VCDNP Vienna Center for Disarmament and Non-Proliferation

September 2025

Revolution or Evolution? How AI is Reshaping the WMD Proliferation Landscape

Natasha Bajema, PhD
Mara Zarka
With Minh Ly

Disclaimer: This research was sponsored by the Defense Threat Reduction Agency (DTRA) through the Strategic Trends Research Initiative as contract #HDTRA1-24-P-0040 between DTRA and President and Fellows of Middlebury College. The views expressed herein are those of the authors and do not necessarily reflect the official policy or position of DTRA, the U.S. Department of Defense, or the U.S. Government.

Authors

Dr. Natasha Bajema is a Senior Research Associate at the James Martin Center for Nonproliferation Studies at the Middlebury Institute of International Studies. She is an expert on nuclear and biological weapons and risk reduction, artificial intelligence, biotechnology, and other converging national security threats. Previously, she served as Director of the Converging Risks Lab at the Council on Strategic Risks. Dr. Bajema spent over a decade working on nuclear policy and emerging technology issues for the US government, including as a Senior Advisor at the Pentagon and the Department of Energy's National Nuclear Security Administration.

Mara Zarka is a Research Associate and Project Manager at the VCDNP, where her research addresses the intersection of emerging and disruptive technologies with nuclear, the security of nuclear and radiological materials against malicious non-State actors, and the non-proliferation regime and nuclear governance. Her work has also included projects on nuclear safeguards and peaceful uses of nuclear technologies, among others.

About the VCDNP

The Vienna Center for Disarmament and Non-Proliferation (VCDNP) promotes international peace and security by conducting research, facilitating dialogue, and building capacity on nuclear non-proliferation and disarmament.

The VCDNP is an international non-governmental organisation, established in 2010 by the Federal Ministry for European and International Affairs of Austria and the James Martin Center for Nonproliferation Studies at the Middlebury Institute of International Studies at Monterey.

Our research and analysis provide policy recommendations for decision-makers. We host public events and facilitate constructive, results-oriented dialogue among governments, multilateral institutions, and civil society. Through in-person courses and online resources on nuclear non-proliferation and disarmament, we train diplomats and practitioners working in Vienna and around the world.

Acknowledgements

The Defense Threat Reduction Agency (DTRA), through the Strategic Trends Research Initiative, sponsors strategic research projects under Broad Agency Announcement HDTRA1-24-S-0003. DTRA sponsored President and Fellows of Middlebury College to conduct this research examining emerging threats at the nexus of AI and WMD: understanding the implications of data and automation for development pathways. This report analyses the results of the one year technical study and offers reflections and guidance for policymakers at this nexus. The views expressed herein are those of the authors and do not necessarily reflect the official policy or position of DTRA, the U.S. Department of Defense, or the U.S. Government.

VCDNP
Vienna Center for Disarmament
and Non-Proliferation

Andromeda Tower, 13/1
Donau-City-Strasse 6
1220 Vienna
Austria

vcdnp.org
info@vcdnp.org
@VCDNP
VCDNP

Contents

As AI tools advance, possible intersections may emerge between AI and WMD, that need to be monitored and addressed by the international community.

Introduction

In an era where artificial intelligence (AI) is rapidly transforming every aspect of modern life – from how people communicate to how they conduct business – a pressing question emerges: How might these powerful new technologies affect the proliferation of the world's most dangerous weapons?

This paper examines the intersections between AI and weapons of mass destruction (WMD) that may emerge over the next five to ten years and provides an accessible guide for diplomats and policymakers to this nexus. While nuclear, biological, and chemical weapon systems have traditionally relied on physical materials and equipment, specialised facilities, and human expertise in the past, the AI revolution is changing the proliferation landscape in ways that demand attention from policymakers, diplomats, and global security professionals.

Recent advances in AI, particularly tools that can analyse vast amounts of data, automate complex processes, and generate new scientific insights have raised concerns among security experts around the world. For the past several years, technology leaders have warned policymakers that publicly available AI tools such as ChatGPT and others might lower the barriers for developing dangerous weapons. While some of these warnings may be overstated at times, they highlight a critical gap in understanding: diplomats and policymakers do not know enough about how AI might enable new pathways to WMD development or use for a broader range of actors.

The challenge is particularly acute because both AI and WMD are complex, technical fields spanning multiple scientific and engineering domains, some of which evolve rapidly and others which move at slower pace. Moreover, the concept of AI is so broad it encompasses everything from systems that predict weather patterns to chatbots that can write computer code. This wide range of capabilities defies gaining a concrete understanding of AI's full impact, especially when intersecting with the unique characteristics and development requirements of nuclear, biological, and chemical weapons. Understanding where these domains intersect (the AI-WMD nexus) and the implications for national and global security requires careful and comprehensive analysis.

Study Approach and Methodology

This paper is based on a year-long technical study that took a systematic and empirical approach to mapping potential risks at the AI-WMD nexus over the next five to ten years. The study examined the entire lifecycle of WMD development – from the initial research and development (R&D), acquisition, and production stages to weaponisation and delivery – to identify where AI might create new vulnerabilities or enhance capabilities for proliferators. The study was specifically designed to aid policymakers who are directly responsible for prioritising threats and allocating resources in an environment of persistent uncertainty and constraint. Thus, the project team examined near-term developments at the AI-WMD nexus that could affect national and global security rather than engage in speculative scenario building.

The role of data represented the key insight guiding the analysis: AI models require large volumes of relevant and representative training data to function properly. In other words, any element related to developing WMD must first be converted into raw data that computers can process and analyse for today's AI tools to have direct impact. Alternatively, manual processes related to WMD development can be automated with computers and machines in ways that AI can enhance. These two types of "digitisation" are already occurring in certain areas across the WMD development lifecycle – for example, biotech companies can convert the genomes of pathogens into digital information that can be stored on computers, processed using software, and shared over digital networks.

As a third mode of potential impact, emerging technologies may create unprecedented bridges between the AI and WMD domains. These technologies – including drones, additive manufacturing (3D printing), microfluidics devices, gene sequencing, synthesis, and editing – operate at the intersection of the physical and digital worlds, and can be enhanced by AI.

Unlike traditional manufacturing or equipment, they rely fundamentally on digital architectures. When enhanced by AI capabilities, these technologies could transform multiple stages on the WMD development pathway. For instance, AI-enabled 3D printers might optimise the production of specialised components for nuclear weapons, while gene editing technologies could accelerate the design of novel biological agents. Rather than AI directly enabling WMD development, these technologies may serve as the critical intermediaries – translating AI's digital capabilities into physical outcomes with potential weapons applications. As AI capabilities improve, these tools have the potential to indirectly lower technical barriers and create new pathways for WMD.

To understand how AI might affect the risk of WMD proliferation, the project team devised a structured expert elicitation process that took place over the course of a year. The process began with a focused literature review and preliminary interviews with subject matter experts (SMEs) to map the current landscape and identify key areas of concern at the AI-WMD nexus. This research informed the development of a comprehensive survey distributed to SMEs across the nuclear, biological, and chemical weapons domains.

Next, the survey captured SME assessments of how AI might impact each stage of weapons development (see the appendix for a brief overview of survey results). The project team also convened an interactive workshop where experts from government, academia, and industry explored these findings in depth, using forecasting techniques to envision plausible near-term scenarios and identify critical trends and drivers of change at the AI-WMD nexus. This multi-layered approach helped to build a comprehensive picture of emerging risks.

The methodological choice of prioritising insights from SMEs who actively advise or support government decisionmakers was deliberate: much of the available literature in the AI-WMD domain is currently speculative rather than grounded in empirical research or historical precedence. There are few historical cases or robust datasets on the integration of AI and WMD, a circumstance that constrains the potential of traditional evidence-based analysis. By combining survey data, in-depth interviews, and collaborative scenario-building, the project team developed expert insights that reflect consensus views and identify important areas of divergence among those who best understand the implications of these converging fields.

What is at Stake and What is the Objective of this Paper?

The implications of the AI-WMD nexus extend far beyond technical considerations to real-world consequences. If AI significantly lowers the barriers to WMD development, it could:

- Enable new proliferators, including terrorist groups, to pursue capabilities previously available only to nation-states
- Create new pathways for weapons development that bypass traditional controls and monitoring systems established by global regimes
- Accelerate the speed at which dangerous weapons can be developed and deployed
- Complicate international efforts to prevent proliferation and maintain strategic stability

Understanding the risks at the AI-WMD nexus is essential for developing effective global and national policies, strengthening international cooperation, and ensuring that the benefits of AI advancement do not come at the cost of catastrophic security risks. This paper offers a clear, structured guide for thinking about this obscure and emerging threat space, drawing on the latest insights from scientific, technical, and policy experts. The bottom line is that the AI-WMD nexus does not yet constitute a revolution but rather a critical period of evolution that demands careful attention and systematic monitoring.

The paper begins with a brief primer on what is meant by artificial intelligence (AI) and how to understand the current technological revolution. The main body offers three points of reflection, which serve as caveats to the "sky is falling" hypothesis found in much of the existing literature.

- First, the paper challenges a core assumption made about the utility of AI: Why would actors use AI to develop or use WMD? This fundamental and often neglected question shapes the entire debate yet is rarely examined directly.

- Second, the paper emphasises the distinction between proliferator capabilities and intent: In a resource-limited environment, simply increasing capabilities does not automatically lead to the development or use of WMD. Intent remains a decisive factor and can be mitigated with national and global policies.

- Third, the paper examines where AI is most likely to intersect with WMD – and where barriers persist: The development and use of WMD is a complex process that does not end with successful acquisition of dangerous materials. Delivery, operationalisation, and maintaining control require overcoming additional technical hurdles, many of which remain intact.

- The paper concludes with a near-term outlook of how this space may change over the next ten years and provides policy recommendations for what can be done now to mitigate future threats.

This paper aims to provide diplomats and policymakers with insights needed to navigate this complex challenge and protect global security in an age of rapid technological change.

There are several different types of artificial intelligence models and systems that are used for different purposes.

What is Artificial Intelligence?

Given the complexity of AI, the lack of shared understanding among stakeholders, and the rapid pace of technological change, it is necessary to begin this analysis with a brief primer. The persistent confusion and hype surrounding AI have made it increasingly challenging for policymakers, diplomats, and even technical experts to distinguish between genuine innovation, practical capabilities, and speculative risks. This challenge is particularly acute in the context of WMD where misunderstanding can have significant consequences.

In 2022, the release of OpenAI's ChatGPT unleashed a new era of AI for national and global security by making sophisticated models accessible to everyone, raising new questions about potential AI intersections with WMD, with most of the focus on new risks rather than opportunities. This year marked the introduction of generative AI tools to the public, which are often referred to as “frontier” or “foundation” models that have been trained on vast amounts of Internet data and possess general capabilities across unlimited domains.

Over the past few years, the rapid evolution of generative AI, in particular, has deepened the confusion among policymakers and diplomats about what AI is, what it can do, and what it might mean for the WMD domain. Advanced machine learning tools have been deployed across industries and government agencies for at least a decade, quietly transforming how policymakers analyse data, optimise processes, and make predictions. In other words, the AI revolution was already well underway before ChatGPT’s watershed moment.

The following sections provide the study's definition of artificial intelligence and an overview of different machine learning tools, along with three principles intended to guide thinking about the evolving relationship between AI and WMD.

