

WHEN THE ALGORITHM JOINS THE INSURGENCY

*AI, Terrorism and the Governance Test Nigeria's Leadership Cannot Afford to Fail
A Policy and Leadership Briefing for Nigeria's Government, Security and Institutional Leadership*

Oyewole O. Sarumi PhD

Digital Transformation Architect, Political Economy and Policy Analyst, and Leadership Strategist

Abstract

Artificial intelligence has crossed a threshold that African security planning cannot afford to treat as theoretical. A 2026 Cambridge University field study, drawn from fifty-seven interviews with twenty-seven former Boko Haram and Islamic State West Africa Province fighters, provides the first documented evidence that an active terrorist organisation has embedded frontier AI chatbots into attack planning, weapons troubleshooting, logistics and operational security. This article examines that evidence alongside two other 2026 developments of comparable weight: the demonstrated ability of generative AI to design functional viral genomes at Stanford University and the Arc Institute, and the accelerating autonomous cyber capability tracked by the United Kingdom's AI Security Institute. It situates these findings against the backdrop of Nigeria's seventeen-year insurgency and the joint United States–Nigeria military campaign against ISWAP that began in December 2025 and intensified through May 2026. Applying a disciplined evidentiary framework that separates the established from the speculative, the article argues that AI functions as a force multiplier rather than a creator of entirely new capability, compressing the distance between specialised expertise and operational execution for both insurgents and the legitimate state. It develops the concept of an AI security paradox, in which every security application of AI carries a corresponding vulnerability, and proposes a National AI Security, Counter-Terrorism and Biosecurity Observatory alongside an eight-layer governance framework spanning model testing, synthesis screening, laboratory security and international cooperation. The article closes with differentiated recommendations for government, security agencies, regulators, universities, businesses and civil society, and four scenarios for how this landscape may evolve by 2030. Its central argument is that Nigeria's leadership community faces a narrowing window to govern AI simultaneously as an economic opportunity and a national-security technology, before adversaries who already treat it as both gain a decisive and difficult-to-reverse advantage.

Keywords: artificial intelligence, terrorism, asymmetric warfare, Boko Haram, ISWAP, AI security, biosecurity, Nigeria

Introduction: The New Security Frontier

For most of the past decade, the dominant narrative connecting artificial intelligence to terrorism centred on propaganda: extremist groups using generative tools to produce recruitment videos, translate messaging into multiple languages, or automate social media accounts. That narrative was accurate as far as it went, but it understated the trajectory of the technology and, by extension, the trajectory of the threat. Frontier AI systems released over the past two years are no longer simple content generators. They are increasingly capable problem-solvers, code writers, research assistants and, in their more advanced agentic forms, autonomous actors able to pursue multi-step objectives with limited human supervision.

That shift matters for counter-terrorism, cybersecurity and biosecurity simultaneously, because the underlying capability, an AI system's growing ability to reason across a problem, retrieve relevant knowledge and propose or execute a sequence of actions, transfers freely across domains. A model that helps a legitimate engineer troubleshoot a mechanical fault can, in principle, help an insurgent troubleshoot a weapon. A model that helps a biosecurity researcher design a therapeutic bacteriophage can, in a different hand, illuminate part of the pathway toward more dangerous biological design. This article's task is to examine that transfer honestly, distinguishing what has actually been documented from what is technically plausible but undemonstrated and what remains speculative, before translating that assessment into a strategic response appropriate for Nigeria and the wider African security environment.

For a country in its seventeenth year of confronting an insurgency that has killed more than thirty-five thousand people and displaced over two million, this is not an abstract policy debate conducted in Washington think-tanks or London ministries. It is a live question intersecting with active combat operations, a widening informal technology market, and a national AI adoption drive that Nigeria's own government is, rightly, trying to accelerate. Between December 2025 and May 2026, the United States and Nigeria jointly conducted a series of airstrikes and special-forces raids against ISWAP and Boko Haram in the North-East, killing the senior ISWAP commander Abu-Bilal al-Minuki and, by the Nigerian Defence Headquarters' own count, roughly one hundred and seventy-five militants within days of the May offensive (U.S. Africa Command, 2026). That campaign is being waged against exactly the organisations that Cambridge researchers have now documented as active adopters of frontier AI. The kinetic and the digital dimensions of this insurgency are no longer separate stories; they are one story, and Nigeria's leadership must read them together.

From Propaganda to Operational Capability: What the Cambridge Evidence Shows

The empirical anchor for this analysis is a 2026 field study by Antonia Juelich, a terrorism researcher at the University of Cambridge's Programme on AI Science and Policy, titled "God Has Helped Us, and So

Will AI": How the Terrorist Group Boko Haram Uses Frontier AI (Juelich, 2026). The ninety-three-page study rests on fifty-seven interviews with twenty-seven former members of both Boko Haram factions, the Islamic State West Africa Province and Jama'at Ahl as-Sunnah lid-Da'wah wa'l-Jihad, conducted in north-eastern Nigeria between 2025 and 2026. It is, by Juelich's own characterisation and by the assessment of independent reviewers who examined the report after it was first reported by The New York Times, the first field-based evidence that an active terrorist organisation has adopted frontier AI as an organisational capability rather than merely as a messaging tool.

Former fighters told Juelich that both factions had established dedicated AI units, staffed by small teams equipped with paid subscriptions to multiple commercial platforms including ChatGPT, Claude, Gemini, Grok, Meta AI and DeepSeek, to manage access, train fighters and respond to operational queries (Premium Times, 2026). The pattern of use extended well beyond propaganda into what Juelich terms an entire attack chain: operational planning, weapons troubleshooting, explosives-related problem-solving, logistics, operational security and some drone-related activity. One widely corroborated account describes fighters using AI to calculate how to modify motorcycles and time jumps across a defensive trench in order to breach a fortified base, a granular illustration of precisely the kind of tactical problem-solving the study documents (France 24, 2026). This is the pattern that should concern Nigerian and West African policymakers most: not a single dramatic incident, but the quiet routinisation of AI as a consultative layer sitting above an organisation's existing capabilities.

