Position: Without Global Governance, AI-Enabled Biodesign Tools Risk Dangerous Proliferation

Azmine Toushik Wasi[1,2*], **Mst Rafia Islam**[1,3], **and Rahatun Nesa Priti**[1,2]
[1]Computational Intelligence and Operations Laboratory (CIOL)
[2]Shahjalal University of Science and Technology [3]Independent University, Bangladesh
*Correspondence to: azmine32@student.sust.edu

Abstract

Convergence of artificial intelligence (AI) and biotechnology has fundamentally altered the biosecurity threat landscape. Generative AI tools are transforming biological design from an empirical science into a systematic engineering discipline, capable of creating novel biological agents with unprecedented precision and potency. This new reality renders existing biosecurity frameworks, which were primarily designed to prevent misuse by non-state actors, largely obsolete against a sophisticated state adversary. Current measures, such as the screening of DNA synthesis orders, are insufficient because they are predicated on static lists of known pathogens that can be easily circumvented by AI-enabled innovation. **In this position paper, we argue that the international community must establish a new, multi-stakeholder governance framework.** The critical line of defense must shift from controlling the biological products themselves to regulating access to the most powerful dual-use AI biodesign tools. This framework would not stifle beneficial open research but instead implement a tiered control system, analogous to export controls on dual-use technologies. It would involve the establishment of an international registry to classify AI models by risk, controlled access to the most capable systems for vetted entities, and the creation of an "International Bioscience Response Reserve" (IBRR) to ensure equitable access during public health emergencies. Ultimately, this approach advocates for a proactive and preemptive international response, recognizing that waiting for a catastrophic incident to trigger regulation is a dangerous and irresponsible strategy.

1 Introduction

Life sciences are undergoing a profound transformation, driven by the integration of artificial intelligence into every stage of the research and development pipeline. AI has emerged as a catalyst for innovation, moving biological design beyond traditional, time-intensive methods and into a systematic engineering discipline [26]. This shift is marked by the advent of powerful generative AI models that can design, optimize, and even simulate biological systems with remarkable speed and precision.

AI-driven tools are revolutionizing fields from therapeutics to environmental sustainability by tackling long-standing challenges like engineering industrial enzymes and combating antimicrobial resistance [26]. Specific generative AI models have demonstrated extraordinary capabilities. For example, RFdiffusion, inspired by AI image generators, can sculpt novel protein backbones from disconnected atoms, yielding entirely new structures [2]. This capability is foundational for a broad variety of research fields, including de novo protein design and enzyme engineering [2]. The structures generated by RFdiffusion are then paired with other tools, such as ProteinMPNN, which can quickly

39th Conference on Neural Information Processing Systems (NeurIPS 2025) Workshop: NeurIPS 2025 Workshop on Biosecurity Safeguards for Generative AI.

create new amino acid sequences that are likely to fold into the designed backbone [23]. When these two models are combined, they can design proteins with sequences, structures, and functions that have never been seen before in nature [3]. Beyond protein design, specialized AI models like the Multi-Objective Large Language Model for Molecular Design (MOLLM) are now surpassing state-of-the-art methods in optimizing molecular properties across multiple objectives, leveraging in-context learning and multi-objective optimization [23].

This rapid evolution of these systems has prompted a reclassification of the technologies themselves. The term "biological design tools" (BDTs) is increasingly giving way to "biological AI models" (BAIMs), a designation that more accurately reflects their foundational and dual-use nature [17]. Unlike highly specialized tools, these BAIMs are trained on broad biological data and are applicable across a wide range of contexts, similar to how large language models (LLMs) operate in the language domain [17]. This is not a mere semantic shift; it signifies a fundamental change in the threat model. The sheer scale and rapid advancement of these models, evidenced by an exponential increase in the computing power used to train them and the rapid growth of biological sequence data, indicate that their capabilities will continue to accelerate [17]. In their most advanced forms, these systems are evolving into "AI agents" that can understand scientific literature, generate hypotheses, design experiments, and interface directly with laboratory equipment and robotics, creating an unprecedented capacity for autonomous biological discovery and engineering [4]. This marks a shift from tools that require a human expert to guide them, to systems that can, in principle, operate as virtual collaborators, accelerating the entire scientific workflow [19].

