AI-enabled Biological Tool Policy Dashboard[1]

Pierce Manlangit
Asia Centre for Health Security

Georgie Hau Sorensen
University of Bristol

Elena Pedrosa Prats
Comillas Pontifical University

G.M.M Miftahul Alam Adib
Shahjalal University of Science and Technology

Aayushi Vishnoi

With
Apart Research

Abstract

The rapid development of AI-enabled biological tools is creating biosecurity risks that existing governance frameworks are not well designed to address. The Global Risk Index for AI-enabled Biological Tools (GRI) offers a framework for assessing AI-bio tools by misuse-relevant capability, maturity, and availability, but it does not disclose the specific finalist tools assessed or map governance coverage at the national level. This paper develops a proof-of-concept governance mapping approach focused on protein engineering, the GRI category identified as requiring immediate governance attention. We examine the United States and the United Kingdom because of their relevance to frontier AI governance, institutional role in AI safety, and high GRI contribution scores. Consistent with responsible disclosure norms, we assess tools at the category level rather than naming specific high-capability systems. Using a cross-database methodology based on EpochAI datasets, we identify 42 protein engineering AI models associated with US institutions and 11 associated with UK institutions. We then map 18 governance instruments, 14 from the US and 4 from the UK, against the protein engineering

[1] Research conducted at the AIxBio Hackathon, April 2026

category. We find that governance in both countries is fragmented, largely downstream of AI model outputs, and poorly calibrated to AI-generated protein design outputs that precede physical material production or synthesis-provider screening. In the US, the revocation of EO 14110 removed the most directly relevant AI-biosecurity executive instrument without a biosecurity-equivalent replacement. We present these findings through an interactive policy dashboard designed to support future expansion across additional AI-bio tool categories and countries.

1. Introduction

The emergence of AI-enabled biological tools capable of predicting protein structures, designing nucleic acid sequences, and synthesizing biomedical knowledge represents a qualitative shift in the biological risk landscape. Unlike traditional biosecurity threats, which typically involve controlled physical materials and established regulatory categories, AI-enabled tools operate largely in the informational domain, where existing controls are poorly calibrated (Eslami et al., 2025). A researcher can access frontier protein engineering models from a browser (Winn, 2026); the outputs may not require any regulated material transfer and may not trigger any licensing or reporting requirement(Baker & Church, 2024; Kosal, 2024).

The Global Risk Index for AI-enabled Biological Tools (GRI) (Webster et al., 2025) provides a rigorous framework for assessing these tools according to misuse-relevant capability, maturity, and availability. However, the GRI does not present the specific AI tools selected for assessment, nor does it analyze how or whether those tools are governed at the national level. Understanding this governance gap is essential for translating the GRI's risk signal into actionable policy.

This paper addresses that gap. We select the United States and United Kingdom as our initial focus countries for reasons that are both practical and analytically grounded. Practically, English-language legal and policy sources are far more accessible to our team, making a fast, scrappy proof-of-concept feasible within the project's timeframe. Analytically, both countries are home to leading frontier AI safety institutions including NIST, OSTP, and the recently established AI Safety Institutes, making their governance responses particularly consequential as potential models for other nations. They also rank first and third in GRI contributor scores, meaning they are central to any serious governance discussion regardless of scope.

Throughout this paper, we refer to AI-bio tools at the category level rather than naming specific systems. The GRI itself does not disclose its finalist tools by name, which we interpret as a deliberate infohazard mitigation choice. We also adopt the norm we adopt here since naming

specific high-capability tools in a governance gap analysis could function as a roadmap for misuse rather than a resource for policymakers.

Given our timeframe, v1 of this analysis focuses on a single tool category: protein engineering. This category was selected because the GRI flags it as having the highest number of required immediate action items among all categories assessed. The dashboard architecture is designed to extend to remaining categories (i.e. nucleic acid design, pathogen literature synthesis, drug and toxin discovery, and sequence-to-phenotype prediction), in future iterations.