It is equally important to state plainly what this evidence does not show. The study does not establish, and Juelich does not claim, that Boko Haram or ISWAP has developed a biological, chemical or radiological weapon using AI. She is careful to separate documented operational activity from her interviewees' stated attitudes toward mass-casualty weapons, and she acknowledges the study's own limitations: testimony that cannot be independently corroborated given the secrecy surrounding an active insurgency, and the absence of forensic or platform-level evidence linking specific attacks to specific AI interactions. A review of the report by Nigerian journalists at Premium Times similarly found that many of its central claims could not be independently verified and rested largely on the accounts of defectors, a caveat that belongs in any serious policy reading of the research rather than in a footnote (Premium Times, 2026).

"Not a single dramatic incident, but the quiet routinisation of AI as a consultative layer sitting above an organisation's existing capabilities."

Boko Haram, ISWAP and Nigeria's Seventeen-Year War

Boko Haram's insurgency gives this research an urgency that purely theoretical AI-safety debates elsewhere cannot convey. The interviewees told Juelich that foreign trainers, described as coming from Libya, France and Arab countries, had supplied laptops, virtual private networks, encrypted communications software and paid AI subscriptions during training sessions for selected commanders, with knowledge transfer occurring through transnational jihadist networks in a pattern consistent with how the Islamic State has historically disseminated technical capability across its provinces. Independent reporting following the Cambridge study's release, including a follow-up published by Al Jazeera on 18 September 2026 drawing on new material, has broadly corroborated the outline of this diffusion pathway even where specific tactical claims remain difficult to verify independently.

This organisational adoption is unfolding against the backdrop of the most significant escalation in the counter-insurgency campaign in years. Following an initial round of American airstrikes against armed groups in north-western Nigeria in December 2025, the United States and Nigeria launched a major joint offensive in the North-East in May 2026, combining special-forces raids with multiple rounds of airstrikes. The operation killed ISWAP's senior commander Abu-Bilal al-Minuki alongside several other commanders, and Nigerian authorities reported roughly one hundred and seventy-five militants killed within days (U.S. Africa Command, 2026). Foreign Policy's assessment of the campaign noted that, unlike Nigerian-led strikes that have frequently produced civilian casualties, the joint operations had so far avoided that outcome, a distinction with real implications for public trust in a theatre where civilian harm has historically deepened, rather than resolved, the grievances feeding recruitment (Foreign Policy, 2026).

The strategic lesson for Nigerian leadership is not that AI has replaced conventional counter-insurgency, which remains, as this campaign demonstrates, a matter of intelligence, special operations and airpower. The lesson is that the organisations Nigeria is fighting kinetically are the same organisations now documented as active, organised adopters of the world's most advanced consumer AI systems. Any counter-terrorism strategy that treats these as separate problems, one for the Defence Headquarters and one for the National Information Technology Development Agency, is already a strategy one step behind the adversary it is meant to defeat.

Reading the Evidence Honestly: Capability Uplift, Not Capability Creation

The concept that best captures what the Cambridge evidence demonstrates is capability uplift rather than capability creation. Boko Haram already possessed weapons, motorcycles, drones, explosives, communications systems and personnel before any of its members opened a chatbot. What AI appears to have added is not a new class of weapon but a force multiplier that makes existing capabilities faster, cheaper, more accurate and less dependent on scarce human expertise. A human explosives specialist can


train perhaps a handful of people directly; an AI system, accessed simultaneously by multiple fighters across multiple locations, can in principle answer thousands of operational questions without the geographic and capacity constraints that have historically limited the diffusion of specialised knowledge within an insurgent organisation.

This is the essential strategic insight for leaders across every sector, not only counter-terrorism but business and public administration: technology diffuses into adversarial environments as readily as it diffuses into legitimate ones, and often faster, because adversarial organisations accept levels of risk, secrecy and experimentation that no regulated institution could tolerate. A terrorist organisation does not need board approval, a procurement committee or a five-year AI strategy before experimenting with a new tool. That asymmetry in institutional friction, not any inherent technical sophistication among the adopters, is what allows a fragmented insurgency to move faster on AI adoption than many well-resourced corporations and, uncomfortably, than many government agencies.

Leadership Reflection

Nigerian institutions rarely lose a technology race because a rival organisation possesses smarter tools. They lose it because bureaucratic friction, the very governance discipline that legitimate institutions are right to maintain, becomes a competitive disadvantage against adversaries who observe none of it. The task for leadership is not to imitate that recklessness but to remove the friction that serves no genuine safeguarding purpose while strengthening the friction that does.

Figure 1. The Capability Compression Pathway



The pathway above summarises the conceptual model running through this article: AI capability compresses specialised knowledge, producing an uplift in what any actor, legitimate or malicious, can accomplish. That uplift feeds organisational adaptation and, for adversarial actors, an asymmetric

advantage that in turn generates security risk. The only stage in this chain that institutions can reliably control is the last one: governance response. Everything upstream of it is moving at a pace set by the technology itself.

The Rise of Agentic AI and the Compression of Expertise

The distinction between a chatbot and an AI agent is the single most important conceptual shift underlying this entire analysis. A chatbot answers a question. An agent pursues an objective: it can inspect an environment, identify a problem, test a proposed solution, adapt its approach, interact with external tools and continue through multiple stages toward a goal with limited human intervention at each step. The UK AI Security Institute, a directorate of the Department for Science, Innovation and Technology that has spent two years running structured pre- and post-deployment evaluations of frontier models, documented in its inaugural Frontier AI Trends Report that models can now complete autonomous software tasks that would take a skilled human expert more than an hour, and that success rates on internal self-replication evaluations rose from roughly five per cent in 2023 to sixty per cent in 2025, even as the Institute stressed it has found no evidence of such behaviour occurring spontaneously outside controlled testing (AI Security Institute, 2026).