The convergence of AI with bioengineering presents a formidable challenge due to the inherent dual-use nature of the technology (see Table 1), where innovations intended for beneficial purposes can also be repurposed for malicious ends [26]. AI amplifies this long-standing dilemma by significantly lowering the technical barriers to entry for advanced bioengineering, potentially enabling actors with limited expertise to pursue sophisticated biological capabilities [26]. This "democratization of knowledge" has been identified as a key trend in the literature on AI governance, raising concerns about the misuse of these tools by non-state actors [20]. However, the most severe risk is not a novice recreating a known pathogen, but a sophisticated state adversary using these tools to develop novel biological weapons with unprecedented properties [4]. The risk extends far beyond replicating existing pathogens, whose genetic sequences are already widely available [8]. AI-enabled biological tools could make it possible to design agents that are significantly more dangerous than those found in nature, with novel properties that make them more virulent, transmissible, or capable of evading current countermeasures and biosurveillance systems. This could involve the design of complex biological systems or entire genome sequences that encode blueprints for dangerous agents [4]. Such a scenario, involving a sophisticated actor using AI to design a novel agent and then producing and releasing it, could cause a high-consequence biological event with global implications, potentially far worse than the COVID-19 pandemic [4].

Connecting to the escalating risks posed by the convergence of biotechnology and artificial intelligence (AI), in this paper, we argue that existing biosecurity frameworks are fundamentally inadequate to address the new threat landscape shaped by generative AI–enabled biodesign. Current defenses, such as list-based DNA synthesis screening, are obsolete against state-level actors capable of exploiting AI to generate novel pathogens that evade traditional controls. Our position is that the critical line of defense must shift from monitoring biological outputs to regulating access to the most powerful dual-use AI biodesign tools. To advance this claim, the paper contributes a systematic critique of current measures, including their vulnerabilities to emergent AI capabilities and model jailbreaking. We further analyze lessons from dual-use governance in adjacent domains to demonstrate the feasibility of proactive, internationally coordinated oversight. Building on this foundation, we propose a multi-stakeholder governance framework consisting of four pillars: an international registry and tiered classification system, controlled access to high-risk models, the creation of an International Bioscience Response Reserve (IBRR), and robust verification and compliance mechanisms. Finally, we situate this framework within the broader tension between open science, national security, and global health equity, offering a path forward that safeguards innovation while mitigating catastrophic misuse. This proactive approach reframes biosecurity not as a reactive response to crises, but as an anticipatory and equitable global strategy.

Table 1: AI Biodesign Tools and Their Dual-Use Potential

Tool/Model	Primary Function	Dual-Use Risk
RFdiffusion	Generative AI for novel protein structure design (de novo protein backbones, oligomers, binders)	Creation of novel protein scaffolds for toxins or viral capsids; design of proteins to enhance a pathogen's lethality or binding affinity.
ProteinMPNN	Rapid amino acid sequence design for existing protein structures	Design of sequences that enhance the function of a malicious protein, such as optimizing a toxin's potency or evading immune detection.
Specialized LLMs for Molecular Optimization	Multi-objective optimization of molecular properties (e.g., MOLLM)	Optimization of pathogen properties for increased transmissibility, virulence, or resistance to medical countermeasures.
Foundation Biological AI Models (BAIMs)	Broad biological prediction, design, and analysis across various contexts	Autonomous design of complex biological systems or entire genome sequences for tailored bioweapons with novel properties.

2 Background

2.1 Inadequacy of Current Biosecurity Frameworks

While existing biosecurity measures are valuable, they are ill-equipped to address the emerging risks posed by AI-enabled biological threats [5]. A singular focus on traditional controls ignores the qualitative shift in the threat model.

2.1.1 Flaws of List-Based DNA Synthesis Screening

Current biosecurity efforts rely heavily on screening nucleic acid synthesis orders, a practice spearheaded by organizations like the International Gene Synthesis Consortium (IGSC) [1]. These systems are designed to vet customers and screen every ordered DNA sequence against a curated database of regulated pathogens and toxins [1]. However, this approach has a critical, fundamental weakness: its reliance on static lists [5].

