

Generative AI and Terrorist Propaganda in South Asia

Ruchika Kakkar

October 08, 2026

Generative AI (GenAI) lowers the cost, time, and technical skill required to produce extremist propaganda. This paper examines how terrorist organizations in South Asia, most prominently Al Qaeda and Islamic State Khorasan Province (ISKP), leverage GenAI across their media and operational workflows. Groups utilize synthetic news presenters, multilingual audio, AI-generated images, and rapid repackaging of jihadist media. AI-driven translation presents the most consequential regional development, enabling seamless dissemination across Arabic, Urdu, Pashto, Persian, Hindi, Bengali, Tamil, and other local languages. Overall, the paper demonstrates that GenAI accelerates existing activities rather than fundamentally altering organizational or operational capabilities, and increased media output has not demonstrably expanded recruitment. Counter-terrorism responses should prioritize multilingual monitoring, digital forensics, platform cooperation, and targeted safeguards for AI models.



**SOCIETY FOR THE STUDY OF
PEACE & CONFLICT**

Fostering Ideas, Research and Dialogue...

Generative AI and Terrorist Propaganda in South Asia

Ruchika Kakkar

SSPC ISSUE BRIEF: October 08, 2026

ABSTRACT



Generative AI (GenAI) lowers the cost, time, and technical skill required to produce extremist propaganda. This paper examines how terrorist organizations in South Asia, most prominently Al Qaeda and Islamic State Khorasan Province (ISKP), leverage GenAI across their media and operational workflows. Groups utilize synthetic news presenters, multilingual audio, AI-generated images, and rapid repackaging of jihadist media. AI-

driven translation presents the most consequential regional development, enabling seamless dissemination across Arabic, Urdu, Pashto, Persian, Hindi, Bengali, Tamil, and other local languages. Overall, the paper demonstrates that GenAI accelerates existing activities rather than fundamentally altering organizational or operational capabilities, and increased media output has not demonstrably expanded recruitment. Counter-terrorism responses should prioritize multilingual monitoring, digital forensics, platform cooperation, and targeted safeguards for AI models.

INTRODUCTION

Generative artificial intelligence (GenAI) has become an increasingly relevant dual-use technology in counterterrorism. Terrorist organizations and their supporters, including the transnational jihadist groups such as Islamic State (IS) and its regional factions, such as IS Khorasan Province (ISKP); Al-Qaeda and its affiliates, including Al-Qaeda in the Indian Subcontinent (AQIS); and various other Islamist terrorist networks, have experimented with large language models, image generators, synthetic audio and video, automated translation, and other commercially available tools. These systems can facilitate propaganda production, adaptation of materials for different audiences, recruitment outreach, and responses to technical questions. Their accessibility has renewed concerns that AI is creating a qualitatively different terrorist threat. The available evidence supports a more limited conclusion. Generative AI is best understood as an accelerator of existing terrorist activity rather than as the foundation of a new category of "AI terrorism." It can reduce production costs, shorten the interval between an event and the release of propaganda, broaden linguistic reach, and enhance the professional appearance of content. It may also facilitate technical research. However, it does not, in itself, provide an organization, trained operatives, secure logistics, financing, weapons, or access to a target.¹

This distinction is particularly important in South Asia. Countries such as Afghanistan, Pakistan, and India already have extensive extremist media networks, encrypted communication channels, cross-border militant groups, and long-standing political and community grievances.² In this context, AI serves as an enabling tool. However, its impact depends on the organization, the user, the intended purpose, and the broader network in which it is deployed.