This transition from tool to agent is precisely what converts a capability that is merely useful into one that is strategically significant, because an agent's value to an adversary scales with the amount of autonomous, unsupervised action it can take before a human needs to check its work. For a terrorist organisation operating with limited technical staff and a strong incentive toward operational secrecy, an agent able to work through a multi-step technical problem with minimal supervision offers exactly the kind of capability compression that most threatens existing counter-terrorism assumptions, assumptions that have historically relied on the scarcity of technical expertise as an implicit barrier to sophisticated attacks. Nigerian security institutions that continue to plan around that scarcity as though it still holds are planning against a threat model that no longer describes reality.

AI and the Escalating Cyberwarfare Frontier

The evidence on autonomous cyber capability is now considerably more concrete than most public discussion suggests, and it is also more specific than the claim, unsupported by the evidence, that AI models are already hacking systems independently at scale. What the evidence does support is a clear and accelerating trend. The UK AI Security Institute's cyber-capability tracking found that the length of cyber tasks frontier models can complete autonomously, measured against the time such tasks would take a skilled human, has been doubling roughly every few months since late 2024, with the doubling rate itself

accelerating from an eight-month estimate in November 2025 to a 4.7-month estimate by February 2026, a pace the newest models had already begun to exceed (AI Security Institute, 2026a). Separately, the Institute's evaluation of frontier models against purpose-built cyber ranges simulating multi-stage corporate and industrial-control-system attacks found that model performance scaled steadily with the computing resource applied at inference time, with no observed plateau, a finding with direct implications for how cheaply a well-resourced attacker could scale an operation (AI Security Institute, 2026b).

For Nigeria, where financial-sector digitisation and government service delivery are both advancing rapidly under the National Artificial Intelligence Strategy, this trend deserves the same institutional seriousness as the physical insurgency. Banks, telecommunications operators and government platforms that have not yet experienced an AI-accelerated intrusion should not read that absence as evidence of safety; it is at least as plausible evidence of a defensive posture that has not yet been tested against the current generation of tools. Independent testing cited by industry security bodies has found that AI safety commitments across major developers have, if anything, eroded rather than strengthened over the past year, and that models still refuse a meaningful share of terrorism-relevant prompts only inconsistently (ASIS International, 2026). Nigerian institutions should not assume that safeguards built elsewhere will hold reliably on their behalf.

Drones, Robotics and the Falling Cost of Asymmetric Violence

Low-cost commercial drones adapted for reconnaissance and attack have already reshaped asymmetric warfare in conflicts from Ukraine to the Middle East, and the integration of AI-enabled computer vision, autonomous navigation and target recognition into these platforms is accelerating that shift further. The Cambridge study's reference to drone-related AI use within Boko Haram and ISWAP, though less developed in the available evidence than the group's use of AI for weapons troubleshooting and planning, situates Nigeria squarely within this global trend rather than apart from it. The strategic logic is straightforward: a commercially available drone, already inexpensive relative to conventional military hardware, becomes still more effective when paired with AI systems that assist with flight planning, targeting or evasion of countermeasures, lowering the technical bar for an insurgent group to conduct reconnaissance or deliver a payload with meaningful precision.

This article does not provide, and should not be read as endorsing, any operational detail regarding the construction or deployment of weaponised drones. The analytical point is strategic rather than technical: AI is reducing the cost of asymmetric warfare across the drone and robotics domain in the same way it appears to be reducing the cost of cyber operations and, in the Boko Haram case, of conventional tactical planning. Nigeria's own counter-drone and airspace-monitoring capacity, still developing relative to the

scale of its security challenges, needs to keep pace with an adversary environment in which low-cost autonomous systems are becoming steadily more capable and more accessible.

Propaganda, Psychological Operations and the Battle for Belief

Even as this article argues that AI's more consequential near-term risk lies in operational uplift rather than propaganda, the information domain remains a significant battlefield in its own right. Generative AI has made it materially cheaper to produce multilingual extremist content, synthetic voices and personalised recruitment messaging at a scale earlier propaganda operations could not sustain. The United Nations Office on Drugs and Crime's 2025 symposium on AI and counter-terrorism highlighted precisely this dual character: the same generative capability that fuels radicalisation through rapidly produced propaganda can, with appropriate safeguards, be redirected toward detecting and countering that same content (UNODC, 2025).

The UN Security Council's Counter-Terrorism Committee Executive Directorate has likewise briefed member states that while AI is not yet regarded as an imminent, transformative threat in the propaganda domain specifically, it is being actively monitored, alongside a parallel effort by technology platforms and civil-society organisations to use AI-based content moderation to identify and remove terrorist material before it spreads (UN Security Council CTC, 2026). Nigeria's own counter-disinformation capacity, whether in government communications, security agencies or civil society, should be understood as a genuine component of counter-terrorism policy rather than a separate media-literacy initiative, precisely because propaganda and operational planning increasingly draw on the same underlying AI infrastructure.