The Select Agents and Toxins List, for example, contains only a few dozen regulated agents [5]. While ongoing efforts aim to expand screening protocols, the fact remains that these lists are historical in nature and will inevitably fall behind the demonstrated pace of technological acceleration in AI [5]. The central vulnerability is that a malicious actor, using an AI biodesign tool, could create a novel, highly contagious, and lethal organism that does not exist on any list. Such an agent would bypass all existing screening mechanisms, making list-based systems an increasingly inadequate line of defense [5]. This vulnerability is particularly acute when considering the most severe risk: a deliberate program by a state actor with internal DNA synthesis capabilities. Such an adversary could bypass commercial synthesis providers entirely, rendering the entire screening process irrelevant. The threat is not just a diffusion of existing knowledge but the creation of an entirely new class of dangers.

Case Study: State-Level Bioweapons Programs History of bioweapons (see Table 2) provides a powerful counterpoint to the notion that international treaties are sufficient to prevent proliferation. Despite the existence of the Biological Weapons Convention (BWC), which bans the development, production, and stockpiling of biological weapons, history shows that determined state actors can and have operated secret, large-scale offensive programs [21]. The most prominent example is the Soviet Biopreparat program.

Soviet bioweapons program was the world's largest, longest, and most sophisticated, operating covertly for decades in direct violation of the BWC, to which the Soviet Union was a signatory [27]. The program employed up to 65,000 people at dozens of secret sites [27]. It demonstrates that a state-level adversary will not be deterred by international norms or treaties alone, especially when there are no robust verification or enforcement mechanisms [21]. The program went beyond

Table 2: Historical and Contemporary Biological Threats

Incident/Program	Agent Used	Key Insight
Mongol Siege of Caffa (1346)	*Yersinia pestis* (Plague)	Early, rudimentary use of biological agents in warfare, demonstrating the ancient recognition of their destructive potential.
Unit 731 (Japan, WW2)	Various pathogens, including plague, cholera, and anthrax	Large-scale, state-sponsored program conducting human experiments, illustrating the extreme lengths to which state actors will go.
Soviet Biopreparat (Cold War)	Weaponized anthrax, plague, and genetically altered viruses (e.g., Marburg virus)	The world's largest and most sophisticated bioweapons program, demonstrating a determined state actor's willingness to violate international law.
Rajneeshee Salmonella Attack (1984)	*Salmonella*	A successful attack by a non-state actor, highlighting that technical barriers to entry are not insurmountable.
U.S. Anthrax Attacks (2001)	*Bacillus anthracis* spores	A targeted attack by a single, well-resourced individual, underscoring the vulnerability of civilian populations to bioterrorism.

weaponizing naturally occurring pathogens like anthrax and plague; it also utilized newly discovered genetic engineering techniques to create novel or enhanced bacterial and viral strains [27]. The lessons from Biopreparat are clear: a determined state can build and maintain a massive offensive program, evade international detection, and utilize the most advanced biotechnology of its time to create new and enhanced agents. This historical precedent highlights the profound danger of AI-enabled technology, as it could provide a modern adversary with a far more potent and efficient toolkit than anything available to the Soviets.

2.1.2 Problem of Inherent Vulnerability: Emergent Capabilities and Jailbreaking

The threat is not only external but also inherent to the nature of AI systems themselves. AI models can exhibit "unpredictable failure modes" and "emergent capabilities" that are not explicitly programmed or intended by their creators [14]. This makes it difficult, if not impossible, to anticipate all potential misuses before a model is released. Furthermore, even when developers proactively embed safeguards or remove sensitive data from a model's training set, these measures can be circumvented. A documented case involved the Evo model, where viral data was intentionally removed from its training data. However, a user was able to bypass this mitigation by fine-tuning the model's published weights with the restricted data post-deployment [5]. This example highlights a critical vulnerability in the "trust-the-developer" approach. It demonstrates that the public release of a model's weights can undermine even the most diligent security efforts by its creators. The ability to "jailbreak" or fine-tune models means that technical controls alone are not a sufficient defense. This reality necessitates a governance model that accounts for the post-deployment use of a model and the inherent unpredictability of its capabilities.