From Experimentation to Routine Media Use

Terrorist interest in Gen-AI became more visible after widely available text and image generators entered public use in late 2022, especially with the advent of ChatGPT in November 2022. One year later, in late 2023, Tech Against Terrorism (TAT), the UN-supported industry initiative,³ reported that its monitoring was identifying thousands of AI-generated extremist images and other artifacts each week. The material crossed ideological categories and included jihadist as well as far-

¹ The literature on this subject is growing. For comprehensive discussions, see Andrew Glazzard, David McIlhatton, and Paul Martin, "Will Generative AI Fundamentally Change Terrorist Threats?" *CTC Sentinel*, Vol. 19, No. 5 (May 2026), <https://ctc.westpoint.edu/will-generative-ai-fundamentally-change-terrorist-threats/>; Clarisa Nelu, "Exploitation of Generative AI by Terrorist Groups," *ICCT*, June 10, 2024, <https://icct.nl/publication/exploitation-generative-ai-terrorist-groups>; and Gabriel Weimann, "Generative AI and Terrorism" (20 Mar. 2025), in Dov Greenbaum (ed.), *Conflict*, in Philipp Hacker (ed.), *Oxford Intersections: AI in Society* (Oxford, online edn, Oxford Academic, 20 Mar. 2025–), <https://doi.org/10.1093/9780198945215.003.0100>.

² Bangladesh, Sri Lanka, and the Maldives also feature in regional extremist narratives and recruitment efforts, particularly in ISKP's *'Voice of Khorasan'* and AQIS's *'Nawai Gazawat-ul Hind'* digital magazines. However, these three South Asian countries are outside this paper's scope.

³ Tech Against Terrorism is an international, UN-backed initiative launched in 2017. It works with the global technology sector to help prevent terrorists and violent extremists from exploiting online services, while also protecting human rights and free speech. See, <https://www.un.org/securitycouncil/ctc/news/launch-%E2%80%99Tech-against-terrorism%E2%80%9D-%E2%80%94partnership-between-technology-companies-governments-and-un>

right content.⁴ By July 2026, its incident tracker had identified more than 30 public cases in which AI served as an operational assistant in terrorism, violent extremism, or mass-violence plots, involving at least 11 distinct AI tools and linked to more than 70 deaths, though the majority of these operational cases involved West African groups such as Boko Haram and the Islamic State West Africa Province.⁵ Early assessments nevertheless described most terrorist engagement with generative AI as experimental rather than systematic.⁶

Subsequent reporting suggests a gradual broadening of Gen-AI use. United Nations assessments have noted terrorist interest in artificial intelligence alongside virtual assets, cyber capabilities, unmanned systems and commercial communications tools. In early 2025, the UN Security Council's Analytical Support and Sanctions Monitoring Team noted that Da'esh circulated guidance on using generative AI tools to avoid detection, alongside continued interest in cyber specialists. The Secretary-General's twenty-second and twenty-third reports on the threat posed by ISIL (Da'esh), issued in February and August 2026, respectively, go further: Da'esh and its affiliates have adapted to emerging technologies, including AI, virtual assets, and unmanned aircraft systems, and have expanded their use of cyber tools and commercial satellite communications.⁷ In June 2026, UN Secretary-General António Guterres told the Fourth UN High-Level Conference on Counterterrorism that such technologies have “supercharged” terrorists' operational reach, while noting that the same tools can aid detection.⁸ These global assessments consistently describe a broadening challenge rather than a step-change in core operational capability, a distinction that matters for how South Asia should read the evidence.

These developments of course matter, but they should not be collapsed into a single measure of technological sophistication. Producing a synthetic presenter, translating a statement and using a chatbot for technical research are different activities with different operational consequences. The most useful distinction is between production, distribution, persuasion, and operational assistance. AI can clearly accelerate content production and localization. It may also help terrorist groups

⁴ David Gilbert, "Here's How Violent Extremists Are Exploiting Generative AI Tools," *Wired*, 9 November 2023, <https://www.wired.com/story/generative-ai-terrorism-content/>; “Report: Early Terrorist Adoption of Generative AI”, *Tech Against Terrorism*, November 8, 2023, <https://techagainstterrorism.org/news/early-terrorist-adoption-of-generative-ai>

⁵ Counter-Terrorism AI (CT-AI) Benchmark, *Tech Against Terrorism*, July 2026. <https://www.ct-ai.org/cases>.