The AI–Biology Convergence: Sixteen Viruses and a Warning from Johns Hopkins

The most scientifically significant development informing this article occurred only weeks before it was finalised. On 6 August 2026, a team led by Stanford bioengineering graduate student Samuel King, working with Brian Hie of Stanford and the Arc Institute, published in the peer-reviewed journal *Science* the first demonstration of generative AI designing complete, functional viral genomes from scratch (King et al., 2026). Using genome language models called Evo 1 and Evo 2, trained on roughly two million bacteriophage genomes and fine-tuned on the family containing the well-studied phage ΦX174, the team generated approximately seven hundred thousand candidate genomes, selected three hundred and two for laboratory synthesis, successfully built two hundred and eighty-five, and found that sixteen were fully functional viruses capable of infecting and killing the bacterium *Escherichia coli*, with some variants suppressing drug-resistant bacterial strains more effectively than the natural template they were modelled on (Stanford News, 2026).

The qualification that matters most, and the one most likely to be lost in casual retelling, is that these are bacteriophages, viruses that infect bacteria, not humans, animals or plants, and the researchers deliberately excluded eukaryote-infecting viruses from the training data, worked with a non-pathogenic host, and conducted the work in a secure laboratory. Several of the designed phages were explored specifically for their therapeutic potential against antibiotic-resistant infections, a genuinely promising application given the scale of the global antimicrobial-resistance crisis. But the underlying capability is unambiguously dual-use, and the companion Perspective article published in the same issue of *Science*, by Thomas Inglesby and Moritz Hanke of the Johns Hopkins Center for Health Security, drew the line explicitly: the ability to compose viral genomes using generative AI now exists, they wrote, while the governance to safely steer it does not, and extending the method to viruses capable of infecting humans, animals or plants could in principle produce pathogens falling outside what current countermeasures can contain (as summarised in BetaNews, 2026).

A finding with more immediate policy relevance than the science itself is the governance gap it exposed. Biosecurity specialists who reviewed the study warned that AI-generated genomes can evade existing DNA-synthesis screening tools, which are built primarily on databases of known pathogens and struggle to flag genuinely novel sequences that were never catalogued in the first place. This is not a hypothetical regulatory footnote. It is the precise mechanism through which a dual-use scientific advance could, over time, outrun the screening infrastructure every credible biosecurity framework, including Nigeria's own, currently depends upon.

Could AI Lower the Barrier to Biological Weapons? An Honest Answer

The honest answer, grounded in the evidence reviewed here, is: potentially, over time, but not in a manner yet demonstrated, and not without overcoming steps that AI alone does not solve. The relevant chain runs from AI-assisted design, through physical DNA or RNA synthesis, laboratory validation, scaling, and ultimately delivery and successful transmission in a target population. AI is increasingly capable of accelerating the first step of that chain, and the Stanford-Arc Institute results confirm that the design-to-function pathway can now be closed experimentally for at least one class of relatively simple biological entity. It does not follow that the same pathway is currently closeable for a human pathogen, a vastly more complex undertaking involving immune evasion, transmissibility and stability requirements that remain, on the present evidence, beyond what generative biological design has demonstrated.

The OECD's technology assessment on synthetic biology, AI and automation reaches a complementary conclusion at the level of governance rather than pure science: the convergence of AI, automated laboratory equipment and synthetic biology is progressively democratising access to biological design and

experimentation capability that was previously concentrated in a small number of specialised institutions, a trend the OECD frames explicitly as lowering barriers to entry with corresponding implications for misuse involving toxins and pathogens, even as it also expands legitimate access to life-saving biotechnology (OECD, 2025). The appropriate policy response, as both the National Academies of Sciences, Engineering, and Medicine and the OECD argue, is systemic: synthesis-order screening, laboratory-automation access controls and biosecurity governance addressing the entire ecosystem, not a narrower approach that assumes the risk can be managed simply by instructing AI chatbots to refuse dangerous prompts (National Academies of Sciences, Engineering, and Medicine, 2025).

What Is Established, What Is Emerging and What Remains Speculative

A forensic discipline requires the analyst to state uncertainty as clearly as certainty. On the current evidence, Boko Haram and ISWAP's operational adoption of frontier AI for planning, troubleshooting and logistics is field-documented and strongly supported, even as specific tactical claims within that evidence carry the corroboration limits described earlier. Rapidly advancing autonomous cyber capability is demonstrated under controlled evaluation by a credible government institution. AI-designed, functional novel biological entities have been demonstrated experimentally, though confined so far to bacteriophages rather than human pathogens. The claim that most AI models are already autonomously hacking real-world systems at scale is not supported and should be set aside as exaggeration. The claim that terrorist groups have already used AI to create biological weapons is not supported and should likewise be set aside. The claim that AI could, over a longer horizon, meaningfully compress the technical distance between extremist intent and biological, cyber or conventional operational capability is a credible, evidence-grounded concern that merits sustained institutional attention precisely because it does not require exaggeration to justify that attention.

Evidentiary Tier	Definition	Illustrative Finding in This Article
Established / Confirmed	Field-documented, corroborated, or peer-reviewed	Boko Haram/ISWAP organisational adoption of frontier AI for planning, troubleshooting and logistics (Juelich, 2026); AI-designed functional bacteriophage genomes (King et al., 2026)
Emerging / Strongly Supported	Demonstrated under controlled evaluation or by credible	Accelerating autonomous cyber-task completion tracked by the

Evidentiary Tier	Definition	Illustrative Finding in This Article
	institutions, real-world consequence not fully settled	UK AI Security Institute (2026a, 2026b)
Plausible but Unproven	Reasonable extrapolation from current trajectories, requiring further evidence	AI-assisted compression of the pathway toward more capable biological or cyber weapons over the coming decade
Speculative / Not Established	Insufficient evidence; should not inform policy as fact	Claims that terrorist organisations have already produced biological, chemical or radiological weapons using AI

This discipline, distinguishing confirmed field evidence from strongly supported scientific trend from genuinely speculative extrapolation, is not an academic nicety. Security institutions that base policy on exaggerated claims lose credibility and risk misallocating scarce resources toward the wrong threat model; institutions that ignore well-supported trends because the most sensational version of a claim later proved false make the opposite and equally costly error. Nigeria's policy response, developed in the sections that follow, is designed to track the evidence at each of these levels rather than the loudest version of the headline.