Our main contributions are:

- A replicable cross-database methodology for identifying country-attributed protein engineering AI models, demonstrated using two EpochAI databases
- A structured governance assessment mapping 18 policy instruments across the US and UK against the protein engineering tool category, with explicit coverage ratings
- An interactive policy dashboard making these findings navigable by country and instrument type, designed to scale to additional countries and tool categories

2. Related Work

The GRI is the most direct predecessor to this work. It establishes the tool scoring methodology we build on and our contribution is the governance layer it does not address. Importantly, the GRI's decision not to name finalist tools is a disclosure choice we interpret as intentional and replicate here.

Work on AI biosecurity risk assessment, including studies examining the uplift potential of large language models and generative biology tools (Anthropic, 2025), has examined the capability side of this problem. Work from the Johns Hopkins Center for Health Security (Pannu et al., 2024) and the Nuclear Threat Initiative (Nuclear Threat Initiative, 2025) has examined biosecurity governance more broadly, though not with AI-tool-level granularity or grounding in a scored tool taxonomy.

On the governance side, the revocation of EO 14110 (Exec. Order No. 14,110, 2023), and its replacement with EO 14179 (Exec. Order No. 14,178, 2025) , which explicitly reorients US AI policy toward removing regulatory barriers rather than managing biosecurity risks, represents a significant recent shift that existing governance analyses have not yet captured. The UK's AI Safety Institute and its ongoing frontier model CBRN evaluations represent the most active government-level response to this category of risk, but have not yet been mapped against a specific protein engineering tool set.

The closest analogues to our approach are DURC policy analyses, but these predate the current AI-bio tooling landscape and do not address models operating upstream of the physical triggers

those frameworks were designed to catch. Our work differs by grounding the governance assessment in a specific, cross-validated tool set identified through a structured database methodology, rather than assessing AI capabilities in the abstract.

3. Methods

Our analysis was structured into multiple stages, as detailed below:

Stage 1: Tool identification via research affiliation.

Rather than relying on the GRI's undisclosed finalist list, we developed an independent method for identifying relevant protein engineering AI models attributed to US and UK institutions. Our analysis was based on a dataset compiled by EpochAI for protein engineering-specific benchmarking (Atanasov et al., 2026; Epoch AI, n.d.). The dataset included 361 protein engineering AI models, of which 103 models were manually attributed to a country of origin. This attribution was based on the affiliation of the corresponding author for each of the research publications describing each of the AI models. This method attributed a total of 42 AI models to US-based institutions, while 11 AI models were attributed to UK-based institutions.

Stage 2: Governance landscape review. For the United States and United Kingdom, we conducted a structured review of governance instruments across three tiers: (1) binding legal instruments (statutes, regulations, and executive orders), (2) official policy and strategic frameworks, and (3) agency guidance, notices, and standards. We identified 14 US instruments and 4 UK instruments as potentially relevant to the protein engineering tool category. For each instrument, we assessed: its legal status, its applicability to AI-generated biological outputs, and whether it creates any enforceable trigger point (licensing, reporting, or restriction) upstream of physical material production.

Dashboard. Findings were rendered in an interactive web-based policy dashboard enabling filtering by country and instrument type, with explicit gap flagging and links to primary sources. The architecture is designed to scale to additional tool categories and countries as the analysis expands.

4. Results

4.1 Tool landscape

Cross-database matching identified a markedly asymmetric tool landscape between the two countries. The US accounts for 42 exact-match protein engineering models in our working set,

reflecting its position as the leading GRI contributor country. The UK accounts for 11 models. This asymmetry is itself a relevant finding: the governance burden in the US is proportionally higher given the volume and diversity of tools being developed there.

Notable model family patterns: the RFdiffusion family (RFDiffusion2, RFdiffusion3, Motif RFdiffusion) appears in the protein engineering database but only the original RFdiffusion appears in the general tracking database, suggesting that specialized protein engineering tools are undercounted in broad AI model registries. This has implications for any governance approach that relies on model registries as a monitoring mechanism.