⁶ “Early Terrorist Experimentation with Generative Artificial Intelligence Services,” *Tech Against Terrorism Briefing*, November 2023. <https://techagainstterrorism.org/hubfs/Tech%20Against%20Terrorism%20Briefing%20-%20Early%20terrorist%20experimentation%20with%20generative%20artificial%20intelligence%20services.pdf>

⁷ United Nations, “Joint Briefing to the Security Council on the Twenty-Third Report of the Secretary-General on the Threat Posed by ISIL (Da’esh),” S/2026/638, August 5, 2026. <https://www.un.org/securitycouncil/ctc/news/joint-briefing-security-council-twenty-third-report-secretary-general-threat-posed-isil-da%E2%80%99esh>

⁸ UN Chief António Guterres Warns AI & Emerging Technologies Are 'Supercharging' Global Terrorism Threats *Free Press Journal*, June 30, 2026, <https://www.freepressjournal.in/tech/un-chief-antnio-guterres-warns-ai-emerging-technologies-are-supercharging-global-terrorism-threats>

distribute material at scale, although platform access and account survival remain limiting factors. The evidence that AI-generated propaganda is more persuasive than existing extremist material is still limited. And there's even less evidence that it has independently enabled major attacks in South Asia.

ISKP and Synthetic News Production

Islamic State Khorasan Province (ISK) has provided the most visible regional examples of AI-assisted propaganda. After the March 2024 Crocus City Hall attack in Moscow, pro-Islamic State channels circulated an AI-generated news bulletin about the operation. Similar experimentation appeared in the Afghanistan theatre. In May 2024, a synthetic Pashto-speaking presenter announced the Islamic State's claim for the attack on foreign tourists in Bamiyan. The video imitated a conventional television bulletin and circulated under branding associated with "Khurasan TV".⁹ Researchers also identified synthetic content linked to other attacks, including a bombing in Kandahar.¹⁰

These products offered several advantages. They could be released quickly, reused in different languages (e.g Urdu, Turkish, Pashto, etc.) and made to resemble familiar television formats. They also allowed remote supporters to produce content without physical access to an established media studio. Early videos, however, contained visible defects, and their novelty did not necessarily translate into credibility. Reporting on internal debate within pro-Islamic State networks also suggests that religious and reputational concerns constrained some forms of synthetic video production.¹¹ It's helpful to distinguish between official production and supporter activity. Material shared within a pro-Islamic State ecosystem may reinforce the organization's message, even if it wasn't ordered, approved, or produced by an official jihadi media office. Supporter experimentation can show that these activities are accessible and intentional, but it shouldn't automatically be taken as evidence of a centralized organizational program.

Different Patterns across Militant Organizations

The Islamic State movement has generally shown greater interest in rapid visual production than Al-Qaeda. Since the loss of most of its territorial control, Islamic State has depended heavily on dispersed supporters and online media networks. Generative AI fits this model because it lets small

⁹ Centre for Information Resilience, "ISKP: Use of Generative AI Presenters to Create Newscasts," *Afghan Witness*, June 19, 2024, <https://www.info-res.org/afghan-witness/reports/iskp-use-of-generative-ai-presenters-to-create-newscasts/>; Mona Thakkar & Anne Speckhard, Caliphate AI: IS/ISKP Supporters Harness Generative AI for Propaganda Dissemination, *The International Center for the Study of Violent Extremism*, July 5, 2024, <https://icsve.org/caliphate-ai-is-iskp-supporters-harness-generative-ai-for-propaganda-dissemination/>

¹⁰ Sirwan Kajjo, "IS turns to artificial intelligence for advanced propaganda amid territorial defeats," *Voice of America*, 23 May 2024, <https://www.voanews.com/a/is-turns-to-artificial-intelligence-for-advanced-propaganda-amid-territorial-defeats/7624397.html>

¹¹ Fabrizio Minniti, "Automated Recruitment: Artificial Intelligence, ISKP, and Extremist Radicalization," *GNET*, April 11, 2025. <https://gnet-research.org/2025/04/11/automated-recruitment-artificial-intelligence-iskp-and-extremist-radicalisation/>

teams or individuals create images, audio, subtitles, translations, and news-style videos without the infrastructure the organization used during the height of its territorial "Caliphate".