Nigeria at the AI–Terrorism Frontier: A Convergence of Vulnerabilities

Nigeria's exposure to this emerging threat landscape is compounded by a specific combination of factors that few other African states share simultaneously. Active insurgencies, principally Boko Haram and ISWAP, already operate within Nigerian territory and have now been field-documented as frontier AI adopters, even as they face the most intensive joint counter-insurgency campaign in years. Nigeria's borders sit within a wider, porous West African and Sahelian security ecosystem through which the Cambridge study's reported transnational knowledge transfer, from Libya, from Islamic State networks in the Sahel, and from further afield, can continue to flow. A large informal technology market makes drones, electronics and computing resources increasingly accessible outside formal regulatory channels.

Nigeria's own AI adoption is expanding rapidly among young people, universities, businesses and government under the National Artificial Intelligence Strategy 2025–2029, a genuine economic opportunity that simultaneously widens the population of potential users, both legitimate and otherwise,

with access to frontier tools (Federal Ministry of Communications, Innovation and Digital Economy, 2025). Institutional cybersecurity capacity across much of the public and private sector remains, by most independent assessments, still maturing relative to the sophistication of the threats now being documented. And Nigeria's public-health system, already stretched, would face acute strain from any biological incident, however unlikely, that outpaced existing surveillance and response capacity.

None of this should be read as suggesting Nigeria is uniquely or fatalistically vulnerable. It should instead be read as an argument that AI governance in Nigeria cannot be treated purely as an economic-development or productivity question, the frame that has understandably dominated the National Artificial Intelligence Strategy's public presentation to date. It is also, and increasingly, a national security question, and the institutions responsible for economic AI policy and the institutions responsible for security policy need structured mechanisms to coordinate rather than operate on separate tracks that happen to share the same underlying technology.

AI as a Weapon of Counter-Terrorism: The Underused Opportunity

The same convergence of capabilities that concerns this article also represents a genuine opportunity for the Nigerian state, currently under-exploited. AI-enabled intelligence fusion can help analyse far larger volumes of open-source and lawfully collected intelligence than human analysts can process unaided. AI-enabled geospatial and satellite analysis can help identify suspicious movement patterns across difficult terrain that has historically limited conventional surveillance in the North-East. AI-enabled predictive logistics analysis can help identify anomalous supply-route activity, while AI-enabled cybersecurity tools can help detect intrusions against government and critical infrastructure networks before they succeed. AI-enabled multilingual analysis, deployed under appropriate lawful oversight, can help process the volume of communications relevant to counter-radicalisation and financial-intelligence work far faster than existing capacity allows, and AI-enabled counter-disinformation tools can help identify coordinated extremist influence operations as they emerge rather than after they have spread.

Each of these applications, as with the risks described earlier in this article, carries a corresponding governance requirement, addressed in the section that follows. But the strategic objective for Nigeria should not be framed as stopping AI. It should be framed as ensuring that Nigeria's legitimate security, intelligence and public-health institutions possess AI capability at least as sophisticated as what is now available to the adversaries documented in the Cambridge research, an objective the country is currently far from meeting, given the gap between the informal, ad hoc AI access described among Boko Haram's dedicated units and the institutional capacity of most Nigerian security agencies.

The AI Security Paradox

Every application of AI to national security carries within it the seed of a corresponding vulnerability, and Nigerian policymakers should confront this paradox directly rather than allow enthusiasm for AI-enabled counter-terrorism to obscure it. AI built for cybersecurity defence can, in the wrong hands, be repurposed for cyberattack. AI built to accelerate biological discovery can, as the Stanford-Arc Institute results demonstrate, also expand the frontier of biological design in ways that are not automatically benign. AI built for drone-based surveillance edges uncomfortably close to autonomous systems capable of lethal targeting. AI built for intelligence fusion can, without careful constraint, slide into mass surveillance that erodes the civil liberties the state is nominally trying to protect. AI built to detect disinformation can be redirected to suppress legitimate dissent if institutional safeguards are weak.

This paradox does not counsel inaction. It counsels governance built into the architecture of AI-enabled security tools from the outset rather than added retroactively after a scandal or a wrongful-targeting incident forces the issue. The single most important safeguard, discussed further below, is that AI in national-security applications should be restricted to detection, analysis and recommendation, with humans retaining accountability for any consequential decision, particularly any decision involving the use of force.

"Every application of AI to national security carries within it the seed of a corresponding vulnerability."

Human Oversight, Ethics and the Law of Armed Conflict

The principle that AI may recommend but humans must decide is not a rhetorical flourish; it is the operational safeguard against the specific failure modes AI systems are known to exhibit, including hallucination, algorithmic bias and false-positive identification. An AI intelligence system that misidentifies an innocent civilian as a combatant, or a drone-vision system that misclassifies a target, does not merely produce an isolated error; in an active counter-insurgency environment, such an error can produce a wrongful death and, cumulatively, can undermine the legitimacy of the entire counter-terrorism effort among the civilian population whose cooperation that effort ultimately depends on. International humanitarian law's existing requirements around distinction, proportionality and precaution were not written with autonomous decision-support systems in mind, and Nigeria, alongside other troop-contributing and conflict-affected states, has a direct interest in ensuring that emerging state practice on human-in-the-loop requirements for lethal decisions keeps pace with the underlying technology rather than trailing it by a decade, as regulatory frameworks so often do.