Defining Artificial Intelligence

AI represents a general-purpose technology, much like electricity or the Internet before it. Its effects will ripple through every sector of human activity, creating non-linear and often unexpected impacts that are difficult to predict. The digital revolution has already laid the groundwork for this rapid transformation, creating interconnected systems and data flows that AI can now leverage and accelerate.

For the purpose of this analysis, AI refers to the science and engineering of making intelligent machines and software. At its essence, AI automates tasks previously done by humans. Over many decades now, computers have accomplished a growing number of tasks both faster and more accurately than their human counterparts, using various forms of AI embedded in operating systems and software.

Three mindset shifts on AI can help improve clarity in discussions about the nexus of AI and WMD.

AI is Not New

Though it seems counterintuitive, it is helpful to stop thinking of AI as a new phenomenon but rather as a natural continuation of past efforts to create intelligent machines and software. The field of AI emerged alongside the development of computers in the 1950s. For many decades since, computer programmers have been making machines more intelligent using various programming approaches, allowing them to perform complex tasks previously done by humans, including within the WMD domain. These developments are simply not called AI anymore.

Machine learning is a subfield of AI that emerged in the 1950s as a data-driven approach for making machines and software more intelligent by having them learn from data, identify patterns, and make predictions without being explicitly programmed to do so – representing a fundamental shift in how computers are programmed to solve complex problems:

- **Traditional programming**: In traditional programming, developers explicitly specify a set of logical rules and data structures for the computer to follow and produce desired outputs. To do this, the rules of a specific domain must be known in advance. That means humans need to have a full (or at least good) understanding of the problem before they can program computers to solve it.

- **Machine learning**: With machine learning, the programmer specifies the desired outputs of a model, provides huge volumes of training data, and offers basic instructions for data processing. However, in contrast to traditional programming, they do not give the computer an explicit set of instructions for generating specific outcomes. With only a few instructions on procedure, the AI model analyses patterns in the data and identifies statistical relationships; it determines rules that allow it to generate the requested outputs from new inputs. This approach is enormously powerful because humans do not need to understand every aspect of a complex problem in advance to program a computer to solve it.

However, machine learning advanced slowly due to major technological limitations. Given today's growing volumes of data, falling costs of data storage, and increased computing power, machine learning approaches have now gained significant traction, solving old and new problems alike. To be effective, massive volumes of relevant data are needed to support machine learning tools. In most cases, the quality of the training data is more important for predicting accurate outcomes than the algorithm itself. The effectiveness of a machine learning tool will depend on how relevant the data is for solving a problem and how well the data represents the problem's operational environment.

Diversity of AI Models and Approaches

The second mindset shift is to move beyond talking about the “impact of AI” on WMD to talking about specific tool types, model architectures, training data, and desired outcomes. This requires policymakers and diplomats to dive deeper and understand the basics of machine learning. Belying simplistic definitions, the AI tools being deployed today are incredibly diverse. These nuances are critical for understanding the risks of AI for WMD.

To assess the implications for WMD development and proliferation, it is helpful to think about machine learning tools as being conceptually categorised into two broad types of tools – predictive AI (or task-based AI) and generative AI. These two groups have some similarities: both are machine-learning tools built on deep neural networks, and both are trained on massive volumes of data.

However, they are trained on different data sources, built for different purposes, and thus, they work differently.

Predictive AI tools or "narrow AI" are specialised machine learning tools that exceed human capabilities on specific, well-defined tasks. They analyse patterns found within the data and produce evidence-based outcomes with a probabilistic level of certainty. This is similar to applying advanced statistics to large datasets to determine correlative and causal relationships using the scientific method. Such AI tools have been deployed in commercial and national security settings for more than a decade.

Predictive models are trained on large volumes of relevant, curated, and often proprietary data. To function correctly, such data must be of high-quality, broadly representative of the domain, and relevant to the problem being solved. These tools are expensive to develop in terms of infrastructure, data curation, and expertise, and their availability is often limited to proprietary stakeholders. Alternatively, they can be made available for commercial use, but often at high access prices for most users.

Generative AI tools are trained on Internet data to produce new data similar to a given dataset. These "general" or broadly capable tools include large language models (such as ChatGPT), diffusion models (image and video generators), and multimodal models (many models are now capable of handling different data types). These AI models identify patterns and trends within the training dataset and generate a set of rules about the relationships among the data (e.g., text, images, video, audio). Then, they produce novel outputs based on their training data. Sometimes, these outputs can be entirely fake or inaccurate (often called hallucinations). These tools are expensive to develop but are made available for public use for free or low fees.

Generative AI tools are often referred to as “frontier models”, which are defined as general-purpose models that demonstrate cutting-edge capabilities.

The key to getting frontier models (such as ChatGPT) to reduce their hallucinations and perform better in a specific domain is to provide them with external data sources (e.g., a proprietary database). This approach, often called retrieval augmented generation (RAG), helps the model focus on the most relevant areas of its training data when answering a query or prompt. Fine-tuning is another way to customise a generative AI model for a specific domain by retraining it with curated, domain-specific datasets. Both approaches require significant volumes of curated data and technical expertise (see Figure 1).

Understanding the difference between predictive and generative AI approaches is crucial for assessing proliferation risk because they enable different types of actors in different ways.

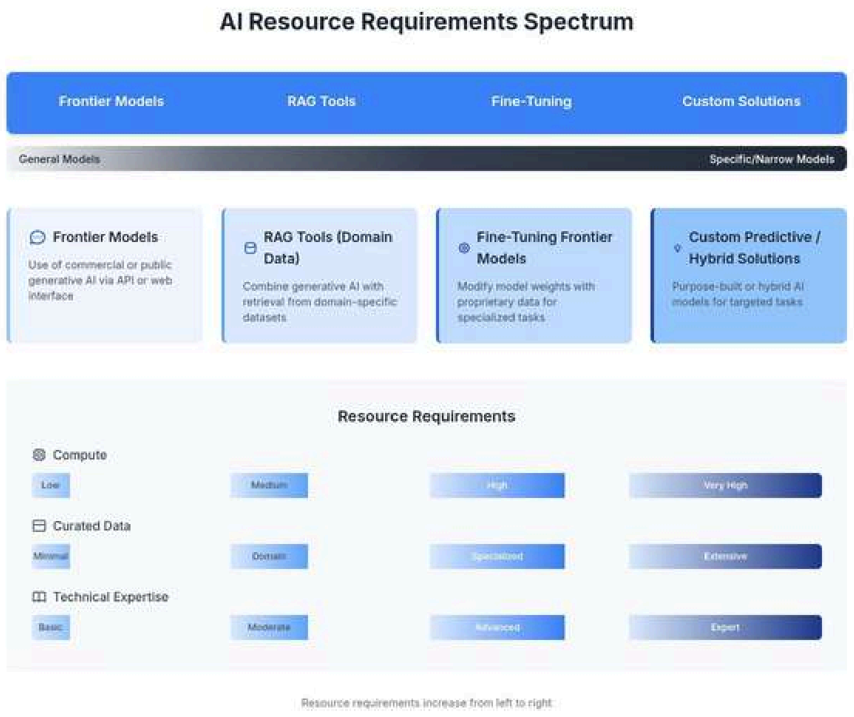
AI Resource Requirements Spectrum
Frontier Models
RAG Tools
Fine-Tuning
Custom Solutions
General Models
Specific/Narrow Models
Frontier Models
Use of commercial or public generative AI via API or web interface
RAG Tools (Domain Data)
Combine generative AI with retrieval from domain-specific datasets
Fine-Tuning Frontier Models
Modify model weights with proprietary data for specialized tasks
Custom Predictive / Hybrid Solutions
Purpose-built or hybrid AI models for targeted tasks
Resource Requirements
Compute
Low
Medium
High
Very High
Curated Data
Minimal
Domain
Specialized
Extensive
Technical Expertise
Basic
Moderate
Advanced
Expert
Resource requirements increase from left to right

Fig. 1: Overview of General AI Models and Resource Requirements for Different Approaches to Customise Models for Specific Domain Purposes[1]

Advanced nation-states with substantial resources are best positioned to develop custom AI models (using generative, predictive AI, or hybrid approaches) that require access to computing power, curated datasets, and technical expertise. These states can afford to collect specialised data, employ expert teams, and develop infrastructure, and build targeted systems that could accelerate specific aspects of weapons research, from materials science to warhead design optimisation (see Figure 1).

In contrast, non-state actors and less resourced states will likely rely on publicly available generative AI models and commercial off-the-shelf predictive AI tools. While these tools may lower technical barriers, they also have important limitations. For example, publicly available generative AI models are prone to hallucination, meaning they often produce plausible sounding, but factually incorrect or misleading responses. Effectively navigating and applying the outputs of these models requires a significant level of domain expertise. Without it, users may struggle to distinguish between accurate guidance and erroneous or even dangerous misinformation.

The conceptual distinction between generative and predictive AI is already blurring with new technological advancements. Recently, DeepMind announced the development of AlphaFold 3 that can make predictions with unprecedented accuracy in predicting a protein's 3D structure from its amino acid sequence. The system incorporates both predictive and generative AI architectures. Thus, a coming space to watch will involve innovative and powerful combinations of both categories of generative and predictive tools, leveraging their individual synergies.

[1] Figure 1 was created with the help of Canva AI Image Generator on 4 August 2025 using specific prompts on information to include and how to present it in a figure.

Perpetual Beta Mode

Despite their impressive capabilities, today’s AI models are prototypes and represent the earliest stage of development. Given rapid advancements in AI models, reaching meaningful assessments about the differential impact of AI on nuclear, biological, and chemical weapons presents another unprecedented challenge for risk mitigation. For example, the contrasting speed of the AI and nuclear domains can lead to dangerous complacency about AI risks among nuclear experts – especially given the current geopolitical environment and potential consequences of a nuclear war.

Generative AI will continue to advance rapidly with an increasing focus on agentic AI or autonomous agents. Large language models such as advanced versions of ChatGPT have already exhibited impressive abilities in reasoning, making plans, engaging in self-reflection, and refining their processes. Developers are interested in harnessing these abilities for implementing autonomous actions in the digital world – the capability to perform tasks on applications independently without human intervention. At a basic level, agentic AI refers to when agents powered by large language models iterate, engage with the Internet, and access external web applications, and leverage these tools to carry out actions autonomously on behalf of humans.

The beta nature of AI development and rapid pace of advancement makes it difficult to assess their implications for the WMD domain. What is certain, however, is that the landscape will continue to evolve rapidly, requiring policymakers and diplomats to develop flexible frameworks for understanding and responding to emerging risks.
With this foundation in AI concepts established, a more nuanced examination of the AI-WMD nexus becomes possible. The following sections offer three critical points of reflection that challenge simplistic narratives about AI's impact on WMD proliferation. These perspectives provide essential context for policymakers and diplomats navigating this complex intersection of emerging technology and global security.

It is important to monitor how AI might lower barriers or create new pathways for WMD proliferation as well as clarify the limits of what AI can currently achieve.

Assessing the Value Proposition of AI for WMD Development

The first critical point in evaluating the AI-WMD nexus requires stepping back from an attitude of technological determinism to examine fundamental questions about utility and purpose. Rather than assuming AI will automatically accelerate proliferation, a more sophisticated analysis must consider why and under what circumstances proliferators might choose to deploy these technologies in weapons development programmes.