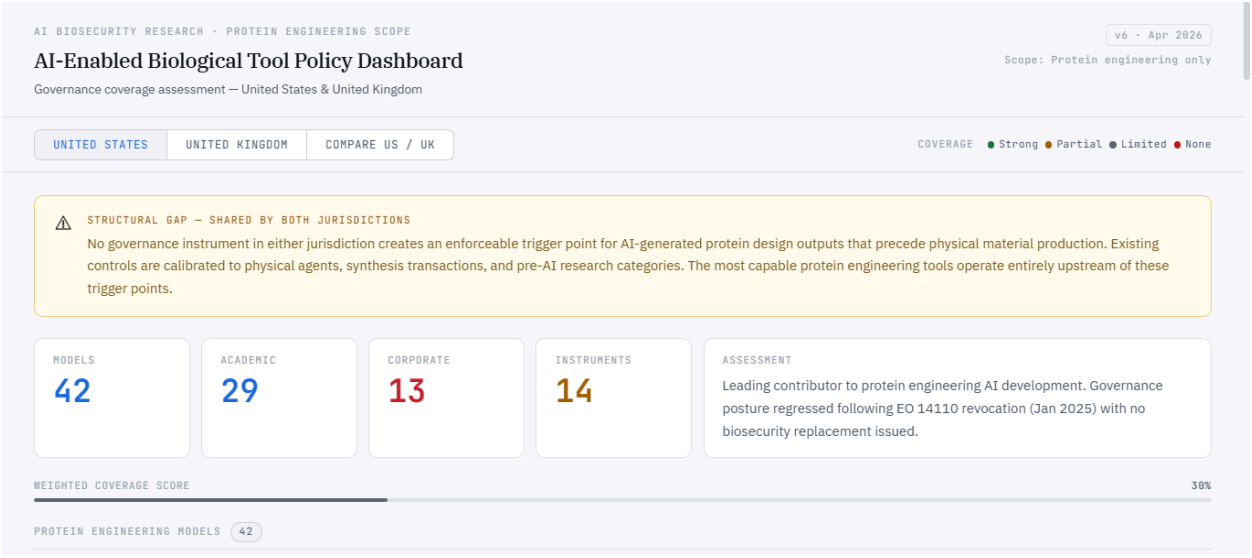

Figure 1. Screenshot of the AI-enabled biological tool policy dashboard used in this analysis. The complete tool, including the full set of entries and classifications, is provided in Appendix A.

4.2 US governance assessment

We identified 14 US governance instruments relevant to the protein engineering category. Table 1 summarizes the coverage classifications.

The most important change in the US policy baseline was the rescission of EO 14110. EO 14110 directed federal agencies, including OSTP and NIH, to assess AI-related biosecurity risks and develop nucleic acid synthesis screening standards (Exec. Order No. 14,110, 2023). It was revoked in January 2025 by EO 14148, not by EO 14179 (Exec. Order No. 14,148, 2025). EO 14179 then set a new federal AI policy direction, emphasizing US AI leadership and the removal of policies viewed as barriers to AI development (Exec. Order No. 14,179, 2025). EO 14292 later introduced a separate biological research safety agenda, including revision or replacement of the 2024 DURC/PEPP policy and the 2024 nucleic acid synthesis screening framework (Exec. Order No. 14,292, 2025). However, it does not create an AI-specific trigger for protein design model

outputs. We therefore treat the repeal of EO 14110 as a material weakening of the AI-biosecurity governance pathway, rather than as a complete removal of US biosecurity governance.

The OSTP Framework for Nucleic Acid Synthesis Screening and the HHS Screening Framework Guidance for Providers and Users of Synthetic Nucleic Acids are the closest federal instruments to the protein engineering risk category (Office of Science and Technology Policy, 2024; Screening Framework Guidance for Providers and Users of Synthetic Nucleic Acids, 2023). Both, however, operate downstream of AI model use. They apply when synthetic nucleic acids are procured or submitted for synthesis, not when an AI system generates a protein design, sequence, or candidate output. This leaves a gap for AI-generated outputs that are analyzed, shared, refined, or used in workflows that never pass through commercial synthesis.