This is also where Nigeria's broader human-rights obligations intersect directly with its counter-terrorism posture. AI-enabled surveillance and intelligence tools, deployed without clear legal authorisation, oversight mechanisms and redress procedures, create exactly the kind of unaccountable state capability that erodes public trust and, in the specific context of an insurgency that has partly fed on grievances against state institutions, can be strategically counterproductive even when the underlying security objective is legitimate. The Foreign Policy assessment of the 2026 joint strikes noted that their careful targeting, so far avoiding the civilian casualties that have historically accompanied Nigerian-led operations, has helped sustain public support (Foreign Policy, 2026); that same discipline, applied to AI-enabled decision-making, is what will determine whether the tools described in this article strengthen or erode the state's legitimacy over time.

Toward a National AI Security, Counter-Terrorism and Biosecurity Architecture

This article recommends that Nigeria establish a National AI Security, Counter-Terrorism and Biosecurity Observatory, convened jointly by the National Information Technology Development Agency, the Office of the National Security Adviser, the Department of State Services, the Nigeria Police Force, the Defence Space Administration, the Armed Forces of Nigeria, the Nigeria Centre for Disease Control, the Federal Ministry of Health, relevant universities, and appropriate private-sector cybersecurity and biotechnology partners. Its mandate should include monitoring AI-enabled terrorism and extremist adoption of frontier AI; monitoring AI-enabled cyber threats against Nigerian critical infrastructure; studying the AI-biology convergence as it applies to Nigerian public-health preparedness; building indigenous AI red-teaming capability rather than relying entirely on foreign assessments; establishing national AI-incident reporting mechanisms; training security and health personnel; supporting university research programmes in this space; and coordinating intelligence-sharing partnerships with regional and international counterparts, given that AI-enabled terrorism, like the underlying technology itself, is inherently transnational.

Why an Observatory, Not Another Committee

Nigeria does not lack committees. It lacks a standing institution with the authority, budget and technical depth to track a threat that moves faster than the annual reporting cycles most inter-agency committees are built around. An Observatory succeeds only if it is funded and staffed as a genuine operational capability, with dedicated technical personnel who spend their time on this problem rather than a rotating cast of officials attending it as one item on a crowded agenda.

The Eight-Layer Governance Framework

Beneath that coordinating body, this article proposes an eight-layer governance framework spanning the full pathway from model design to physical deployment. The logic is cumulative: no single layer is sufficient on its own, and the framework's value lies precisely in the fact that a threat which slips past one layer should still be caught by another. The layers move from the most abstract, testing of frontier models themselves, through access controls, physical-world safeguards for biological and cyber capability, and finally to the institutional and international cooperation without which no single layer can hold indefinitely.

Layer	Governance Focus	Threat Addressed
1	Model-level security testing of frontier systems	Biological, cyber, chemical and autonomous-behaviour misuse at the model stage
2	Tiered access controls and authentication	Unmonitored access to higher-risk capabilities by unvetted users
3	DNA/RNA synthesis-order screening	Physical realisation of AI-designed biological sequences, including designs that evade legacy screening tools
4	Laboratory security requirements	Connection of automated laboratory equipment to unrestricted autonomous agents
5	AI-agent cybersecurity controls	Agents operating outside bounded, monitored environments
6	Drone and robotics controls	Autonomous systems gaining unrestricted access to lethal physical capability
7	Government–university–private-sector intelligence sharing	Fragmented awareness of emerging misuse patterns
8	International cooperation	Cross-border diffusion of AI-enabled threats that no single state can govern alone

The third and fourth layers deserve particular emphasis given the Stanford-Arc Institute findings discussed earlier. If AI-generated biological sequences can evade synthesis-screening tools built on databases of known pathogens, then screening infrastructure that Nigeria has not yet built cannot be allowed to inherit the same blind spot from the outset. This is a rare case where a developing biosecurity system has the advantage of building its screening architecture after, rather than before, the vulnerability was identified, provided the Nigeria Centre for Disease Control and the Federal Ministry of Health treat this as a design requirement now rather than a retrofit later.

Universities and the Talent Nigeria Is Not Yet Producing

Nigerian universities have an unusually consequential role to play in this landscape, distinct from the important but different task of teaching students to use AI productively. Universities should establish structured programmes in AI safety, AI security, AI governance, AI biosecurity and AI-enabled national-security research, producing a smaller but highly specialised cadre of graduates who understand not only how to deploy AI but how to secure it. The risk, already visible in the broader pattern of AI adoption across Nigerian higher education, is that the country trains a large population of AI users while producing far too few AI security and safety specialists, precisely the imbalance that leaves institutions structurally unprepared for the risks described throughout this article.

This is not a call for another generic AI curriculum. It is a call for Nigerian universities to build formal research partnerships with international counterparts already engaged in this specific work, following the model Juelich herself recommends of involving conflict and terrorism researchers directly in AI safety assessment rather than relying on technical testing alone. A handful of Nigerian institutions with the ambition to lead in this space could, within a five-year horizon, become the reference point for AI-security research across West Africa, an outcome that would strengthen the country's hand in exactly the international cooperation the eight-layer framework depends upon.

Business, Boards and the Twin Maturity Model

Nigerian businesses face a parallel and complementary challenge. Many organisations remain at the earliest stages of AI maturity, using generative tools primarily for content creation, while the genuine competitive frontier has already moved toward AI-enabled workflows, autonomous agents and organisational intelligence. This article argues that businesses need a second, parallel maturity model running alongside their AI productivity ambitions: an AI security maturity model. Every organisation deploying AI should be asking two questions simultaneously rather than sequentially, how much value is being extracted from AI and how much risk is being created by deploying it, questions that remain

inseparable in any serious AI strategy but are still treated as separate workstreams in most Nigerian organisations today, a gap that mirrors, at smaller scale, the gap this article identifies at the national level.