One of the most persistent misconceptions in the media is that AI functions as a kind of magic wand: an all-purpose tool that can instantly solve any problem, including enhancing and accelerating the development of WMD. This assumption, fuelled by headlines and speculative commentary, risks obscuring the reality of how AI tools might be used, what their limitations are, and the practical challenges that would confront any proliferator – state or non-state actor – attempting to harness AI to support their WMD activities.

What Problem Is AI Actually Solving?

To move beyond the hype, it is crucial to ask a simple but overlooked question: What problem is AI being applied to, and is it the right tool for the job? As discussed above, machine learning, the dominant AI technique today, offers a fundamentally new approach to problem-solving with computers and machines, but its real-world value depends on the nature of the problem to be solved and the existence of relevant data.

Many complex scientific and engineering challenges associated with WMD development have already been solved and honed with traditional methods over many decades. In some cases, existing solutions are highly effective, and improvements with AI may offer only marginal gains. In others, the core barriers are not technical, but practical or political: acquiring rare materials, building specialised facilities, or recruiting experts. Before assuming that AI will automatically accelerate WMD proliferation, it is necessary to ask:

- Do proliferators need better solutions?
- Are existing methods good enough?
- Or, are there genuine knowledge or capability gaps that AI could help fill for different proliferators?

For example, would a state actor or non-state actor use AI to develop a novel design, or simply use tried-and-true approaches? Would AI be deployed to ideate new solutions, close knowledge and capability gaps, or merely improve the efficiency of established processes? The answers depend on the proliferator's resources, risk aversion, the availability of data, and the specific challenges they face.

Assessing the Costs and Benefits of Using AI for WMD Development

Before investing in AI, any potential proliferator – like any rational actor – would weigh the costs and benefits of different approaches. Key questions would include:

- Have traditional methods already solved the problem?
- If not, can the desired outcome be achieved without AI?
- Would AI offer a meaningful improvement to an existing solution, or just a costly complication?
- Is the "return" on AI investment – whether in novelty, speed, secrecy, or effectiveness – worth the effort?
- What risks are involved with achieving a new or enhanced solution using AI?
- What types of technical expertise will be required to achieve an AI-enabled solution?
- If proliferators depend on publicly available AI models, are results from the frontier models reliable?

Crucially, the effectiveness of AI depends on the quality and relevance of available data. As discussed earlier, even the most sophisticated machine learning model is only as good as the training data it receives. If the data does not accurately represent the domain or task in question, the AI's outputs will be unreliable – or even dangerously misleading. In the WMD context, relevant, high-quality data is often scarce, proprietary, classified, or simply unattainable for most proliferators.

Moreover, the technical expertise required to select, curate, and label data, to train and fine-tune models, and to interpret AI outputs, is not trivial. As with any advanced technology, the skill of the user will determine the extent of AI's utility.

As proliferators move from frontier models and off-the-shelf tools toward more customised and hybrid AI solutions, the resource costs and requirements for a return on investment rise sharply. Key enablers – including computational infrastructure, access to curated data, technical expertise, physical infrastructure, and even reliable energy sources – become deciding factors in reaching this decision. In many cases, these resource requirements may act as real-world barriers, limiting which proliferators can realistically exploit AI for WMD purposes.

Possible Use-Cases of AI for WMD Development

If proliferators choose to leverage AI for WMD development, the available use-cases – and the value they derive – will depend heavily on whether they employ predictive AI, generative AI, or hybrid approaches (see appendix for a comprehensive listing of potential use cases):

- **Predictive AI models** excel at recognising patterns in existing data, forecasting evidence-based outcomes, and classifying information. Their architectures make them well-suited for optimising design trade-offs, detecting anomalies in acquisition processes, simulating production and delivery, and enhancing the reliability and safety of complex systems across the nuclear, biological, and chemical domains. For example, predictive models can streamline fissile material production and increase yields, anticipate bioreactor failures in biological agent production, or forecast the environmental stability of chemical agents – of course, provided that high-quality, representative data is available.

- **Generative AI models**, by contrast, can create entirely new content: from novel weapon designs to synthetic research protocols, deepfake communications, and alternative synthesis routes for dangerous agents. These models are particularly potent at the ideation stage (i.e., through digital simulations), generating new pathogen variants or novel toxins in the biological realm, or designing chemical compounds engineered to evade controls and detection. Generative AI also introduces new risks in acquisition and delivery, such as crafting persuasive phishing campaigns, producing false credentials, or scripting disinformation operations. While generative AI models can accelerate innovation and lower barriers to entry, they are also prone to “hallucinations” or misleading outputs, especially in data-scarce, high-stakes environments like WMD development. Proliferators will also have to overcome established guardrails that protect frontier models from releasing dangerous information related to WMD.

As discussed in the AI primer, the line between predictive and generative AI continues to blur, with many contemporary tools displaying both capabilities depending on how they are prompted or fine-tuned.

However, all AI use-cases share a critical caveat: without access to quality, relevant, and representative data, even the most advanced AI systems may be ineffective or dangerously unreliable – a case of “garbage in, garbage out”. Ultimately, the impact and risks of AI in WMD proliferation are determined not just by model architecture, but by the data and expertise underpinning their application at each stage of the weapons development lifecycle.

Reality Check: Data, Expertise, and Resources Still Matter

The assumption that AI will immediately empower nefarious actors to develop WMD overlooks the practical constraints that shape real-world decision-making. Access to sophisticated AI frontier models does not erase the need for domain data, expertise, infrastructure, and investment to support the entire development lifecycle of nuclear, biological, and chemical weapons – from R&D, acquisition, and production, to weaponisation and delivery. For many proliferators, traditional methods may remain more attractive, cost-effective, or feasible than attempting to use, build, or adapt advanced AI systems for WMD development. Thus, in the near-term, they may not turn to AI at all.

While it is vital to monitor how AI might lower barriers or create new pathways for WMD proliferation, it is equally important to clarify the limits of what AI can currently achieve and not achieve – and to avoid magical thinking that exaggerates risks. Only by anchoring analysis in these practical realities can policymakers and diplomats craft effective strategies for managing the intersection of AI and WMD in the years ahead.

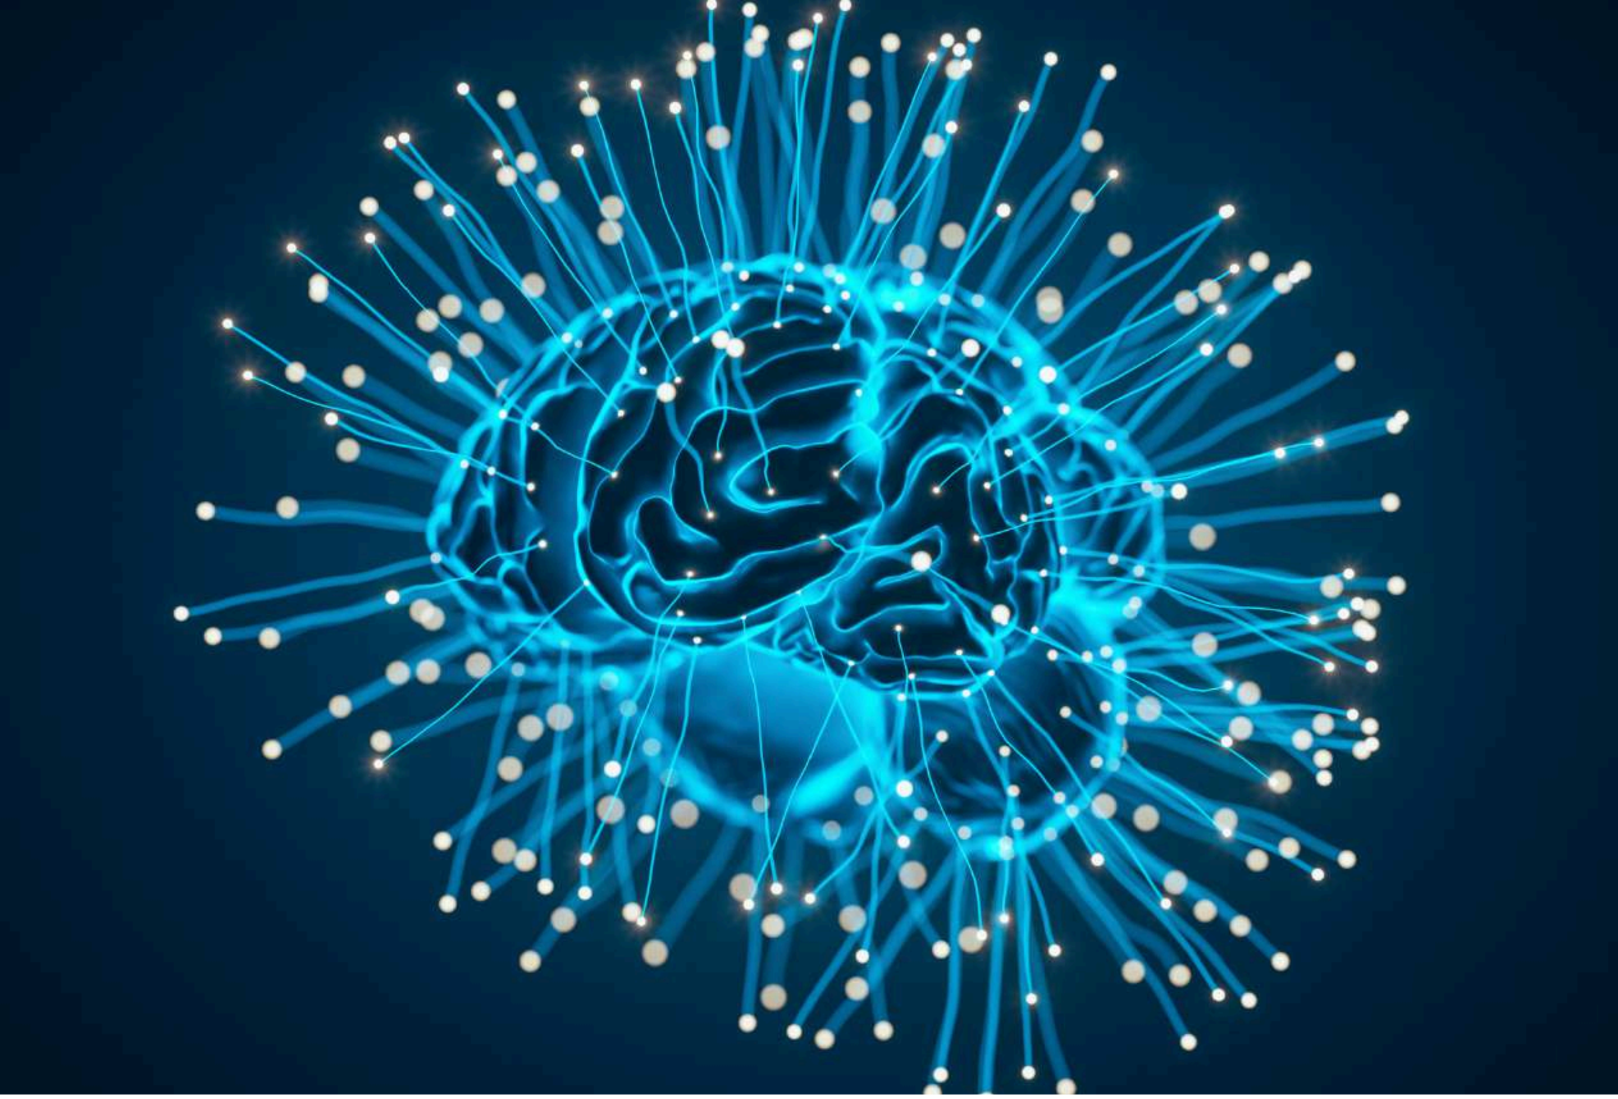

It is important to look beyond technological risks and consider the actors capabilities and intent for using technologies for nefarious purposes.