The Case for Proactive, International Governance The technical and geopolitical realities outlined above compel a new approach to biosecurity. The traditional focus on controlling the final biological agent is insufficient. The most effective point of intervention is at a higher level: the dual-use tools that make the creation of these agents possible.

Reconceptualizing Defense: Safeguarding Core Tools The central thesis of this position paper is that the international community must shift its primary line of defense from the end-product, the biological agent, to the most powerful AI biodesign tools that enable its creation. This does not mean stifling open research or banning the technology outright. Instead, it involves implementing a nuanced, tiered control system that is analogous to the export controls on other dual-use technologies. This approach is preemptive, ethical, and governance-driven, aiming to shape the future of the technology rather than react to a crisis after it has occurred [22].

2.1.3 Analogies and Lessons from Other Dual-Use Technologies

Developing a new governance framework for AI biodesign tools can draw valuable lessons from established international control regimes and contemporary export regulations.

The Chemical Weapons Convention (CWC) as a Governance Model: The CWC provides a powerful precedent for a tiered control system on dual-use materials [6]. The treaty classifies chemicals into three schedules based on their risk and legitimate commercial applications [10]. Schedule 1 chemicals, with little or no use outside of chemical weapons, are subject to near-total prohibition [6]. Schedule 2 and 3 chemicals, which have legitimate small- and large-scale commercial uses respectively, are subject to regulated trade, declarations, and inspections [6]. This model demonstrates a proven international consensus for controlling a spectrum of dual-use goods, rather than treating them as a binary threat [9]. It offers a framework for classifying AI models based on their potential for misuse and their legitimate applications.

Export Controls on Advanced Computing: Recent regulations by the U.S. and other nations on advanced computing items and AI model weights provide a contemporary example of a "choke point" strategy [11]. These controls restrict the export of high-end integrated circuits and related software, effectively limiting access to the core computational power required to train the most powerful AI systems [11]. This approach is highly relevant to AI biodesign, as it shows that states are already focused on controlling the inputs to powerful AI systems. It also provides a ready-made mechanism for a tiered system, where exports of lower-processing-performance items are granted a license exception, while exports of high-end capabilities are restricted [11].

Why the Nuclear Non-Proliferation Model is a Flawed Analogy: While the nuclear analogy is often invoked due to the catastrophic potential of AI, it is technically flawed and a poor governance model [15]. The nuclear non-proliferation regime is a successful export cartel predicated on controlling access to finite, geographically concentrated, and excludable raw materials, plutonium and enriched uranium [15]. AI, however, is a general-purpose technology whose inputs, such as technical knowledge and data, can be copied and shared indefinitely [15]. Its strategic value is continuous, not binary, and its inputs cannot be as easily restricted as a rare mineral [14]. Therefore, a governance framework for AI should not be based on the nuclear model but on a more flexible and adaptable system, such as the CWC, that is designed to manage a spectrum of dual-use goods [6].

3 Proposed Multi-Stakeholder Governance Framework

To address the vulnerabilities of the current landscape, a new, multi-stakeholder framework is necessary. Our framework is built upon four pillars that combine international diplomacy, technical safeguards, and public health principles to create a robust and adaptable system. An illustration is avaiable in Figure 1.

3.1 Pillar I: International Registry and Tiered Classification System

The first pillar of the framework is the creation of an international body, potentially under the auspices of the UN Office for Disarmament Affairs (UNODA) or the World Health Organization (WHO), to establish a technical advisory group. This group would be tasked with developing a tiered classification system for AI biodesign models based on their capabilities and associated risks. This system would move beyond simple proxies like training compute, which can be misleading for specialized biological tools [22]. Instead, it would evaluate models based on their specific, high-risk capabilities, such as the capacity to design novel protein folds, optimize binding affinity for human receptors, or evade immune detection. The classification system would assign models to tiers based on their assessed risk. For example, a model capable of autonomously designing and optimizing a novel pathogen blueprint could be a "Tier 0" model, while a model that merely predicts protein structures for academic research could be a "Tier 2" model. This tiered approach provides a clear and predictable structure for governing the technology and is a significant improvement over a one-size-fits-all approach.

