 - Morning Session." National Archives Web Archive. https://webarchive.nationalarchives.gov.uk/ukgwa/20120216080106/http://7julyinquests.independent.gov.uk/hearing_transcripts/02022011am.htm.
- United States Secret Service, National Threat Assessment Center. *Investigating Ideologically Inspired Violent Extremists: Local Partners Are an Asset: A Case Study on Abdulhakim Mujahid Muhammad*. Department of Homeland Security. December 2015. <https://www.secretservice.gov/media/82/download?inline=true#page=7>.
- U.S. Department of Homeland Security. *Report on the Use of Artificial Intelligence for Chemical, Biological, Radiological, and Nuclear (CBRN) Threats*. April 26, 2024. https://www.dhs.gov/sites/default/files/2024-06/24_0620_cwmd-dhs-cbrn-ai-eo-report-04262024-public-release.pdf.
- U.S. Department of Justice. "Former Fort Riley Soldier Sentenced For Distributing Info on Napalm, IEDs." August 19, 2020. <https://www.justice.gov/usao-ks/pr/former-fort-riley-soldier-sentenced-distributing-info-napalm-ieds>.
- U.S. Department of Justice. *United States of America v. Omar Faraj Saeed Al Hardan*. Indictment. January 2016. <https://www.justice.gov/opa/file/811696/dl?inline=#page=8>.

- U.S. Department of Justice. “Iraqi Refugee Convicted of Attempting to Provide Material Support to ISIL.” U.S. Attorney’s Office, Southern District of Texas. October 17, 2016.
https://www.investigativeproject.org/documents/case_docs/3107.pdf.
- U.S. Department of State. “Terrorist Designations of the Terrorgram Collective and Three Leaders.” Office of the Spokesperson. January 2025.
<https://2021-2025.state.gov/office-of-the-spokesperson/releases/2025/01/terrorist-designations-of-the-terrorgram-collective-and-three-leaders/>.
- United States v. Daniel Jongyon Park. U.S. District Court for the Central District of California. Criminal Complaint. 5:25-mj-00400-DUTY. June 4, 2025.
https://static.foxnews.com/foxnews.com/content/uploads/2025/06/ed25mj00400duty-lodged-complaint_redacted.pdf.
- United States v. Michael Gann. U.S. District Court for the Southern District of New York. Criminal Complaint. 1:25-cr-00331-DEH. June 6, 2025.
<https://storage.courtlistener.com/recap/gov.uscourts.nysd.646218/gov.uscourts.nysd.646218.1.0.pdf#page=5>.
- United States v. Aws Mohammed Younis Al-Jayab. U.S. District Court for the Eastern District of California. Criminal Complaint. 2:16-MJ-1-EFB. January 6, 2016.
<https://www.justice.gov/opa/file/811696/dl?inline=#page=8>.
- United States v. Jarrett William Smith. U.S. District Court for the District of Kansas. Criminal Complaint. 5:19-cr-40091-DDC-ADM. September 21, 2019.
<https://extremism.gwu.edu/sites/g/files/zaxdzs5746/files/Jarrett%20William%20Smith%20Criminal%20Complaint.pdf#page=3>.
- Winter, Tom and Jonathan Dienst. “Helped by AI, man built bombs he planned to detonate in Manhattan, officials say.” NBC News. July 23, 2025.
<https://www.nbcnews.com/politics/national-security/helped-ai-man-built-bombs-planned-detonate-manhattan-officials-say-rcna220693>.



GovAI