Actor Capabilities, Intent, and AI-WMD Risk Assessments

The second point of reflection challenges a prevalent analytical shortcoming in current discussions of the AI-WMD nexus: the tendency to focus almost exclusively on technological capabilities while neglecting the critical role of intent. In a world of limited resources at both the national and global levels, diplomats and policymakers are compelled to make difficult decisions about how to prioritise among a growing array of threats. For this reason, it is vital to systematically analyse all potential risks at the AI-WMD nexus rather than to become overly consumed by only one angle.

Based on recent literature, it seems prudent to direct most resources and attention toward the powerful intersection of AI and biological weapons – the so-called AI-bio nexus. This rapidly growing body of work frequently highlights the proliferation risks posed by advanced AI models, the accessibility of frontier models, and the potential for novices to rapidly enhance their capabilities to develop biological weapons. But this raises a critical question: Can policymakers and diplomats be sure they are allocating limited resources effectively? Are they potentially overlooking equally significant risks at the intersections of AI with nuclear and chemical weapons domains?

Understanding Capabilities versus Intent

Much of the discourse at the AI-bio nexus is preoccupied with two interconnected concepts related to technical, scientific, and operational capabilities: “uplift” and “raising the ceiling”. Uplift refers to AI’s potential to elevate less skilled or novice proliferators – giving them access to skills and knowledge that would otherwise be out of reach. Raising the ceiling, on the other hand, describes AI’s ability to enhance the capabilities of already skilled, sophisticated, or well-resourced proliferators, such as advanced nation-states or highly proficient non-state groups. The focus on the capability spectrum enabled by AI drives much of the anxiety about how AI might reshape the threat landscape for biological weapons.

Current assessments of the AI-bio nexus tend to blur the essential distinction between capabilities (what proliferators can do with AI's assistance) and intent (what proliferators choose to do, driven by their motivations and objectives). This distinction is not merely academic, it is fundamental for meaningful threat assessments and the development of effective policy. Without careful attention to both capability and intent, there is a danger of misjudging the true nature and scale of emerging WMD risks in the age of advanced AI.

By ignoring the role of intent, the current literature distorts risk assessments in two important ways: first, it amplifies concerns about the AI-bio nexus to sometimes sensational levels; and second, it diverts attention from other equally consequential intersections between AI and other WMD domains, such as nuclear and chemical weapons. To clarify the logic behind potential gaps that emerge, it is helpful to examine the drivers behind the current emphasis on biological compared to nuclear weapons.

The biological weapons domain is fundamentally different from the nuclear weapons domain. Unlike nuclear weapons development, which depends on weapons-usable fissile material, expansive infrastructure, and highly regulated supply chains, biological weapons development is far more enmeshed with commercial sectors. The biotechnology sector is advancing rapidly, propelled by commercial innovation, a thriving ecosystem of startups, academic research, and service providers. This environment creates substantial dual-use overlap, as many of the same technologies that are essential for legitimate medical or agricultural progress can, in theory, be repurposed for malicious ends.

These dynamics have fuelled narratives that the life sciences are “becoming easier”, “more automated”, and “more plug and play”. Some experts argue that this trend increases the likelihood that non-state or lone-wolf actors might attempt to develop bioweapons. The perceived democratisation of advanced science feeds concerns that AI will further accelerate this trend, enabling individuals or small groups to conduct complex experiments, synthesise dangerous pathogens, or bypass traditional forms of oversight with unprecedented ease.

The experts interviewed for this study widely believe that AI does not fundamentally alter the calculus for biological weapons proliferation (i.e., the intent to develop and use) for most actors, even if it lowers some barriers for less sophisticated proliferators. Documented cases of biological weapons proliferation are rare, and for state actors, they are almost always shrouded in secrecy.

A common critique among the experts interviewed for this study is that risk analysts tend to assume above-average novelty on bad actors, when history shows that most non-state actors and many state actors tend to rely on tried-and-true methods, sticking with what they already know rather than taking on the risks and uncertainties of untested approaches. Innovation, especially for nefarious purposes, is extraordinary rather than routine.

In contrast to bio, the nuclear domain is more tightly regulated, less commercialised, and advances more slowly due to its immense physical, technical, and financial requirements. Yet this does not mean the intersections between AI and nuclear weapons are negligible. Unlike for bioweapons, several nuclear-armed states are already experimenting with or deploying AI-enabled systems to support critical aspects of command and control, such as early warning and targeting. These developments could have profound real-world implications, increasing the risk of crisis escalation or miscalculation under uncertainty.

Focusing most attention on the AI-bio nexus thus risks undermining a broader understanding of the impact of AI on WMD threats, potentially leaving other risks underexplored. This is not to downplay the significance of the AI-bio nexus – especially in light of lessons from the recent global pandemic, which underscored the importance of vigilance in the life sciences. Rather, it is a call for analytical balance: while the AI-bio intersection is undeniably important, it is essential to ensure that prevailing biases do not obscure or neglect equally consequential risks in the nuclear or chemical domains.

Bottom Line: Capabilities Are Advancing, But Intent Remains the Linchpin

When reaching decisions, policymakers and diplomats need to assess threats by looking beyond mere technical risk – namely, what is possible with a given technology – to consider both the capability of potential proliferators, and their intent to develop and use those capabilities for harm. This distinction is crucial in the WMD context, where the consequences of misjudging the threat can be catastrophic, but the costs of overreaction can also be significant.

While AI may make it technically easier to design a novel toxin or optimise a warhead, the actual threat materialises only when proliferators choose to pursue these ends. Historically, this convergence of WMD capability and intent has been rare. Even in the case of biological weapons, where the information to inflict harm has grown more accessible over the past three decades, few non-state actors have shown an inclination to cross the line from interest to tangible action.

The decision to develop WMD is also rarely considered in isolation, especially by non-state actors. Proliferators often weigh the utility of nuclear, biological, or chemical weapons against more conventional alternatives, which are typically much easier to acquire, use, and control for specific effects. This competition with conventional options can limit the perceived utility, and thus the intent to pursue WMD, even in a world of rapid AI advancement.

The gap between capability and intent is not uniform across all types of proliferators. Experts interviewed for this study found that state actors with robust technical expertise and resources stand to gain the most from integrating AI into WMD development pathways, especially when leveraging custom solutions and proprietary or classified datasets. For these actors, AI can accelerate research timelines, optimise production processes, and refine delivery mechanisms – potentially reducing detection risks, especially when paired with complementary technologies like robotics or additive manufacturing. However, the need for a return on investment for using AI remain in play.

Yet, even for states, intent is not dictated by technology alone. It is profoundly shaped by political will, strategic calculation, and risk tolerance. Human factors can serve as persistent barriers to the effective integration of AI into WMD programs. Bureaucratic inertia, organisational silos, and the challenges of coordinating between cyber and physical science disciplines may slow or even derail ambitious projects. The expertise required to develop advanced AI systems is distinct from that which is needed in chemistry, biology, or nuclear engineering, and these skillsets are rarely housed within the same teams or institutions. Such differences in technical background, work culture, and institutional priorities make the seamless convergence of the AI and WMD domains much more difficult in practice than it may appear on paper.

For non-state actors and less capable states, the landscape is even more constrained. Here, AI's main value lies in lowering informational and basic capability barriers – helping proliferators to quickly access open-source knowledge or understand standard protocols.

However, the leap from online information to tangible action remains daunting. Practical hurdles – such as obtaining specialised equipment, developing tacit hands-on skills, ensuring safety or security, and navigating complex supply chains – remain critical challenges. AI cannot “teach” the practical expertise or remove operational risks inherent in WMD development. As many experts interviewed for this study observed, if an actor needs AI to answer basic technical questions, they are probably not in a position to implement the answers effectively in the real world.

Effective threat assessments require linking a particular proliferator’s intent to a specific weapon type and a concrete objective or mission. It is not clear what objectives WMD may fulfil in today’s evolving security environment. Even among state actors, motivations for WMD development and use have shifted over time. In the biological domain, for instance, there has been a move away from Cold War visions of battlefield use toward more covert, targeted, or deniable operations. Authoritarian regimes might see AI as a tool for domestic repression or for grey-zone tactics that fall short of open warfare, rather than for building arsenals aimed at achieving mass destruction.