The USG Policy for Oversight of DURC and PEPP and its implementation guidance are more directly relevant to biological research oversight (Office of Science and Technology Policy, 2024; The White House, 2024). Their scope, however, remains anchored to biological agents, toxins, pathogens, and defined research activities. They do not directly regulate AI-generated protein designs unless those outputs become part of covered research involving specified agents or activities. Because EO 14292 ordered the 2024 DURC/PEPP framework to be revised or replaced, we classify this area as relevant but unsettled.

The Commerce Control List under the Export Administration Regulations covers controlled commodities, software, and technology, including some dual-use biological items. It does not, however, clearly designate AI model weights or general protein-design software as controlled items for the purposes of this analysis (The Commerce Control List, 2026).

The NIST AI RMF provides a general, voluntary framework for AI risk management, while the NIST RFI on chemical and biological AI models signals federal interest in this specific risk category (Safety Considerations for Chemical and/or Biological AI Models, 2024; Tabassi, 2023). Neither instrument is binding, and neither creates a sector-specific control point for AI-generated protein engineering outputs.

California's SB 53 introduces obligations for large AI developers at the state level, but it is not biosecurity-specific and does not directly address protein engineering model deployment or AI-generated biological design outputs (Artificial Intelligence Models: Large Developers, 2025).

4.3 UK governance assessment

We identified four UK governance instruments relevant to the protein engineering category.

The UK Biological Security Strategy and its 2025 implementation report provide the strongest strategic recognition of AI-biosecurity risk in the UK policy landscape (Cabinet Office, 2025; HM Government, 2023). The Strategy frames biosecurity around understanding, preventing, detecting, and responding to biological risks, while the implementation report explicitly recognizes that

AI-enabled scientific tools and engineering biology create both opportunities and risks of misuse. However, these documents remain strategic rather than regulatory. They do not create an enforceable trigger for AI-generated protein design outputs.

The National AI Strategy addresses AI governance and innovation broadly, but it does not contain biosecurity-specific provisions for protein engineering or AI-generated biological outputs (Department for Science, Innovation and Technology et al., 2021).

The Code of Practice for the Cyber Security of AI is also relevant only indirectly. It sets out voluntary cyber-security measures for AI systems and organisations that develop or deploy them, but it addresses the security of AI systems rather than biosecurity risks arising from biological model outputs (Department for Science, Innovation and Technology, 2025).

4.4 Key shared gap: upstream outputs

The structural finding across both countries is consistent: no governance instrument in either jurisdiction creates an enforceable trigger point for AI-generated protein design outputs that precede physical material production. Existing controls are calibrated to physical agents, synthesis transactions, and defined research categories, all of which occur downstream of what protein engineering AI tools actually produce. The gap is not merely an absence of rules; it reflects a categorical mismatch between the risk architecture of AI-enabled protein engineering and the regulatory logic of existing biosecurity frameworks.

5. Discussion

Our findings suggest that the governance gap for AI-enabled protein engineering is structural rather than incidental. The revocation of EO 14110 in the US and the non-binding status of the UK's most relevant instruments mean that as of mid-2025, neither country has an enforceable governance mechanism specifically calibrated to AI-generated protein design outputs. The DURC and nucleic acid synthesis screening frameworks in the US come closest, but both operate at trigger points that AI design tools can precede entirely.

The asymmetry in tool counts (42 US models versus 11 UK models in our cross-database working set) suggests that governance urgency is disproportionately concentrated in the US, even as the UK has produced stronger strategic framing. This is a pattern worth monitoring as the country with the highest development activity currently has the weakest formal governance posture following the EO 14110 revocation.

The UK's AI Safety Institute remains the most institutionally promising development in either country, being a government body with an explicit frontier model CBRN evaluation mandate. Whether (a) evaluation produces binding governance recommendations and (b) on what timeline, are the open questions our dashboard is designed to track over time.

Limitations.

This analysis is bounded by several constraints. First, our source base is limited to English-language, publicly available documents; classified or informally-held guidance is not captured. Moreover, the governance landscape is actively evolving. The EO 14110 revocation occurred within our research window and illustrates how rapidly the picture can shift.