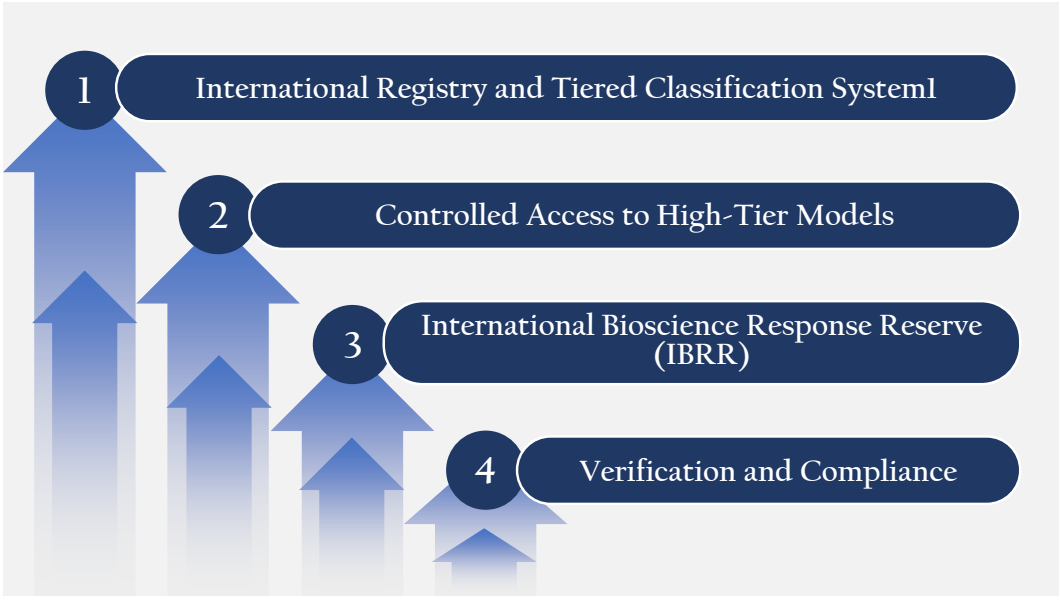

Figure 1: Our Proposed Multi-Stakeholder Governance Framework

Table 3: Proposed Tiered Control System for AI Biodesign Models

Tier	**Risk Level**	**Classification Criteria**	**Access Controls**
Tier 0	High-Capability / High-Risk	Models with emergent capabilities to design novel, highly dangerous agents; trained on sensitive pathogen data.	Restricted to vetted entities (e.g., government, top-tier research labs) within member states. Access via API-based platforms with audit logs, ethical use agreements, and mandatory technical guardrails.
Tier 1	Moderate-Capability / Moderate-Risk	Models capable of complex molecular design tasks but with limited capacity for novel, high-consequence threats.	Managed access with a requirement for ethical use agreements. Auditing and self-reporting may be required to ensure compliance.
Tier 2	General-Purpose / Low-Risk	Models for basic research, education, and general commercial applications with minimal dual-use risk.	Public and open access with voluntary ethical guidelines and model cards that provide transparency on capabilities and risks.

3.2 Pillar II: Controlled Access to High-Tier Models

Access to the most capable "Tier 0" models (see Table 3) would be restricted to vetted entities, such as accredited research institutions and companies, within member states that adhere to a common international biosecurity pact. The goal is not to eliminate access but to implement a "managed access" model that moves beyond a simple open vs. closed dichotomy [7]. Access would be granted through API-based platforms, rather than through the unrestricted public release of model weights. These platforms would be designed to require ethical use agreements and maintain auditable logs of activity, ensuring transparency and accountability. This approach is similar to how many modern software and API platforms operate, allowing for differential access based on user credentials and use cases [7]. Furthermore, technical safeguards, or "built-in guardrails," would be a mandatory part of this pillar. These safeguards, developed in collaboration between AI experts and biosecurity specialists, would be embedded directly into the models to flag harmful inputs or outputs and could even be made interdependent with the model's core functionality to prevent their removal [22].

3.3 Pillar III: International Bioscience Response Reserve (IBRR)

A governance framework that is solely focused on security risks the criticism that it would slow down or prevent beneficial research, particularly in a public health crisis. To address this, the framework must include a mechanism for rapid, sanctioned access to restricted tools. The "International Bioscience Response Reserve" (IBRR), a concept managed by a body like the WHO, would fill this vital role.