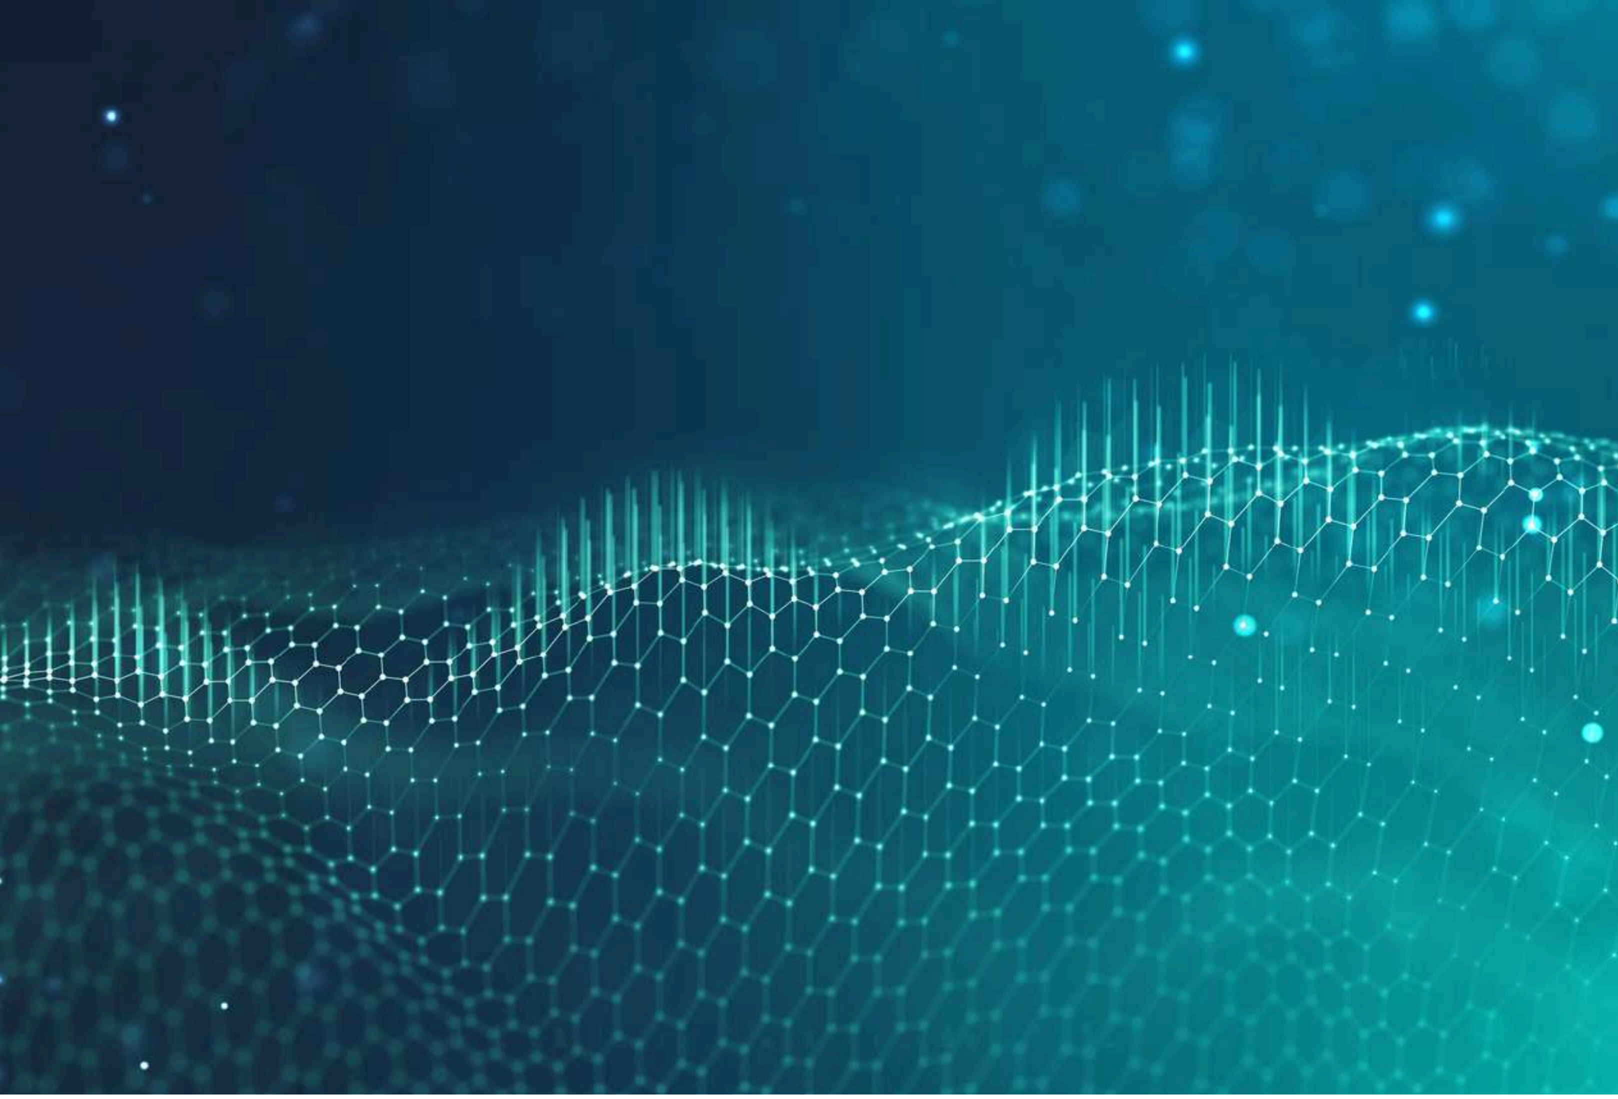

The impact of AI varies across the different stages of WMD development.

Divergent Impact of AI Across the WMD Development Lifecycle

The third point of reflection examines how AI's impact will vary significantly across different stages of WMD development, challenging the notion that these technologies will enable proliferators to develop and use WMD because AI tools provide them critical assistance at the early stages. To understand the true impact of AI on the proliferation landscape, it is essential to look beyond broad generalisations and examine how AI might affect each stage of the WMD development process. Unlike many fields where digital technologies have rapidly transformed practices, the development of nuclear, biological, and chemical weapons remains deeply rooted in the physical world. These processes rely not only on specialised materials and equipment, but also on the tacit knowledge and hands-on expertise of scientists, engineers, and technicians.

The influence of AI will vary both by weapon type – nuclear, biological, or chemical – and by the stage in the development pathway. This distinction is crucial: successful proliferation requires mastery of all stages, not just one. Each phase of the pathway demands a different combination of scientific, technical, logistical, and engineering competencies. Moreover, leveraging AI for any of these stages brings its own specialised requirements and expertise, which are not easily substituted for other more traditional skills.

While it is common to assume that the addition of AI will make proliferation easier, a closer look is needed to test this assumption. To provide clarity, it is helpful to first define the major stages along the WMD development pathway:

- **Research and Development**: This stage encompasses the creative work aimed at expanding the body of knowledge relevant to nuclear, biological, and chemical weapons. It includes both fundamental research and the practical application of new findings toward the pilot and commercial-scale production of materials, devices, and processes necessary for WMD development.

- **Acquisition**: At this stage, proliferators seek to obtain the source materials, specialised equipment, enabling technologies, and production infrastructure required for WMD development. Overcoming barriers to acquisition is often a significant challenge.

- **Production**: At this stage, the focus shifts from merely acquiring materials and equipment to scaling up production activities. This may involve moving from pilot-scale operations to the mass production of critical materials and components.

- **Weaponisation**: In this stage, produced materials are converted into a usable weapon. The specific technical requirements for weaponisation differ considerably across the nuclear, biological, and chemical domains, and may involve steps such as agent stabilisation, encapsulation, purification, weapon design, fabrication, assembly, and testing.

- **Delivery**: In the last stage, proliferators must develop or acquire reliable means to deliver the weapon to its intended target. The complexity and sophistication of delivery systems vary according to weapon type and the desired operational effect.

Examining the intersection of AI with each of these stages – and the enduring barriers that exist – offers a more nuanced and realistic picture of how, where, and if AI might alter the WMD proliferation landscape.

Opportunities of AI for WMD and Persistent Barriers

In the near-term, the impact of AI will be uneven across the development stages for nuclear, biological, and chemical weapons. To understand what is – and is not – changing, it is essential to look at how AI may assist each stage of WMD development, and where fundamental barriers remain.

Nuclear Weapons: Efficiency, Not Accessibility

Nuclear weapons remain the hardest to acquire and develop. The core barriers are physical, expensive, and deeply technical:

- **Research and Development**: AI can help design components, run simulations, and analyse large datasets, potentially lowering the knowledge barrier for some aspects of nuclear weapons design. It can also assist with optimising enrichment processes or improving the safety and reliability of existing stockpiles.

- **Acquisition and Production**: AI may help find efficiencies in component manufacturing, maintenance, and supply chain management, and can aid in covertly sourcing or diverting materials. Additive manufacturing (3D printing) enhanced by AI could help streamline component production for new proliferators.

- **Weaponisation and Delivery**: AI can enhance the design of the “physics package” (the core of a nuclear weapon), improve missile guidance and targeting, and support decision-makers by evaluating scenarios and response options more quickly. AI can also integrate data from multiple sources for early warning, intelligence analysis, and operational planning, giving nuclear-armed states potential advantages in time-sensitive situations.

For nuclear weapons, the fundamental challenge remains acquiring and processing fissile material—a task that requires vast resources, infrastructure, and technical expertise. AI cannot produce weapons-usable uranium or plutonium, nor can it replace the need for hands-on learning and experience in assembling and maintaining complex nuclear systems. Human oversight and risk aversion remain paramount, and the overall cost and risk of a nuclear weapons programme are not significantly reduced by AI.

Biological Weapons: New Efficiencies, Old Hurdles

Biological weapons are unique for their reliance on living organisms, making their development process especially complex and unpredictable. AI's greatest potential for bioweapons is in the early stages of the development pathway:

- **Research and Development**: AI can help generate new ideas, design novel proteins or toxins, and sift through vast scientific databases. For example, AI tools can predict how a virus might mutate or help design proteins with specific properties. This makes scientific exploration faster and more accessible, potentially lowering the expertise needed to get started.

- **Acquisition and Production**: AI can help less-experienced proliferators identify what materials or equipment they need and where to obtain them, consolidating information and even suggesting procurement strategies. AI-enabled automation can optimise laboratory processes and scale up production.

- **Weaponisation**: The practical challenges of growing, stabilsing, and weaponising living organisms remain significant. Notably, AI does not solve the problem of keeping agents alive and effective for dispersal. Given the lack of dual-use applications, little relevant data is available to support this phase.

- **Delivery**: AI can assist with planning—predicting weather, choosing optimal times, or even piloting drones for delivery. Still, the unpredictable nature of biology and environmental factors means that real-world deployment is fraught with risk and uncertainty. As an important factor for consideration, AI could generate disinformation, spreading panic, manipulating news narratives, and amplifying the psychological impact of even small-scale attacks.

Despite these advances, several barriers remain. Tacit knowledge, i.e., hands-on experience that cannot be learned from a textbook, remains essential for success. Data quality is often poor, and AI tools can make mistakes or suggest ideas that are technically unworkable or dangerous to the user. Physical acquisition of dangerous materials remains tightly controlled, and scaling up production is challenging. Most experts interviewed for this study agree: AI may make it easier for sophisticated proliferators to push the boundaries, but it does not fundamentally transform the risk landscape for less-skilled groups.

Chemical Weapons: Incremental Gains, Not a Revolution

Chemical weapons are less reliant on living systems but still depend on access to precursor chemicals, industrial equipment, and safe handling:

- **Research and Development**: AI can accelerate the discovery and design of novel chemical agents, predict toxicity, and even find alternative synthesis routes using open-source information. This can help actors bypass some traditional knowledge barriers, and AI can also help search historical records for forgotten and discarded methods.

- **Acquisition and Production**: AI can identify suppliers and design synthetic routes to circumvent export controls or legal barriers. It may also help actors source equipment more efficiently. AI, especially when combined with robotics, can automate lab tasks and optimise reactor conditions. It can improve process efficiency but cannot eliminate risks associated with handling toxic chemicals.

- **Weaponisation**: AI can help design and optimise devices or methods for turning chemicals into effective weapons. For example, generative AI tools can assist in engineering dispersal devices, determining the best way to aerosolize an agent or modify commercial equipment for weaponisation. However, achieving the right purity, stability, and dispersal characteristics still requires significant hands-on expertise and experimentation.

- **Delivery**: AI can support targeting and operational planning by analysing weather conditions, crowd density, and vulnerabilities in target sites. It can also assist in optimising routes for drones or other delivery systems and in predicting the spread of chemicals in various environments. As an important factor for consideration, AI could produce disinformation, spreading panic, manipulating news narratives, and amplifying the psychological impact of even small-scale attacks.

Tacit knowledge and hands-on skills are still crucial for handling dangerous chemicals. Acquiring large quantities of precursors remains challenging, and the risks of accidental exposure or detection are high. For many traditional agents, existing methods are already effective, so there is little incentive to innovate. AI may marginally lower the barrier for new proliferators, but it does not fundamentally change the calculus for most.

The Limitations of AI in the Physical World of WMD

While AI has dramatically accelerated the pace and scope of digital research and development (R&D) for nuclear, biological, and chemical weapons, the transition from virtual simulation and design to physical reality of producing and deploying weapons remains a major limiting factor. AI is exceptionally powerful in the digital domain – modelling, simulating, and optimising potential designs with unprecedented speed and precision. However, the journey from a digital concept to a functional, deployable weapon is long and fraught with challenges that AI alone cannot overcome.