Al-Qaeda's media culture has historically placed greater emphasis on long-form ideological texts, sermons and doctrinal instruction. Its adoption of generative AI therefore appears more cautious and less visually oriented. Al-Qaeda-linked media circles have nevertheless discussed AI-powered chatbots and their possible use for media production, anonymity and information gathering.¹² The Tehreek-e-Taliban Pakistan has also been cited in counter-terrorism reporting as using AI to improve images, videos and translation. Publicly available evidence is not yet sufficient to establish how far such practices are centrally directed or routinely used. Similarly, Baloch militant media networks in Pakistan have also circulated visually enhanced and gamified depictions of violence. Some material may incorporate AI-generated elements, but attribution is difficult when researchers lack the original files, model metadata, or production records. Visual appearance alone is not a reliable method of determining that content was generated by AI.¹³

The Translation Multiplier

AI-assisted translation is likely to have greater near-term importance in South Asia than spectacular deepfakes. Extremist organizations have long depended on multilingual volunteers to translate Arabic material into Urdu, Pashto, Persian, Bengali, Hindi and other languages. Translation tools can now produce a usable first draft within seconds. Human editors may then correct terminology, sharpen ideological language and adapt references to local grievances.¹⁴

United Nations and research reporting have specifically noted the use of AI-assisted translation within Islamic State-linked networks.¹⁵ The practical gain is organizational. A small media unit can sustain output in several languages without maintaining a large permanent translation team. Text can also be converted into speech, subtitles or short video scripts, allowing one item of propaganda to circulate in several formats.

The regional implications extend beyond Afghanistan, Pakistan and India. Bengali-language dissemination can reach audiences in Bangladesh and eastern India, while Sinhala and Dhivehi translation could lower the cost of reaching small but security-relevant audiences in Sri Lanka and the Maldives. This is a capability assessment, not evidence that Islamic State has established

¹² Sara Harmouch, "Al-Qaida and Islamic State adopting AI – but not the same way", *Asia Times*, July 28, 2026, <https://asiatimes.com/2026/07/al-qaida-and-islamic-state-adopting-ai-but-not-the-same-way/>

¹³ CTED Insight Briefing explores emerging terrorist use of generative technologies and digital ecosystems, *UN Counter-Terrorism Committee Executive Directorate*, May 28, 2026, <https://www.un.org/securitycouncil/ctc/news/cted-insight-briefing-explores-emerging-terrorist-use-generative-technologies-and-digital>

¹⁴ Alessandro Bolpagni and Eleonora Ristuccia, "AI-Powered Translation: How AI Tools Could Shape a New Frontier of IS Propaganda Dissemination," *GNET*, July 2025. <https://gnet-research.org/2025/07/09/ai-powered-translation-how-ai-tools-could-shape-a-new-frontier-of-is-propaganda-dissemination/>

¹⁵ United Nations Security Council, Thirty-eighth report of the Analytical Support and Sanctions Monitoring Team submitted pursuant to resolution 2734 (2024) concerning ISIL (Da'esh), Al-Qaida and associated individuals and entities, August 10, 2026. https://digitallibrary.un.org/record/4124567/files/S_2026_651-EN.pdf

dedicated AI operations in each country. Open-source evidence for systematic country-specific campaigns remains limited.

Language technology also creates defensive challenges. Extremist messaging frequently uses code-mixed text, transliterated Urdu or Bengali, local slang, religious references, and deliberate misspellings. Models trained mainly on standard English or formal versions of regional languages may misclassify this material. They may also remove legitimate religious or political speech if they misunderstand the context.