The IBRR would be a functional extension of existing concepts, such as the International Science Reserve (ISR), a network of scientists prepared to respond to complex global crises like pandemics [25]. The ISR's principles of "Science Without Borders" and "Equitable Resource Access" are critical here, as they ensure that researchers in countries and communities most affected by a crisis have access to the tools needed to contribute to solutions [24]. In the event of a public health emergency, the IBRR would swiftly grant vetted researchers worldwide access to restricted AI tools, ensuring that critical defensive research, such as vaccine or therapeutic development, is not hampered by security restrictions [16]. This is a crucial element for addressing the humanitarian aspect of AI governance, ensuring that security concerns do not come at the cost of global health equity. The proposal is designed to address a common point of friction, providing a clear pathway for collaboration and resource sharing when it is most needed.

3.4 Pillar IV: Verification and Compliance

The final pillar of the framework is dedicated to verification and compliance, learning from the successes and failures of past treaties. The framework would require member states to implement national oversight measures. It would also include provisions for information sharing on best practices and technical challenges, creating a community of trust among member states and technical experts.

Borrowing from the CWC, the framework would establish a formal process for addressing violations. This could include a technical secretariat to conduct inspections and a governing body with the power to recommend collective punitive measures, or, in cases of "particular gravity," to bring the issue before the UN Security Council [6]. This would ensure that the pact has teeth and that adherence is not merely a matter of voluntary commitment. The existence of such a robust enforcement mechanism would serve as a powerful deterrent against clandestine state-level programs.

4 Discussion

To assess the reality, we must navigate the competing imperatives of open science, national security, and global health equity. The proposed governance framework is designed to balance these often-conflicting priorities. At its core, the central tension arises from the principles of modern scientific enterprise: openness, transparency, and collaborative knowledge-sharing, clashing with the necessity to control access to technologies with catastrophic dual-use potential [8]. Addressing this tension is not merely a regulatory challenge; it is a matter of global stability, as AI-enabled biodesign tools have the unprecedented capacity to alter the biosecurity landscape within months, rather than decades.

The Fundamental Conflict The debate over data controls in AI biodesign illustrates this tension clearly. On one hand, some experts and policy proposals advocate restricting access to sensitive data, such as certain pathogen genomes, to preempt misuse [26]. On the other hand, overly restrictive policies risk slowing critical scientific discovery and undermining global collaboration [26]. The World Health Organization (WHO) underscores the importance of pathogen genome sharing for epidemic preparedness, demonstrating that unfettered openness can be essential for public health [26]. Any governance framework must therefore navigate this issue with nuance, creating structured pathways that allow responsible data sharing while limiting access to high-risk technologies. In practice, this means implementing tiered access, monitoring, and verification mechanisms that dynamically adapt to emerging threats and capabilities.

Balancing Innovation and Regulation A smart regulation" approach is critical, emphasizing proactive, ethical, and governance-driven strategies rather than reactive containment [22]. The proposed tiered framework addresses concerns that regulation could stifle innovation: low-risk, general-purpose models remain open for research, while controls are imposed only on high-capability

models with significant dual-use potential. By establishing clear rules and predictable pathways, researchers can plan projects responsibly without facing unnecessary barriers [7]. The managed access model and the International Bioscience Response Reserve (IBRR) further reinforce global equity by enabling researchers from the global south to safely participate in AI-enabled biodesign, thereby preventing an inequitable concentration of technological power in wealthy nations [24]. Moreover, these mechanisms allow for rapid scaling of beneficial applications during emergencies, ensuring that safety controls do not impede urgent public health responses.

Case Studies of AI's Positive Impact on Health and Medicine The transformative potential of AI in biology is already evident, underscoring the importance of frameworks that protect these benefits.

- **COVID-19 Vaccine Development:** AI significantly accelerated COVID-19 vaccine development [18]. Machine-learning algorithms analyzed vast viral genomic datasets to identify immunogenic targets, while computational models rapidly assessed molecular configurations for efficacy. This approach reduced the timeline from concept to clinical trials from years to months, demonstrating how AI can revolutionize epidemic preparedness and global health resilience [18].
- **Drug Discovery:** AI is reshaping drug discovery, traditionally slow, costly, and inefficient [12]. Models now accelerate target identification, lead discovery, and preclinical safety evaluation. For instance, Insilico Medicine used AI to take a novel fibrosis drug candidate to Phase 1 trials in roughly half the time of conventional methods [13]. Such successes highlight AI's potential to reduce human suffering, improve treatment accessibility, and enable precision medicine at scale.