Digital Capabilities: Where AI Excels

Within a fully digitised research environment, AI can revolutionise the early stages of WMD development. Its strengths include:

- **Modelling and Simulation**: AI can create highly sophisticated models of everything from protein folding to nuclear detonations, allowing for detailed exploration of possibilities before any physical work begins.

- **Rapid Iteration and Optimisation**: Algorithms can quickly test countless design variations, identifying optimal solutions far faster than traditional methods.

- **Efficient Analysis and Virtual Prototyping**: AI can sift through vast amounts of data, propose novel approaches, and support virtual tests, all within a seamless digital ecosystem.

These capabilities allow for faster, cheaper, and more creative research than ever before. In theory, many early barriers to entry are lowered for proliferators who can operate in this digital space.

Physical Realities: Where AI Hits Its Limits

However, digital capabilities only go so far. The physical world imposes its own set of challenges on developing biological, chemical, and nuclear weapons – ones that are not easily solved by algorithms, no matter how advanced:

- **Simulation ≠ Experimental Validation**: A promising simulation may not survive the harsh realities of an actual experiment in the real world. Biological systems behave unpredictably, chemicals may react in unforeseen ways, and nuclear processes are notoriously sensitive to minute errors.

- **Virtual Testing ≠ Real-World Behaviour**: The leap from virtual to physical often reveals flaws that digital models cannot predict. Environmental variables, equipment tolerances, and the sheer complexity of real-world systems frequently cause digital designs to fail in practice.

- **Design Capability ≠ Production Ability**: Having a theoretically perfect design is not the same thing as being able to manufacture an agent or material at scale, safely and reliably.

- **Theoretical Knowledge ≠ Practical Execution**: Knowing how a process should work is different from making it work in a laboratory or on a production line, especially under constraints like secrecy, resource limits, and time pressure.

The Enduring Importance of Tacit Knowledge

A central limitation for AI is its inability to fully capture or replace tacit knowledge – the hands-on expertise developed through years of practical experience. Tacit knowledge encompasses:

- Skills and judgment that cannot be fully written down or digitised.
- Intuitive problem-solving developed through repeated trial and error.
- The "art" of adapting to unexpected setbacks or improvising solutions with imperfect information.

AI can assist by codifying expert insights and even suggesting troubleshooting steps. Yet, it cannot replicate the dexterity, intuition, or adaptive problem-solving that experienced technicians, scientists, and engineers bring to the table. Most critically, AI cannot substitute for the physical skills needed to carry out delicate or complex tasks under uncertain conditions.

Currently, AI may reduce – but not eliminate – the need for human expertise. However, according to the experts interviewed for this study, its greatest value is realised when paired with skilled human operators who can interpret results, spot errors, and bridge the gap between digital output and physical realities. In some cases, the introduction of AI may even create new demands for tacit knowledge, as teams learn to work with advanced systems and adapt to new workflows.

Looking ahead, circumstances may evolve rapidly. As AI further integrates with robotics and potentially functioning humanoid machines, its ability to perform hands-on tasks could increase. If AI frontier models eventually master the foundational laws of physics, biology, and chemistry, their role in WMD development could expand dramatically, encompassing more stages of the pathway. Finally, if frontier models match or exceed human knowledge in all domains, AI may begin to advance science on its own, accelerating the reduction in tacit knowledge required for developing WMD. However, until then, the limitations of AI in the physical world and tacit knowledge remain critical barriers.

Given the rapid pace of technological change it is important to remain vigilant, adaptable, and informed on the AI-WMD nexus.

Future Outlook and Recommendations

The Evolutionary, Not Revolutionary, Impact of AI on WMD

The analysis presented throughout this paper leads to a critical conclusion: the impact of AI on WMD development is not yet revolutionary, but rather evolutionary. While AI technologies are indeed transforming many aspects of the development lifecycle for nuclear, biological, and chemical domains, they have not fundamentally altered the core barriers to proliferation. Instead, AI is gradually reshaping the proliferation landscape, introducing new efficiencies and capabilities that require careful monitoring and management.

Over the next decade, it will be critical to monitor and track key trends and drivers – in particular, the creation of curated datasets and AI's advancement of relevant commercial sectors. Commercial interests will continue to drive AI innovation in ways that create dual-use technologies, but this overlap will be greatest in the chemical and biological domains and significantly less pronounced in the nuclear sector.

The rapid pace of AI development means that today's assessment could change quickly. Policymakers and diplomats must remain vigilant, adaptable, and informed as they navigate this complex intersection.

Policy Recommendations

Based on the findings of this study, this paper offers the following recommendations for policymakers, diplomats, and security professionals:

1. Enhance Education and Training
Sponsor comprehensive training and education programmes on AI for policymakers and diplomats focused specifically on the risks and opportunities AI presents for proliferation and non-proliferation efforts. These programmes should demystify AI technologies while providing practical frameworks for risk assessment and management.

2. Foster Cross-Sectoral Collaboration
Form cross-sectional communities of interest that include participation from industry, academia, government, and civil society. These collaborative networks can help identify emerging risks, share best practices, and develop responsible innovation frameworks that balance security concerns with technological advancement.

3. Establish Baselines and Conduct Regular Assessments
Support research efforts to establish a baseline for WMD-relevant capabilities across diverse frontier models and evaluate a wide range of models frequently to measure potential "uplift" in capabilities. This ongoing assessment will help identify concerning trends before they become acute security risks.

4. Develop Guardrails for AI in Nuclear Systems
Assess the risks of AI integration into nuclear weapons systems and build international norms around appropriate guardrails for the role of AI. This should include consensus-building around meaningful human control, verification measures, and crisis stability mechanisms.

5. Strengthen Public Health Infrastructure
Reinforce public health systems and pandemic response capabilities as a critical defence against potential biological threats, whether naturally occurring or deliberately engineered. Robust detection, surveillance, and response systems provide resilience regardless of the source of biological threats.

6. Enhance Nuclear Safeguards and Export Controls
Strengthen existing nuclear safeguards and export controls to account for the ways AI might be used to circumvent traditional barriers. This includes updating technical controls and monitoring systems to detect AI-enhanced proliferation efforts.

7. Leverage AI for Non-Proliferation
Explore how AI tools can support non-proliferation efforts through enhanced monitoring, verification, and compliance activities. The same technologies that could aid proliferators may also provide powerful new capabilities for the non-proliferation community.

8. Develop Technical Solutions
Invest in research to develop technical solutions that can help mitigate specific risks at the AI-WMD nexus, such as tools to detect AI-generated disinformation related to WMD threats or systems that can identify suspicious patterns in procurement activities.

Conclusion

The nexus of AI and WMD presents complex challenges that require nuanced and adaptive approaches. By focusing on the practical realities of how AI might influence WMD proliferation pathways – rather than speculative worst-case scenarios – policymakers and diplomats can develop more effective strategies for managing emerging risks.

The good news is that many of the traditional barriers to WMD development and use remain intact, even in an era of advanced AI. Physical constraints, resource requirements, and the need for tacit knowledge continue to limit how quickly and easily dangerous weapons can be developed and deployed.

Nevertheless, the rapid evolution of AI capabilities demands sustained attention and coordinated international action. By implementing the recommendations outlined above, the global community can work to ensure that advances in AI enhance security rather than undermine it. Through careful stewardship of these powerful technologies, policymakers and diplomats can navigate the AI-WMD nexus with wisdom, foresight, and a commitment to preserving international peace and security.

Experts broadly agreed that AI is an accelerator, not an equaliser when assessing its impact on WMD development.

Appendix: AI and WMD Development Pathways - Expert Assessments

The following appendix synthesises findings from the one-year study on the nexus of AI and WMD. The study provides the foundation for the analysis presented in the main paper, and this brief overview offers policymakers and diplomats detailed insights into how subject matter experts view the evolving risk landscape.

Overview of Survey Results

The subject matter expert (SME) survey was structured to assess the perceived impact of AI across all stages of WMD development pathways for nuclear, biological, and chemical weapons. The survey comprised approximately 150 questions, including both quantitative Likert-scale items and open-ended prompts to capture qualitative insights, and generated over 200 pages of results. Each section included questions on R&D, acquisition, production, weaponisation, delivery, operational planning, and disinformation for each weapon type. Respondents were routed to relevant sections based on their expertise and interests, ensuring depth without overburdening any single participant. Optional comment boxes throughout the survey provided rich context and allowed experts to elaborate on their reasoning, note uncertainties, or flag nuances. Respondents were allowed to skip for items outside their domain, ensuring high-quality responses for each topic.

Respondent Demographics

- Sample size: Survey responses varied by question and weapon type. For most core questions:
    - Nuclear: 20–42 respondents per question
    - Biological: 29–53 respondents per question
    - Chemical: 17–34 respondents per question

- Expertise: Respondents included a mix of technical experts, policy specialists, and practitioners from government, academia, and the private sector, with backgrounds in nuclear engineering, biosciences, chemistry, non-proliferation, defence, and artificial intelligence.

- Geography: Majority from the United States and allied countries, with broad representation across academia, industry, and government-affiliated research institutions.

Bottom Line Up Front

Respondents broadly agreed that AI can accelerate, optimise, or support many phases of WMD development, with the greatest near-term impact in planning, operational support, and information operations. However, substantial practical, technical, and knowledge barriers remain – AI is an accelerator, not an equaliser. The greatest risks are likely to accrue to proliferators with existing expertise and resources, with disinformation and operational planning support standing out as the most immediate and cross-cutting threats.

Nuclear Weapons

Key Insights and AI Touchpoints

- Design and Optimisation:
    - Most respondents agreed that predictive AI tools can aid in nuclear weapons design, particularly for simulation, numerical codes, and 3D modelling, provided quality data is available.
    - Generative AI, especially for 3D modelling, was seen as a modest aid (e.g., for components or enrichment tech), but less transformative for specialised or classified designs.

- Production and Engineering:
    - Additive manufacturing (3D printing) and robotics are expected to offer incremental advantages for non-nuclear components, select weapon parts, and manufacturing efficiency – but scepticism remained about their utility for highly specialised or sensitive items (e.g., plutonium pits).
    - AI for enrichment and reprocessing: Respondents saw AI as potentially helpful for optimising and troubleshooting enrichment processes, but data and complexity are limiting factors.

Noted Limitations

- Most respondents stressed that AI's impact is bounded by data availability, persistent technical/tacit knowledge barriers, and the high complexity and security around nuclear weapons production.

Biological Weapons

Key Insights and AI Touchpoints

- Pathogen Design and Enhancement:
    - A majority agreed that predictive AI can leverage sequence data to assist in enhancing pathogen characteristics, strain selection, and even propose novel pathogen "prototypes" – if high-quality, annotated data is available.
    - Limits: Many respondents highlighted that biology remains "messy"; even with AI, experimental validation, tacit knowledge, and host/pathogen interactions are major hurdles.

- Ethnic Targeting and Biomarkers:
  - Most agreed AI can identify population genetic markers but noted that “unique” ethnic biomarkers are rare and targeting is scientifically problematic; broad discrimination is possible but not precise targeting.