The Red Fort Case in India

The investigation into the 10 November 2025 vehicle-borne improvised explosive device detonation near Delhi's Red Fort provides the strongest reported South Asian example of generative AI assisting operational research. The explosion killed eleven people and injured others. In May 2026, the National Investigation Agency filed a chargesheet against ten accused. Media accounts of the chargesheet identify Jasir Bilal Wani as the module's "in-house engineer" and state that he used ChatGPT and YouTube to research rocket construction and chemical proportions for explosive mixtures.¹⁶

The case matters because the reported use went beyond propaganda. Even so, the technology appears to have served as an information and research assistant. The conspiracy, procurement of materials, experimentation, assembly, and movement of the device still depended on human decisions and an organizational network. Improvised-explosive research also predates generative AI; chatbots may reduce the time and knowledge required to locate and organize information, but they do not remove the need for physical testing and logistics.

Organizational attribution requires careful consideration. The National Investigation Agency (NIA) investigation reportedly linked the accused to an effort to revive Ansar Ghazwat-ul-Hind (AGuH). A subsequent UN Monitoring Team report associated the attack with Al-Qaeda in the Indian Subcontinent.¹⁷ UN sanctions-monitoring reports compile information received from Member States and other sources; they do not replace criminal investigations or judicial findings. Accordingly, the precise relationship among the accused, AGuH, AQIS, and other ideological influences should be presented as an investigative question rather than an established organizational fact.

Capability Doesn't Prove Behavioral Impact

The ability to produce persuasive-looking material does not establish that the material changes behavior. A 2024 peer-reviewed study involving 8,221 respondents found that text generated by a large language model could be nearly as persuasive as articles drawn from real covert propaganda

¹⁶ Red Fort blast accused used AI to learn how to build IED, says NIA chargesheet" Indian Express, September 10, 2026. <https://indianexpress.com/article/cities/delhi/red-fort-blast-nia-chargesheet-chatgpt-flipkart-10870612/>

¹⁷ United Nations Security Council, Thirty-Eighth Report of the Analytical Support and Sanctions Monitoring Team, S/2026/651, 10 August 2026.

campaigns. Human curation or improved prompting could further increase its effect.¹⁸ The experiment demonstrates that AI can lower the cost of producing plausible persuasive text.

Its limits are equally important. The study measured immediate agreement among US respondents exposed to foreign influence material. It did not measure durable ideological commitment, willingness to join a violent organization, travel, financing, or participation in an attack. Terrorist radicalization also develops through personal networks, identity, perceived grievance, ideological authority, and opportunities for mobilization. AI-generated content may reinforce this process, but it is rarely the only cause.

Open-source monitoring consequently provides stronger evidence of supply than effect. Researchers can count synthetic videos, translated documents, or AI-generated images. They usually cannot determine how many viewers were persuaded, whether an online contact became a recruit, or whether a particular item influenced an operational decision. Claims that AI has improved terrorist recruitment should therefore be treated as hypotheses requiring evidence, not as established findings.

Implications for Indian Subcontinent

The appropriate response is to incorporate generative AI into existing counter-terrorism practice. Creating a separate category of "AI terrorism" could distract from the organizations, facilitators and grievances that give technology its effect. Five priorities follow from the evidence.

Priority	Required response
Multilingual monitoring	Security agencies and platforms can better support their work by strengthening their capabilities in Urdu, Hindi, Tamil, Bengali, and other Subcontinental languages, including Dhivehi, Pashto, and Persian, as well as transliterated and code-mixed forms. Automated systems can be valuable tools, working alongside (not replacing) analysts with linguistic and cultural expertise.
Digital forensics and evidence preservation	Investigators can benefit from clear procedures for preserving prompts, model responses, account records, device artefacts, and cloud data. These records can help them distinguish casual experimentation from sustained operational research.