Another limitation of this analysis is the use of the EpochAI dataset, wherein multiple protein-specific frontier LLMs are categorised as "General purpose". Influential protein models such as ESM-2 or ESM-3 are therefore excluded from the analysed protein engineering AI models in this study. Further influential protein engineering AI tools, such as ProteinMPNN, are completely missing from the EpochAI dataset.

Rather than relying on direct GRI access, our manually curated dataset of AI models is a proxy for the GRI finalist pool, not identical to it. Fourth, our mapping is qualitative rather than scored; we have not assigned formal coverage ratings or conducted independent legal validation of instrument scope. Fifth, focusing on protein engineering means the full scope of the AI-bio governance gap is not yet visible. Sixth, focusing on protein engineering means the full scope of the AI-bio governance gap is not yet visible.

Future work.

A natural extension of this work would be to develop a formal coverage rating framework which assesses each instrument-category pairing on dimensions such as explicit scope, enforceability, trigger points, and institutional capacity. Such a rating framework would require validation by national biosecurity and AI governance legal experts, which is beyond the scope of this v1 analysis. Additional priorities include expanding coverage to the remaining four GRI tool categories; extending the country comparison set to China, Singapore, and EU member states; developing a formal confidence scoring layer to capture rating uncertainty; and building a living update mechanism as governance instruments evolve.

6. Conclusion

This paper presents an initial governance assessment of AI-enabled protein engineering tools in the United States and United Kingdom, grounded in a cross-database tool identification methodology and mapped against 18 governance instruments. We find that both countries have significant coverage gaps, particularly for AI-generated outputs that precede existing regulatory trigger points. Moreover there is no binding instrument in either country that specifically addresses AI-assisted protein engineering as a distinct governance category. The revocation of EO

14110 in the US represents an active regression in the governance posture of the world's leading protein engineering AI contributor country.

The interactive policy dashboard we present is designed as a scalable, living artifact: a foundation for ongoing monitoring as both the tool landscape and governance responses continue to evolve. We focus on protein engineering as our initial domain, but the framework applies across the full GRI tool spectrum. Closing the governance gap will require not just new instruments, but a reconceptualization of where in the AI-bio development pipeline regulatory scrutiny should apply. A cross-database methodology like the one we demonstrate here offers a tractable, replicable way to ground that reconceptualization in actual tools rather than abstract capabilities.

Code and Data

This project includes a public code repository hosted on GitHub [https://github.com/GeorgieHauSorensen/AI-enabled-Biological-Tool-Policy-Dashboard]. While the source code for the analysis and dashboard is publicly available, we have maintained a tiered disclosure approach due to the dual-use nature of the subject matter. To mitigate infohazard risks, the specific underlying data files containing sensitive tool-risk correlations have been withheld from the public repository. The interactive policy dashboard is available at [https://georgiehausorensen.github.io/AI-enabled-Biological-Tool-Policy-Dashboard/]. All governance instruments cited are linked directly within the dashboard and in the references section below.

Author Contributions

All authors contributed equally.

References

1. Anthropic. (2025, September 5). *Biorisk*. Red Teaming. https://red.anthropic.com/2025/biorisk/
2. Artificial Intelligence Models: Large Developers, Cal. S.B. 53, Chapter 138 (2025). https://leginfo.legislature.ca.gov/faces/billTextClient.xhtml?bill_id=202520260SB53
3. Atanasov, D., Zanichelli, N., & Denain, J.-S. (2026, April 26). *Expanding our analysis of biological AI models*. Epoch AI. https://epoch.ai/blog/expanding-our-analysis-of-biological-ai-models