- Production and Automation:
  - AI and automation were seen as major “uplifts” for scaling up production, optimising bioreactors, and automating routine lab functions – especially for well-resourced proliferators.
  - Benchtop DNA synthesis: Viewed as an emerging risk, though many noted that DNA synthesis is only one step in a complex chain.

- Acquisition, Procurement, and Social Engineering:
  - Generative AI tools can help identify researchers, companies, and facilities, and assist in impersonation/social engineering (e.g., for ordering genes or accessing cloud labs). However, impact is limited by existing screening protocols and the need for hands-on skill.

- Operational Planning and Delivery:
  - Experts agreed that AI tools can aid in operational planning (weather prediction, crowd density, drone route optimisation), but effectiveness depends on data and real-time adaptability.
  - Disinformation: Overwhelming consensus that generative AI can produce disinformation to complicate public health responses, amplify psychological impact, and increase casualties.

- Equipment and Additive Manufacturing:
  - AI-assisted 3D modelling and printing can help design and produce bespoke lab and delivery equipment, but actual utility depends on access to suitable hardware and expertise.

Noted Limitations

- Experts consistently cautioned that AI cannot replace tacit biological expertise, troubleshooting skill, or the complexity of living systems. In other words, “design” is not “capability”.

Chemical Weapons

Key Insights and AI Touchpoints

- Molecule Design, Synthesis, and Toxicity Prediction:
  - Experts showed a strong consensus that predictive AI can assist in toxicity prediction, novel agent design, and property screening, given sufficient data. Many noted that such tools already exist and are improving rapidly.
  - Generative AI can help design new molecules, but actual weaponisation still requires expert chemists and suitable facilities.

- Production and Automation:
  - AI tools were seen as valuable for process optimisation, reactor control, and automating chemical production (especially when combined with robotics), but “most stages” being automated is viewed as an overstatement—complexity and tacit knowledge still matter.

- Acquisition and Illicit Procurement:
  - AI tools can assist in identifying feedstock, developing procurement strategies, and posing as legitimate companies; generative AI is especially useful for information retrieval and deception, but hallucinations and data quality limit effectiveness.

- Delivery and Dissemination:
  - AI and 3D modelling tools can help design sprayer nozzles/aerosolisation devices; additive manufacturing can enable custom equipment, but practical impact is greater for chemical than biological agents.
  - Operational planning: AI can aid in weather analysis, targeting, and drone route planning.

- Disinformation:
  - Experts showed a strong consensus that generative AI can amplify chaos, impede response, and increase impact through tailored disinformation campaigns, paralleling findings for biological weapons.

Noted Limitations

- While AI can accelerate many steps, production and weaponisation of chemical agents still require advanced skill, infrastructure, and safety protocols.

Cross-cutting Themes

- AI is a force multiplier for well-resourced proliferators and those with baseline expertise; “uplift” for novices is limited by persistent technical and practical barriers.
- Data availability and quality are critical – AI’s impact is capped by gaps in training data, screening, and real-world feedback.
- Disinformation and operational planning support are the most mature and near-term risks across all WMD types.
- Automation and additive manufacturing are seen as converging risks for WMD.
- Tacit knowledge, hands-on skill, and organisational capacity remain core barriers to WMD proliferation, even with advanced AI.

Trends and Drivers

To better understand how AI might affect WMD proliferation over time, the project team brought together experts to identify important patterns and underlying forces that could influence this relationship.

The analysis focused on two key elements:

- Trends: Current patterns that can be observed and measured today
- Drivers: Deeper forces driving these changes

This approach helped distinguish between what is happening on the surface and the fundamental shifts occurring beneath.

The study’s conclusions suggest that while AI has not dramatically changed WMD development yet, the field stands at an important moment. The next five to ten years could see AI evolve from simply making existing processes more efficient to actually enabling new capabilities as technologies combine in powerful ways, data resources improve, and protective measures struggle to keep pace.

By watching these patterns carefully, policymakers and diplomats can spot early warning signs before AI crosses important thresholds in weapons development.

The tables that follow outline the specific trends and drivers identified for nuclear, biological, and chemical weapons, highlighting the unique concerns for each weapon type.

Nuclear Weapons

While AI may not transform nuclear weapons development as dramatically as it might for biological or chemical weapons, it introduces important changes that deserve close attention from policymakers and diplomats.

The table below examines key trends emerging at the intersection of AI and nuclear weapons, along with their implications. These include varying effects across different development stages (with potentially greater impact for countries newly pursuing nuclear capabilities than for established nuclear powers), enhancements to actor capabilities, automation possibilities, improvements in targeting, and risks related to information operations.

Particularly noteworthy is how AI might speed up nuclear development timelines when combined with other emerging technologies like additive manufacturing (3D printing). These implications are further complicated by the unique information landscape surrounding nuclear weapons—where basic principles are widely available in public sources, but critical design details remain highly classified, creating an uneven landscape for AI applications.

Understanding these trends and their implications is crucial for developing appropriate safeguards and international frameworks that address proliferation concerns.

Trend	Implications
Development stage	AI could assist in sorting through large, complex datasets for relevant nuclear design knowledge or to aid in reverse engineering for less experienced states. Limited use for advanced nuclear states that have already optimised designs and for states without access to classified data.
Actor capability enhancement	AI could help developing states enhance production, identify vulnerabilities in the supply chain and target them for acquisition, and refine nuclear designs if, for example, proprietary datasets are acquired. For non-state actors, AI could improve reconnaissance, especially relevant to identifying and targeting key personnel, supply chain vulnerabilities, or uncovering unidentified physical security vulnerabilities.
Automation	AI could automate repetitive and error-prone manufacturing processes, especially for developing states, and in convergence with other technologies, like additive manufacturing and robotics. AI could be used for predictive maintenance, assisting in diagnostics of reactor systems and aiding in the life-extension of systems.
Targeting optimisation	AI could contribute to improved navigation, targeting and survivability of delivery systems.

Trend	Implications
Information operations	AI could be used to improve situational awareness and aid early warning, especially given capabilities to analyse multiple data streams in real time. AI could also be used to create misinformation and disinformation about capabilities or interpretation of intent, which could lead to miscalculation, misinterpretation, and inadvertent escalation.
Convergence with other technologies	AI integration with additive manufacturing and robotics can enhance production of nuclear-related components and improve manufacturing processes. Convergence with other technologies could assist in overcoming tacit knowledge barriers, shortening timelines, reducing production footprint, avoiding detection, and producing prototypes.
Data landscape	The science behind nuclear weapons is rather simple and well documented in open-source scientific literature which can easily be mined using AI. Classified information specific to weapons design is more difficult to obtain, however, if AI becomes more integrated into nuclear weapons systems and much of that information is digitised, there is an increased vulnerability of using AI to target systems and obtain classified data via AI-enhanced cyber-attacks.

The table below identifies the key drivers catalysing the integration of AI into nuclear weapons programmes. While AI's effects on nuclear weapons appear more gradual than the potentially transformative impact it may have on biological or chemical weapons in the coming decade, these underlying forces are steadily reshaping nuclear weapons development, deployment, and strategic thinking.

The drivers fall into four categories – technological, scientific, economic, and institutional – reflecting the complex mix of factors influencing the relationship between AI and nuclear weapons. Understanding these fundamental drivers is essential for developing effective international frameworks and risk reduction measures that address proliferation concerns.

Category	Drivers for the AI-Nuclear Nexus
Technological	Ubiquitous integration of AI tools (predictive and generative) Additive manufacturing/3D printing capabilities Advanced robotics and automation systems Digitisation of nuclear design and testing data High-performance computing advancements Digital twin development for complex systems Enhanced modelling and simulation capabilities AI-facilitated knowledge transfer and interpretation Novel approaches to complex problem solving
Scientific	Materials design and optimisation techniques Advanced stockpile stewardship methods Hydrodynamic testing improvements
Economic	Cost reduction and efficiency improvements Shift from traditional to additive/advanced manufacturing Reduced manufacturing footprint requirements Lower infrastructure and personnel costs Elimination of detectable waste streams
Institutional	Cybersecurity vulnerabilities and mitigations Ongoing nuclear modernisation programmes Bureaucratic and political incentives Arms race dynamics and perception gaps Non-proliferation monitoring improvements Integration of AI into existing nuclear processes

Biological and Chemical Weapons

Biological and chemical weapons development shows similar patterns as AI reshapes proliferation risks, which is why they are presented together here. Those seeking these capabilities are moving away from large-scale military applications toward more targeted, smaller attacks, with AI providing the greatest advantage to mid-level actors.

While AI significantly helps in early research and design phases, important barriers remain in practical laboratory techniques and safety protocols that limit its impact on weaponisation and large-scale production. Information operations powered by AI – such as spreading false information about disease outbreaks or fabricated chemical attacks – are becoming as strategically important as actual weapons.

When AI combines with automation, smaller production facilities, and autonomous delivery systems, it creates new risk profiles. However, at present, AI functions more as a tool that improves and accelerates existing processes rather than one that fundamentally transforms capabilities. The table below details these trends across biological and chemical weapons domains.

Trend	Biological Weapons	Chemical Weapons
Shift in attack objectives	From mass casualty battlefield use to targeted assassination, economic disruption, and regime security applications	From large-scale military deployment to small-scale sabotage, targeted killings, and terror/psychological operations
Actor capability enhancement	AI provides greatest uplift to mid-level proliferators (scientists lacking weaponisation knowledge) and state programmes for novel agent design	AI most benefits proliferators with some chemistry background for synthesis planning and those seeking to evade export controls
Development and acquisition stage	AI most useful in early R&D (protein design, target identification) and ideation, but limited impact on production scale-up and weaponisation	AI most valuable for acquisition (precursor identification, route planning) and synthesis optimisation, less for weaponisation
Tacit knowledge evolution	Physical lab skills remain critical despite AI assistance; automation creates new tacit knowledge requirements rather than eliminating them	Hands-on synthesis and safety skills persist as barriers; AI cannot replace practical chemistry experience
Data landscape	Proprietary biological data increasingly siloed in private companies while public databases face defunding	Chemical data more accessible and representation diversity (SMILES, structures, names) enables evasion of AI screening

Trend	Biological Weapons	Chemical Weapons
Design capabilities	In silico design advancing rapidly (proteins, pathogens) but wet lab validation and biological unpredictability remain major hurdles	Retro-synthesis and toxicity prediction tools proliferating but scale-up and safe handling still require human expertise
Automation	Cloud labs and automated equipment reducing some barriers but BSL requirements and biological variability limit impact	Automated synthesis lagging behind bio; complex reaction conditions and safety requirements maintain human involvement
Information operations	AI-generated disinformation about disease outbreaks or attribution becoming major concern alongside kinetic threats	AI-enabled psychological operations and false flag scenarios potentially more impactful than actual chemical attacks
Novel agent development	AI enabling design of modified pathogens and novel toxins, but functional validation remains challenging	AI optimising known agent classes rather than creating fundamentally new chemicals; toxins (bio-chem overlap) most promising
Convergence with other technologies	Integration with automated labs, gene synthesis, and drone delivery creating new risk profiles	Combination with miniaturised synthesis, 3D printing, and autonomous systems enabling new attack vectors

Technological, scientific, economic, and institutional drivers are changing the landscape of biological and chemical weapons development. Unlike the nuclear weapons analysis, which focused on AI integration into weapons systems, this section takes a different approach, which reflects the highly digital and dual-use nature of modern biological and chemical sciences. The project team examined how rapid digitisation in life sciences and chemical industries creates a fundamentally different environment for AI applications. This dual-use reality means that legitimate commercial and scientific advances simultaneously reshape the potential weapons development landscape.