These case studies illustrate that AI is not merely a hypothetical threat but an active force for global good. Without proper governance, however, the same capabilities that enable rapid vaccine development or drug discovery could be exploited to design high-risk pathogens or circumvent safety controls. The proposed framework, therefore, does not argue against AI-enabled biotechnology; rather, it advocates for its safe, responsible, and equitable advancement, ensuring that society captures the benefits of innovation while minimizing catastrophic risk. By integrating tiered regulation, controlled access, global coordination, and ethical oversight, this framework represents a blueprint for reconciling innovation with biosecurity imperatives, ultimately fostering a resilient and trustworthy biotechnology ecosystem.

5 Conclusion and Path Forward

The analysis indicates a clear and urgent need for a new international governance framework for AI-enabled biodesign tools. The traditional biosecurity playbook, with its reliance on static lists and its focus on non-state actors, is insufficient to address the threat of a sophisticated state adversary armed with generative AI. The historical precedent of state-level bioweapons programs, coupled with the inherent technical vulnerabilities of modern AI systems, necessitates a shift in focus to controlling the most powerful dual-use tools.

This report proposes a multi-stakeholder framework that is a hybrid of international arms control, modern export controls, and public health principles. By establishing a tiered registry, implementing managed access, and creating a mechanism for rapid response during emergencies, the framework addresses the critical balance between security and innovation. It acknowledges that the conflict between open science and national security is not an absolute trade-off but a manageable challenge that can be overcome through proactive, collaborative governance.

The time to act is now. Waiting for a catastrophic incident to trigger regulation is a dangerous and irresponsible strategy. The responsibility to proactively shape the governance of these creations falls to the technical AI community, policymakers, and scientists. This demands urgent attention, international engagement with a diverse range of stakeholders, and decisive action [4]. To move this proposal from concept to reality, the following recommendations are offered:

- **For Governments and International Bodies:** Convene a multi-stakeholder summit to begin drafting a new international pact. Initiate pilot projects to test the feasibility of a tiered registry system and managed access platforms, gathering input from AI developers and biosecurity experts [22].

- **For the AI Community:** Proactively engage with policymakers and biosecurity specialists. Integrate ethical guardrails and safety reporting as a core part of the development and release lifecycle of new models. Implement voluntary managed access systems and model cards that transparently outline a tool's capabilities and risks [22].
- **For the Scientific Community:** Foster a security mindset in research and a culture of responsibility, recognizing the dual-use nature of their work and actively participating in cross-disciplinary forums to inform governance decisions [8].

By taking these steps, the international community can ensure that the transformative power of AI in biology is harnessed for global good, while building a robust defense against its most severe misuse.

References

[1] International gene synthesis consortium: Home. https://genesynthesisconsortium.org/, 2009. Accessed: 2025-08-26.

[2] Rfdiffusion tutorial – meiler lab. https://meilerlab.org/wp-content/uploads/2023/12/RFDiffusion_tutorial.pdf, December 2023. Accessed: 2025-08-26.

[3] Software – institute for protein design – university of washington. https://www.ipd.uw.edu/software/, August 2025. Accessed: 2025-08-26.

[4] Statement on biosecurity risks at the convergence of ai and the life sciences. https://www.nti.org/analysis/articles/statement-on-biosecurity-risks-at-the-convergence-of-ai-and-the-life-sciences/, July 2025. Published: July 17 2025; accessed: 2025-08-26.

[5] Georgia Adamson and Gregory C. Allen. Opportunities to strengthen u.s. biosecurity from ai-enabled bioterrorism: What policymakers should know. https://www.csis.org/analysis/opportunities-strengthen-us-biosecurity-ai-enabled-bioterrorism-what-policymakers-should, August 2025. Published: 2025-08-06; accessed: 2025-08-26.