4. Baker, D., & Church, G. (2024). Protein design meets biosecurity. *Science*, 383(6681), 349. https://doi.org/10.1126/science.ado1671
5. Cabinet Office. (2025, July 8). UK biological security strategy implementation report June 2023–June 2025. GOV.UK. https://www.gov.uk/government/publications/uk-biological-security-strategy-implementation-report-june-2023-june-2025/uk-biological-security-strategy-implementation-report-june-2023-june-2025-html
6. Department for Science, Innovation and Technology. (2025, January 31). *Code of practice for the cyber security of AI*. GOV.UK. https://www.gov.uk/government/publications/ai-cyber-security-code-of-practice/code-of-practice-for-the-cyber-security-of-ai
7. Department for Science, Innovation and Technology, Office for Artificial Intelligence, Department for Business, Energy & Industrial Strategy, & Department for Digital, Culture, Media & Sport. (2021, September 22). National AI strategy. GOV.UK. https://www.gov.uk/government/publications/national-ai-strategy
8. Epoch AI. (n.d.). *Parameter, compute and data trends in machine learning*. https://epoch.ai/data/ai-models-documentation
9. Eslami, M., Rose, E., & Nelson, C. (2025). Synthetic biology/AI convergence: Security threats in frontier science and regulatory challenges. *AI & Society*, 1–18. https://doi.org/10.1007/s00146-025-02576-4
10. Exec. Order No. 14,110, 88 Fed. Reg. 75,191 (November 1, 2023). https://www.federalregister.gov/documents/2023/11/01/2023-24283/safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence
11. Exec. Order No. 14,179, 90 Fed. Reg. 8,741 (January 31, 2025). https://www.federalregister.gov/documents/2025/01/31/2025-02172/removing-barriers-to-american-leadership-in-artificial-intelligence
12. Exec. Order No. 14,292, 90 Fed. Reg. 19,611 (May 8, 2025). https://www.federalregister.gov/documents/2025/05/08/2025-08266/improving-the-safety-and-security-of-biological-research
13. Global DNA Synthesis Map. (2026). *Global DNA synthesis map*. https://globalsynthesismap.bio/
14. HM Government. (2023, June). UK biological security strategy. https://assets.publishing.service.gov.uk/media/64c0ded51e10bf000e17ceba/UK_Biological_Security_Strategy.pdf
15. Kosal, M. E. (2024). Security challenges by AI-assisted protein design: The ability to design proteins in silico could pose a new threat for biosecurity and biosafety. *EMBO Reports*, 25(5), 2168–2171. https://doi.org/10.1038/s44319-024-00124-7
16. Lentzos, F., & Invernizzi, C. (2024). Artificial intelligence and biological misuse: Differentiating between information hazards and design hazards. arXiv. https://arxiv.org/pdf/2306.13952

17. National Institutes of Health. (2024, April). NIH guidelines for research involving recombinant or synthetic nucleic acid molecules (NIH guidelines). U.S. Department of Health and Human Services. https://osp.od.nih.gov/wp-content/uploads/NIH_Guidelines.pdf
18. National Institutes of Health. (2024, October 25). Notification of NIH requirements regarding procurement of synthetic nucleic acids and benchtop nucleic acid synthesis equipment (Notice No. NOT-OD-25-012). U.S. Department of Health and Human Services. https://grants.nih.gov/grants/guide/notice-files/NOT-OD-25-012.html
19. Nuclear Threat Initiative. (2025). A framework for managed access to biological AI tools. https://www.nti.org/analysis/articles/
20. Office of Science and Technology Policy. (2024, May 6). Implementation guidance for the United States government policy for oversight of dual use research of concern and pathogens with enhanced pandemic potential. The White House. https://bidenwhitehouse.archives.gov/wp-content/uploads/2024/05/USG-DURC-PEPP-Implementation-Guidance.pdf
21. Office of Science and Technology Policy. (2024, September). Framework for nucleic acid synthesis screening. U.S. Department of Health and Human Services. https://aspr.hhs.gov/S3/Documents/OSTP-Nucleic-Acid-Synthesis-Screening-Framework-Sep2024.pdf
22. OpenAI. (2025). *Building an early warning system for LLM-aided biological threat creation.* https://openai.com/index/building-an-early-warning-system-for-llm-aided-biological-threat-creation/
23. Pannu, J., Bloomfield, D., Zhu, A., MacKnight, R., Gomes, G., Cicero, A., & Inglesby, T. V. (2024, July 23). Prioritizing high-consequence biological capabilities in evaluations of artificial intelligence models. arXiv. https://arxiv.org/abs/2407.13059
24. Safety Considerations for Chemical and/or Biological AI Models, 89 Fed. Reg. 80886 (proposed October 4, 2024). https://www.federalregister.gov/documents/2024/10/04/2024-22974/safety-considerations-for-chemical-andor-biological-ai-models
25. Screening Framework Guidance for Providers and Users of Synthetic Nucleic Acids, 88 Fed. Reg. 70998 (October 13, 2023). https://www.federalregister.gov/documents/2023/10/13/2023-22540/screening-framework-guidance-for-providers-and-users-of-synthetic-nucleic-acids
26. Tabassi, E. (2023, January 26). Artificial intelligence risk management framework (AI RMF 1.0) (NIST AI 100-1). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.AI.100-1
27. The Commerce Control List, 15 C.F.R. § 774 (2026). https://www.ecfr.gov/current/title-15/subtitle-B/chapter-VII/subchapter-C/part-774
28. The White House Office of Science and Technology Policy. (2024, April). Framework for nucleic acid synthesis screening. U.S. Department of Health and Human Services.