Boards and chief executives should treat this as a governance matter rather than a purely technical one. The same logic that requires a board to understand financial risk exposure now requires it to understand AI risk exposure, including the risk that a vendor's model, procured for a legitimate commercial purpose, carries safeguards no more reliable than the fifty-seven per cent refusal rate industry testing has found in some frontier systems confronted with terrorism-relevant prompts (ASIS International, 2026). Chief information security officers and chief information officers should be reporting on both dimensions of AI maturity to the same board committee, not to separate ones that rarely compare notes.

Africa's Shared Exposure and the Case for Regional Governance

Nigeria's exposure sits within a wider African and West African pattern. The Sahel region's overlapping insurgencies, weak border-security architecture in several states, and established transnational jihadist networks mean that the diffusion pattern documented in the Cambridge study, foreign trainers, shared technical resources, knowledge transfer across affiliated groups, is unlikely to remain confined to Nigeria alone. The African Union and regional bodies such as the Economic Community of West African States have a direct interest in developing shared frameworks for AI-enabled counter-terrorism intelligence-sharing, cybersecurity cooperation and biosecurity coordination, given that fragile states with the least developed AI governance capacity are, on the logic developed throughout this article, also the states most exposed to AI-enabled asymmetric threats.

There is a deeper strategic risk here that extends beyond counter-terrorism narrowly defined. African countries that treat AI purely as an imported productivity technology, consuming models, infrastructure and security tools developed and controlled elsewhere, risk a form of strategic dependency that compounds their exposure, dependent not only economically but in the specific capabilities needed to detect and counter AI-enabled threats originating from better-resourced adversaries or from criminal and extremist networks operating with fewer constraints. This is precisely why Nigeria's parallel interest in African AI infrastructure, local-language models and sovereign data capacity carries a national-security dimension that deserves more explicit recognition than it currently receives in economic-development framing alone.

Four Futures: Scenario Analysis to 2030

Four scenarios illustrate how this landscape could plausibly evolve by 2030, offered as a structured way of thinking about uncertainty rather than as precise probability estimates, which the underlying evidence does not support with confidence. In a Controlled Diffusion scenario, governments, AI developers and

international institutions succeed in establishing meaningful model-level safeguards, synthesis-screening infrastructure and cyber-defence capacity that keeps pace with offensive capability, and AI-enabled terrorism remains a persistent but manageable threat, broadly analogous to how cybercrime is managed today. In a Criminal AI Explosion scenario, safeguards lag capability growth, and both cybercriminal networks and terrorist organisations exploit increasingly autonomous AI agents at a scale that outpaces current law-enforcement and counter-terrorism capacity, particularly in states with weaker institutional resilience.

In an AI–Biology Convergence scenario, generative biological design continues advancing toward increasingly complex organisms, and biosecurity governance either keeps pace, through robust synthesis screening and laboratory controls, or fails to, with correspondingly divergent outcomes for global health security. In an Autonomous Asymmetric Warfare scenario, small non-state actors gain sustained access to AI-enabled cyber, drone and intelligence capability that begins to approach what only well-resourced state actors could historically deploy, fundamentally altering the balance between states and non-state armed groups.

For Nigeria specifically, the early-warning indicators worth monitoring across all four scenarios include the scale and sophistication of AI use documented in future field research on Boko Haram, ISWAP and related groups; the pace of AI-enabled cyberattacks against Nigerian financial and government infrastructure; developments in international DNA-synthesis screening standards and Nigeria's own biosecurity infrastructure; and the rate at which frontier AI developers strengthen or weaken safeguards against terrorism-relevant misuse categories, a trend the UK AI Security Institute's ongoing evaluation programme is well positioned to track and that Nigerian institutions should follow closely rather than assume is being handled adequately elsewhere.

Strategic Recommendations for Nigerian Leadership

For government and national security agencies, the priority is to establish the proposed National AI Security, Counter-Terrorism and Biosecurity Observatory as a standing coordination mechanism, with a mandate and budget commensurate with the scale of the threat documented in this article rather than treated as a symbolic gesture. For the armed forces, police and intelligence agencies, the priority is deliberate investment in AI-enabled intelligence fusion, geospatial analysis and cyber-defence capability, paired explicitly with human-in-the-loop safeguards for any application touching lethal decision-making. For the National Information Technology Development Agency and technology regulators, the priority is to integrate national-security and biosecurity considerations formally into the implementation of the

National Artificial Intelligence Strategy, rather than treating economic development and security governance as separate workstreams that happen to share the same underlying technology.

For universities and biotechnology institutions, the priority is to establish the specialised AI safety, security and biosecurity programmes described earlier, and to build formal research partnerships with international counterparts already engaged in this work, including the institutions cited throughout this article. For AI companies and cybersecurity firms operating in or serving the Nigerian market, the priority is to engage directly with Nigerian counter-terrorism and biosecurity researchers in safety testing, following the model Juelich herself recommends of involving conflict and terrorism researchers in AI safety assessment rather than relying on technical testing alone. For healthcare institutions and the Nigeria Centre for Disease Control, the priority is to treat AI-biosecurity preparedness as a genuine extension of existing public-health emergency planning rather than a separate, lower-priority initiative.

For businesses, the priority is to adopt the parallel AI productivity and AI security maturity framework described earlier as standard board-level practice rather than a technical afterthought delegated entirely to information-technology departments. For civil society, the priority is to maintain active scrutiny of AI-enabled state surveillance capability, ensuring the safeguards recommended in this article are genuinely implemented rather than merely announced. And for international partners, the priority is to recognise that supporting Nigerian and West African AI governance capacity is not charity but a direct investment in reducing the transnational diffusion pathways the Cambridge research has now documented, since a weakly governed node anywhere in the region becomes a diffusion point for the entire region.