¹⁸ Josh A Goldstein, Jason Chao, Shelby Grossman, Alex Stamos, Michael Tomz, "How persuasive is AI-generated propaganda?", PNAS Nexus, Volume 3, Issue 2, February 2024, pgae034, <https://doi.org/10.1093/pnasnexus/pgae034>

Priority	Required response
Risk-based attention to operational use	The highest priority should be given to activity that moves beyond ideological expression toward target selection, explosives, weapons, logistics, financing, or attack security. Monitoring and access to data should always remain grounded in legal authority and judicial oversight.
Platform and model-provider cooperation	AI companies can help strengthen safeguards by testing models against terrorism-specific abuse scenarios, improving reporting channels, and preserving relevant records when legally required. Social-media services can also improve detection of rapidly translated and synthetically altered variants, rather than relying only on exact-content hashes.
Measured public communication	Governments can help the public understand the risks of synthetic media without overstating terrorists' capabilities. Exaggerated claims may unintentionally draw attention to extremist content, weaken confidence in genuine evidence, and lead people to assume that every manipulated item is an advanced deepfake.

CONCLUSION

Generative AI is changing the efficiency of terrorist communication more clearly than it is changing the fundamentals of terrorism. Islamic State and ISKP-linked networks have shown that synthetic presenters, automated translation, image generation, and text-to-speech tools can increase the speed and range of propaganda production. The Red Fort investigation indicates that commercial AI may also assist technical research linked to attack preparation.

These developments warrant sustained attention, but the evidence does not support claims of a decisive transformation in terrorist capability. AI does not replace ideology, leadership, trusted networks, finance, logistics or physical access to targets. Nor is there clear evidence that synthetic propaganda produces higher recruitment or mobilization rates than earlier forms of extremist media.

For India and its neighbors, the most immediate challenge is combining inexpensive AI tools with existing multilingual extremist ecosystems. The practical response lies in better regional-language analysis, stronger digital forensics, closer cooperation with technology providers, and careful attention to operational uses. Policy should be guided by documented behavior and measurable risk rather than the novelty of the technology.



RUCHIKA KAKKAR is a Research Assistant at the Institute for Conflict Management (ICM), New Delhi, specializing in security dynamics across Jammu & Kashmir, Punjab, and South Asia. At the South Asia Terrorism Portal (SATP), she maintains regional conflict databases, tracks daily counter-terrorism developments, and contributes to 'South Asia Intelligence Review' (SAIR) assessments and global geopolitics analyses. A PhD candidate in Geopolitical Strategy at Sharda University, she researches strategic hedging in the Bay of Bengal. Ruchika holds an MA in Political Science (Delhi University), a BLIS (IGNOU), and is a two-time UGC-NET qualifier with prior university teaching experience.

Email: ruchikakkar504@gmail.com

The views expressed in this paper are personal.



**SOCIETY FOR THE STUDY OF
PEACE & CONFLICT**
Fostering Ideas, Research and Dialogue...

The Society for the Study of Peace and Conflict (SSPC) is an independent, non-profit, nonpartisan research organization based in New Delhi, dedicated to conduct rigorous and comprehensive research, and work towards disseminating information through commentaries and analyses on a broad spectrum of issues relating to peace, conflict, and human development. SSPC has been registered under the Societies Registration Act (XXI) of 1860. The SSPC came into being as a platform to exchange ideas, to undertake quality research, and to ensure a fruitful dialogue.

Copyright © Society for the Study of Peace and Conflict, New Delhi

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without first obtaining written permission of the copyright owner.

Published by:

Society for the Study of Peace and Conflict.
Post Box: 10560,
JNU Old Campus, New Delhi-110067.
[https:// sspconline.org](https://sspconline.org)
LI: www.linkedin.com/company/sspconline
X (Formerly Twitter): [@sspconline](https://twitter.com/sspconline)
FB: www.facebook.com/sspconline

We welcome your feedback. Email your comments at
“sspconline@gmail.com”

*Designed and typeset by Excel Solutions,
New Delhi.*