The table below presents these drivers, showing where they overlap across both weapons categories and where they appear uniquely within each domain. Understanding these underlying forces provides essential context not just for current patterns but for the deeper factors that will determine future developments in the AI-WMD nexus.

Category	Overlapping	Biological Weapons	Chemical Weapons
Technological	Large language models (LLMs) for knowledge synthesis Cloud computing infrastructure Laboratory automation equipment High-throughput screening systems Drone delivery platforms Remote access/ cloud laboratories Data mining and analytics tools Miniaturisation technologies	Gene synthesis and sequencing CRISPR and gene editing tools Protein folding AI (AlphaFold) Bioreactor automation Cell-free expression systems Metagenomic analysis tools Environmental DNA sampling Automated cell culture systems Bioprinting technologies Desktop synthesisers	Retro-synthesis planning software Chemical structure databases Desktop synthesisers Flow chemistry systems Multiple representation formats (SMILES, etc.) Reaction optimisation algorithms Automated liquid handlers Industrial process control systems
Scientific	Toxicology and pharmacology advances Structure-function relationships Systems modelling approaches Convergence of bio-chem disciplines Data science methodologies	Gene synthesis and sequencing CRISPR and gene editing tools Protein folding AI (AlphaFold) Bioreactor automation Cell-free expression systems Metagenomic analysis tools	Retro-synthesis planning software Chemical structure databases Desktop synthesisers Flow chemistry systems Multiple representation formats (SMILES, etc.) Reaction optimisation algorithms Automated liquid handlers

Category	Overlapping	Biological Weapons	Chemical Weapons
Scientific	Computational modelling Dual-use research expansion	Environmental DNA sampling Automated cell culture systems Bioprinting technologies Desktop synthesizers	Industrial process control systems
Economic	Low cost of computational resources Data as competitive asset Dual-use market growth Cost of disruption vs. mass casualty Private sector R&D investments Reduced team size requirements IP and data monetisation Dark web/ cryptocurrency enabling	Bioeconomy expansion Pharmaceutical industry spillover Venture capital in biotech DNA synthesis cost reduction Proprietary sequence databases Agricultural biotechnology Personalised medicine market	Global chemical trade volumes Lower infrastructure costs Precursor market fragmentation Industrial chemical availability Patent circumvention value Process optimisation savings Regulatory compliance costs
Institutional	Norm erosion State programme revival	Bioeconomy expansion Pharmaceutical industry spillover Venture capital in biotech	Global chemical trade volumes Lower infrastructure costs

Category	Overlapping	Biological Weapons	Chemical Weapons
Institutional	Weakening international controls Intelligence vs. military priorities Gray zone warfare adoption Regime security focus Academic-industry knowledge transfer Open science movements Fragmented regulatory oversight	DNA synthesis cost reduction Proprietary sequence databases Agricultural biotechnology Personalised medicine market	Precursor market fragmentation Industrial chemical availability Patent circumvention value Process optimisation savings Regulatory compliance costs

Use Cases: Predictive and Generative AI

The tables below compare how predictive AI models and generative AI models may affect nuclear, biological, and chemical weapons development with different types of use-cases. Predictive AI models identify patterns in existing data to forecast outcomes or classify information, typically using supervised learning with labelled datasets. Generative AI models create new content, designs, or solutions by learning data patterns and producing outputs not explicitly included in their training data.

This distinction may blur as hybrid approaches emerge. Some models now perform both predictive functions (classifying inputs, making forecasts) and generative tasks (creating text, images, code, or designs) depending on how users prompt or fine-tune them.

All potential use-cases described below depend critically on quality, relevant, and representative data. Without good data, even the most sophisticated AI systems produce flawed or misleading results. Many WMD-relevant applications face significant data limitations due to classification restrictions, few historical examples, or lack of verified information, which can render advanced AI systems ineffective or dangerously misleading in these domains.

Nuclear Weapons

This table identifies specific ways predictive and generative AI may affect nuclear weapons development. While AI will transform nuclear weapons less dramatically than biological or chemical weapons, it enhances capabilities at every stage from research through delivery.

Predictive AI optimises processes, detects anomalies, and supports simulations that accelerate development and improve reliability. Generative AI, though more limited in nuclear contexts, creates new risks particularly in acquisition (through sophisticated social engineering) and information operations (through convincing deepfakes and disinformation). AI's convergence with additive manufacturing and robotics particularly concerns experts, as these combinations could significantly reduce production barriers and enhance weaponisation capabilities.

Unlike biological weapons where AI might create novel agents, AI's nuclear impact primarily accelerates processes, improves efficiency, and enhances delivery precision. Understanding potential applications helps develop targeted governance measures that address proliferation risks.

Stage	Predictive AI	Generative AI
R&D	Optimise design trade-offs through simulations. Identify and optimise viable nuclear weapons designs from historical data. Improve 3D modelling of weapons designs.	Search existing scientific literature on nuclear weapons designs and data (assuming it exists) – assist in evaluation and potential improvements.
Acquisition	Identify anomalies in behavioural patterns of staff (e.g., insider threat detection) and material sourcing (e.g., supply chain vulnerabilities) to gain access to nuclear material.	Search existing scientific literature on nuclear weapons designs and data (assuming it exists) – assist in evaluation and potential improvements.
Production	Optimise manufacturing processes, especially in combination with convergent technologies (e.g., additive manufacturing). Predictive maintenance of equipment or component failures.	Generate new designs to be used in additive manufacturing.
Weaponisation	Optimise explosive and propellant formulas, predict batch quality, and enhance reliability. Use predictive maintenance for stability and quality of nuclear and non-nuclear components.	Potentially generate novel chemical combinations for propellants or conventional explosives.

Stage	Predictive AI	Generative AI
Weaponisation	Automate certain tasks in combination with convergent technology (e.g., robotics) to enhance precision and safety. Enhance simulations that replace physical testing for greater understanding of potential real-world use case. In combination with convergent technologies (e.g., additive manufacturing), optimise production of non-nuclear components and improve quality and time required for prototype development.	
Delivery	Use predictive maintenance for delivery systems. Refine targeting systems and optimise accuracy in delivery systems. Enhance situational awareness through rapidly synthesising real-time data streams. Assess effectiveness of countermeasures and model adversary responses. Improve data integration to aid in early warning and reduce false alarms.	Generate scenarios for future conflict. Generate deepfakes or disinformation to mislead in a crisis.

Biological Weapons

This table maps how predictive and generative AI may affect each stage of biological weapons development. During research and design, AI technologies will transform the landscape most dramatically: predictive models like AlphaFold revolutionise protein structure prediction while generative systems allow for the design of novel toxins and pathogens. These capabilities extend into acquisition, where AI may help identify source materials and generates sophisticated strategies to evade biosecurity controls.

As development advances to production and weaponisation, AI applications become more specialised, potentially solving technical challenges in scaling biological agents and optimising their stability and delivery characteristics. This mapping reveals how AI may compress development timelines, lower technical barriers, and enhance capabilities across the entire biological weapons pathway, challenging non-proliferation efforts and international security frameworks.

Stage	Predictive AI	Generative AI
R&D/Design	Predicting protein structure (AlphaFold) Identifying optimal genetic modifications Testing hypotheses without lab work Anticipating pathogen behaviour	Novel toxin design Creating pathogen variants Designing immune-evading proteins Generating research protocols
Acquisition	Identifying source locations Predicting optimal collection methods Assessing biosecurity vulnerabilities Analysing synthesis providers' screening methods	Generating strategies for accessing pathogens Creating deceptive cover stories for purchases Designing sequence modifications to evade screening Generating false credentials for facility access
Production	Optimising growth conditions Predicting scale-up challenges Monitoring bioreactor conditions Quality control automation	Creating optimised production protocols Designing equipment APIs Generating troubleshooting guides Creating custom manufacturing systems
Weaponisation	Predicting environmental stability Assessing aerosol behaviour	Designing novel delivery systems Creating stabilisation formulations Generating aerosolisation protocols Developing equipment modifications
Delivery	Weather pattern analysis for dispersal Estimating effectiveness Identifying vulnerable targets Modelling impact scenarios	Generating operation timelines Planning multi-pronged attack strategies

Chemical Weapons

This table illustrates how predictive and generative AI may affect each stage of chemical weapons development. During research and design, AI may accelerate development through protein structure-activity relationship analysis and novel compound design, potentially creating agents engineered to evade detection or defeat countermeasures.

As development moves through acquisition and production, AI may shift toward assistance in circumventing controls and optimising processes. Generative AI particularly concerns experts by creating deceptive procurement strategies and alternative synthesis pathways that bypass existing regulations. In weaponisation and delivery stages, AI may enhance both technical aspects of chemical weapons deployment (through stability modelling and optimised formulations) and operational security (through disinformation and attribution obfuscation).

This mapping demonstrates how AI may compress development timelines, lower technical barriers, and enhance capabilities across the entire chemical weapons pathway, challenging non-proliferation efforts and international security frameworks.

Stage	Predictive AI	Generative AI
R&D/Design	Structure-activity relationship analysis Toxicity prediction	Novel toxic compound design Designing compounds to evade detection Generating synthesis routes for novel agents Creating agent formulations with specific properties Designing compounds to defeat countermeasures
Acquisition	Predicting alternative precursor sources Identifying dual-use chemicals Analysing export control vulnerabilities Detecting procurement patterns Predicting supply chain weaknesses	Creating deceptive procurement strategies Generating alternative synthesis routes Designing procurement cover stories Creating false credentials for suppliers Generating smuggling route options

Stage	Predictive AI	Generative AI
Production	Optimising reaction conditions Predicting scaling challenges Process safety analysis Yield optimisation Quality control prediction	Designing novel production processes Creating automated synthesis instructions Generating equipment specifications Designing miniaturised production systems Creating safety protocols for hazardous synthesis
Acquisition	Predicting alternative precursor sources Identifying dual-use chemicals Analysing export control vulnerabilities Detecting procurement patterns Predicting supply chain weaknesses	Creating deceptive procurement strategies Generating alternative synthesis routes Designing procurement cover stories Creating false credentials for suppliers Generating smuggling route options
Production	Optimising reaction conditions Predicting scaling challenges Process safety analysis Yield optimisation Quality control prediction	Designing novel production processes Creating automated synthesis instructions Generating equipment specifications Designing miniaturised production systems Creating safety protocols for hazardous synthesis
Weaponisation	Environmental stability modelling Shelf-life estimation Predicting interactions with protective gear	Designing novel delivery mechanisms Creating agent formulations for dispersal
Delivery	Modelling atmospheric dispersion Predicting optimal delivery conditions Weather pattern analysis Target vulnerability assessment	Generating tactical deployment plans Developing disinformation campaigns Creating attribution obfuscation strategies

VCDNP

Vienna Center for Disarmament and Non-Proliferation

The VCDNP is an international non-governmental organisation that promotes peace and security by conducting research, facilitating dialogue, and building capacity on nuclear non-proliferation and disarmament.

vcdnp.org
@VCDNP
info@vcdnp.org
VCDNP