Conclusion: Governing the Age of AI-Enabled Asymmetry

The central finding running through this article is not that terrorists have acquired some dramatic new category of weapon. It is that the distance separating specialised knowledge from operational execution is narrowing, steadily and measurably, across cyber operations, biological design and now, field-documented for the first time, terrorist organisational planning itself. AI is increasingly capable of functioning as teacher, analyst, engineer and planner simultaneously, and when that capability connects to physical tools, whether weapons, drones or laboratory equipment, the implications compound rather than merely add.

Nigeria did not create this dynamic and cannot resolve it alone, but Nigeria has more reason than most nations to take it seriously, given the insurgency already operating within its borders, now confirmed by field research to be an early adopter of exactly the capability this article examines, and now the subject of the most intensive joint military campaign in years. The appropriate response is neither the alarmist claim that AI will inevitably produce catastrophe nor the complacent claim that AI is simply another tool no different from those that came before it. It is the more demanding, more accurate position developed

throughout this article: AI is simultaneously a productivity revolution, a scientific revolution and a security revolution, and the institutions capable of recognising and governing all three dimensions at once, rather than treating them as separate portfolios, will be the ones best placed to navigate what follows.

The decisive question of the coming decade is not who possesses the most powerful AI, but who can govern, secure and responsibly deploy increasingly autonomous AI capability before those who intend harm learn to exploit it. Nigeria's institutions, universities and leadership community have an opportunity, still open but narrowing, to answer that question proactively rather than in the aftermath of the incident that finally forces the issue.

"The decisive question of the coming decade is not who possesses the most powerful AI, but who can govern, secure and responsibly deploy it before those who intend harm learn to exploit it."

About the Author

Prof. Sarumi is a digital transformation architect, political economy and policy analyst, and leadership strategist with over 40 years of cross-sector experience across Nigeria and the African continent. He writes from Lagos.

References

- AI Security Institute. (2026). Frontier AI trends report. UK Department for Science, Innovation and Technology. <https://aisi.gov.uk/frontier-ai-trends-report>
- AI Security Institute. (2026a). How fast is autonomous AI cyber capability advancing? <https://www.aisi.gov.uk/blog/how-fast-is-autonomous-ai-cyber-capability-advancing>
- AI Security Institute. (2026b). Measuring AI agents' progress on multi-step cyber attack scenarios. <https://www.aisi.gov.uk/research/measuring-ai-agents-progress-on-multi-step-cyber-attack-scenarios>
- ASIS International. (2026, July 10). Terrorist groups use AI as a battle consultant, not just a propaganda tool. Security Management. <https://www.asisonline.org/security-management-magazine/latest-news/today-in-security/2026/july/boko-haram-AI-use/>
- BetaNews. (2026, August 6). AI-designed viruses: Stanford and Arc Institute build first complete AI-generated viral genomes. <https://betanews.com/article/ai-designed-viruses-stanford-arc-institute/>
- Federal Ministry of Communications, Innovation and Digital Economy. (2025). National Artificial Intelligence Strategy (NAIS) 2025–2029. Federal Government of Nigeria, with the National Information Technology Development Agency.
- Foreign Policy. (2026, May 27). Joint U.S.-Nigeria strikes: Can Washington end the African country's insurgency? <https://foreignpolicy.com/2026/05/27/us-nigeria-joint-strikes-military-insurgency-boko-haram-iswap-security-africa/>
- France 24. (2026, July 14). How jihadist groups like Boko Haram use AI for acts of terror. <https://www.france24.com/en/africa/20260714-how-jihadist-groups-like-boko-haram-use-ai-for-acts-of-terror>
- Juelich, A. (2026). "God has helped us, and so will AI": How the terrorist group Boko Haram uses frontier AI. Cambridge Programme on AI Science and Policy, University of Cambridge.
- King, S. H., Driscoll, C. L., Li, D. B., Guo, D., Merchant, A. T., Brixi, G., Wilkinson, M. E., & Hie, B. L. (2026). Generative design of bacteriophages with genome language models. Science. <https://www.science.org/doi/10.1126/science.aec2657>
- National Academies of Sciences, Engineering, and Medicine. (2025). The age of AI in the life sciences: Benefits and biosecurity considerations. The National Academies Press. <https://doi.org/10.17226/28868>

Organisation for Economic Co-operation and Development. (2025). Synthetic biology, AI and automation: A forward-looking technology assessment (OECD Science, Technology and Industry Policy Papers, No. 187). OECD Publishing. <https://doi.org/10.1787/12158721-en>

Premium Times. (2026, July 13). Boko Haram used ChatGPT, other AI tools to plan attacks, build bombs — Report. <https://www.premiumtimesng.com/news/top-news/894830-boko-haram-used-chatgpt-other-ai-tools-to-plan-attacks-build-bombs-report.html>

Stanford News. (2026, August 6). Evo 2: An AI tool and its E. coli-killer bacteriophages. Stanford University. <https://news.stanford.edu/stories/2026/08/evo-2-ai-tool-e-coli-killer-bacteriophages>

U.S. Africa Command. (2026, May 18). U.S.-Nigeria coordinated strike against ISIS fighters. <https://www.africom.mil/pressrelease/36581/us-nigeria-coordinated-strike-against-isis-fighters---may-18-2026>

United Nations Office on Drugs and Crime. (2025). 2025 UNODC symposium explores the role of artificial intelligence in preventing and countering terrorism. <https://www.unodc.org/unodc/en/terrorism/latest-news/2025-unodc-symposium-explores-the-role-of-artificial-intelligence-in-preventing-and-countering-terrorism.html>

United Nations Security Council Counter-Terrorism Committee. (2026). CTED holds briefing on artificial intelligence and the potential risks and opportunities in the context of terrorism and counter-terrorism. <https://www.un.org/securitycouncil/ctc/news/cted-holds-briefing-artificial-intelligence-and-potential-risks-and-opportunities-context>