https://aspr.hhs.gov/S3/Documents/OSTP-Nucleic-Acid-Synthesis-Screening-Framework-508.pdf
29. The White House. (2024, May). United States Government policy for oversight of dual use research of concern and pathogens with enhanced pandemic potential implementation guidance. https://bidenwhitehouse.archives.gov/wp-content/uploads/2024/05/USG-DURC-PEPP-Implementation-Guidance.pdf
30. The White House. (2025, July). Winning the AI race: America's AI action plan. https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf
31. U.S. Food and Drug Administration. (2025, January 7). Artificial intelligence-enabled device software functions: Lifecycle management and marketing submission recommendations. U.S. Department of Health and Human Services. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/artificial-intelligence-enabled-device-software-functions-lifecycle-management-and-marketing
32. U.S. Food and Drug Administration. (2026, January 29). *Clinical decision support software: Guidance for industry and Food and Drug Administration staff*. https://www.fda.gov/media/109618/download
33. Webster, J., Smith, A., Jones, B., & Doe, C. (2025). *Global risk index for AI-enabled biological tools*. The Centre for Long-Term Resilience; RAND Europe. https://doi.org/10.71172/wjyw-6dyc
34. Winn, Z. (2026, April 17). Bringing AI-driven protein-design tools to biologists everywhere. *MIT News*. https://news.mit.edu/2026/bringing-ai-driven-protein-design-tools-everywhere-0417

Appendix

Limitations. This analysis is bounded by several constraints. First, our source base is limited to English-language, publicly available documents; classified or informally-held guidance is not captured. Moreover, the governance landscape is actively evolving. The EO 14110 revocation occurred within our research window and illustrates how rapidly the picture can shift.

Another limitation of this analysis is the use of the EpochAI dataset, wherein multiple protein-specific frontier LLMs are categorised as "General purpose". Influential protein models such as ESM-2 or ESM-3 are therefore excluded from the analysed protein engineering AI models in this study. Further influential protein engineering AI tools, such as ProteinMPNN, are completely missing from the EpochAI dataset.

Rather than relying on direct GRI access, our manually curated dataset of AI models is a proxy for the GRI finalist pool, not identical to it. Fourth, our mapping is qualitative rather than scored; we

have not assigned formal coverage ratings or conducted independent legal validation of instrument scope. Fifth, focusing on protein engineering means the full scope of the AI-bio governance gap is not yet visible. Sixth, focusing on protein engineering means the full scope of the AI-bio governance gap is not yet visible.

Future work.

A natural extension of this work would be to develop a formal coverage rating framework which assesses each instrument-category pairing on dimensions such as explicit scope, enforceability, trigger points, and institutional capacity. Such a rating framework would require validation by national biosecurity and AI governance legal experts, which is beyond the scope of this v1 analysis. Additional priorities include expanding coverage to the remaining four GRI tool categories; extending the country comparison set to China, Singapore, and EU member states; developing a formal confidence scoring layer to capture rating uncertainty; and building a living update mechanism as governance instruments evolve.