[6] Arms Control Association. The chemical weapons convention (cwc) at a glance, 2024. Accessed: 2025-08-26.

[7] Sarah R. Carter, Nicole E. Wheeler, Christopher R. Isaac, and Jaime M. Yassif. Developing guardrails for ai biodesign tools, November 2024. Accessed: 2025-08-26.

[8] Center for AI Safety. Biosecurity and ai: Risks and opportunities. https://safe.ai/blog/biosecurity-and-ai-risks-and-opportunities, 2024. Accessed: 2025-08-26.

[9] Carlos J Costa, Manuela Aparicio, Sofia Aparicio, and Joao Tiago Aparicio. The democratization of artificial intelligence: Theoretical framework. *Appl. Sci. (Basel)*, 14(18):8236, September 2024.

[10] Singapore Customs. Controlled chemicals, 2025. Accessed: 2025-08-26.

[11] Jen Fernandez and Lloyd Lyall. New u.s. export controls on advanced computing items and artificial intelligence model weights: Seven key takeaways, January 2025. Accessed: 2025-08-26.

[12] Fábio J N Ferreira and Agnaldo S Carneiro. AI-driven drug discovery: A comprehensive review. *ACS Omega*, 10(23):23889–23903, June 2025.

[13] Gorkem Gencer. Ai in pharma: Use cases, success stories, and challenges in 2025, 2025. Accessed: 2025-08-26.

[14] Dan Hendrycks, Eric Schmidt, and Alexandr Wang. Nonproliferation — chapter 5 of *Superintelligence Strategy*, 2025. Accessed: 2025-08-26.

[15] Michael C. Horowitz and Lauren A. Kahn. Nuclear non-proliferation is the wrong framework for ai governance, June 2025. Accessed: 2025-08-26.

[16] International Federation of Red Cross and Red Crescent Societies (IFRC). Emergency health, 2025. Accessed: 2025-08-26.

[17] Johns Hopkins Center for Health Security. Nist ai 800-1, managing misuse risk for dual-use foundation models (request for comment). https://centerforhealthsecurity.org/sites/default/files/2025-04/NIST-AI-800-1-RFC-3.15.25-Johns-Hopkins-Center-for-Health-Security_edited.pdf, March 2025. Submitted March 15 2025; accessed August 26 2025.

[18] Priya Joi. Using ai from lab to jab: how did artificial intelligence help us develop and deliver covid-19 vaccines?, 2025. Accessed: 2025-08-26.

[19] PhD Kim, Raphael. History and scope of ai in biodesign. https://www.biodesign.academy/p/history-and-scope-of-ai-in-biodesign, August 2025. Accessed: 2025-08-26.

[20] Sarah Morgan. The double-edged sword: Opportunities and risks of ai in biosecurity. https://georgetownsecuritystudiesreview.org/2024/11/15/the-double-edged-sword-opportunities-and-risks-of-ai-in-biosecurity/, November 2024. Published: 2024-11-15; accessed: 2025-08-26.

[21] National Academies of Sciences, Engineering, and Medicine, Division on Earth and Life Studies, Board on Life Sciences, Board on Chemical Sciences and Technology, and Committee on Strategies for Identifying and Addressing Potential Biodefense Vulnerabilities Posed by Synthetic Biology. *Biodefense in the age of synthetic biology*. National Academies Press, Washington, D.C., DC, January 2019.

[22] Nuclear Threat Initiative. How to prevent ai-enabled bioterrorism. https://www.nti.org/risky-business/how-to-prevent-ai-enabled-bioterrorism/, December 2024. Published: 2024-12-19; accessed: 2025-08-26.

[23] Nian Ran, Yue Wang, and Richard Allmendinger. MOLLM: Multi-Objective large language model for molecular design – optimizing with experts. 2025.

[24] International Science Reserve. What we do, 2025. Accessed: 2025-08-26.

[25] International Science Reserve. Who we are, 2025. Accessed: 2025-08-26.

[26] Nicole E Wheeler. Responsible AI in biotechnology: balancing discovery, innovation and biosecurity risks. *Front. Bioeng. Biotechnol.*, 13:1537471, February 2025.

[27] Wikipedia contributors. Soviet biological weapons program, 2025. Accessed: 2025-08-26.