Dual-use risks. A governance gap analysis that names specific high-capability tools could function as a targeting document for bad actors seeking to identify unregulated development pathways. We mitigate this by assessing tools at the category level only and by not characterizing individual models as high-risk in isolation. The cross-database overlap is used solely to bound the scope of governance analysis.

Responsible disclosure. We have not identified specific vulnerabilities in existing governance instruments beyond what is already publicly documented in policy literature. Our findings are intended to inform policymakers and researchers, not to publicize exploitable gaps. Further, the governance landscape was assessed as of April 26, 2026 and subsequent developments may materially alter the picture.

Ethical considerations. Our methodology relies entirely on publicly available sources. No proprietary model weights, unpublished research, or non-public government documents were accessed.

LLM Usage Statement

Claude (Anthropic) was used to support drafting, structural iteration, and integration of team findings throughout the writing process. All governance instrument citations, tool identification findings, and coverage assessments reflect the team's independent research. The final submission reflects the team's analytical judgments.

AI BIOSECURITY RESEARCH · PROTEIN ENGINEERING SCOPE

AI-Enabled Biological Tool Policy Dashboard

Governance coverage assessment — United States & United Kingdom

v6 · Apr 2026

Scope: Protein engineering only

UNITED STATES | UNITED KINGDOM | COMPARE US / UK

COVERAGE ● Strong ● Partial ● Limited ● None

⚠ STRUCTURAL GAP — SHARED BY BOTH JURISDICTIONS

No governance instrument in either jurisdiction creates an enforceable trigger point for AI-generated protein design outputs that precede physical material production. Existing controls are calibrated to physical agents, synthesis transactions, and pre-AI research categories. The most capable protein engineering tools operate entirely upstream of these trigger points.

MODELS	ACADEMIC	CORPORATE	INSTRUMENTS	ASSESSMENT
42	29	13	14	Leading contributor to protein engineering AI development. Governance posture regressed following EO 14110 revocation (Jan 2025) with no biosecurity replacement issued.

WEIGHTED COVERAGE SCORE 30%

PROTEIN ENGINEERING MODELS 42

MODEL	INSTITUTION / ORGANISATION	SECTOR
ProteinGenerator	University of Washington	ACADEMIC
JAM-2	Nabla Bio	CORPORATE
Proteina 200M (+15M triangle layers)	NVIDIA	CORPORATE
ProteinZero	University of Illinois	ACADEMIC
Proteina 400M	NVIDIA	CORPORATE
Proteina 200M	NVIDIA	CORPORATE
RFAntibody	University of Washington	ACADEMIC
Chai-2	Chai Discovery	CORPORATE
GPT-4b micro	OpenAI	CORPORATE
Germinal	Stanford University	ACADEMIC
RFDiffusion2	University of Washington	ACADEMIC
ProGen3 46B	Profluent Bio	CORPORATE
AF2Ring	University of Missouri	ACADEMIC

CSDesign	Brigham Young University	ACADEMIC
aSAMc	Michigan State University	ACADEMIC
Deep Kernel Inversion (DKI)	Nosis Bio	CORPORATE
ClusPro Team protein–protein docking protocol	Stony Brook University	ACADEMIC
ASR-Max	Michigan State University	ACADEMIC
BindCraft	University of Washington	ACADEMIC
Caliby	Stanford University	ACADEMIC
AF3Score	Stony Brook University / Boston University	ACADEMIC
BayesDesign	Brigham Young University	ACADEMIC
AIDO.Protein-16B	GenBio AI	CORPORATE
AIDO.Protein-IFdiff	GenBio AI	CORPORATE
DFMDock	MIT / Broad Institute	ACADEMIC
BaseFold	Basecamp Research	CORPORATE
ASR-Dist	Michigan State University	ACADEMIC

LASErMPNN	Harvard University	ACADEMIC
Metalic-AuxIF	InstaDeep	CORPORATE
EvoSeq-ML	Michigan State University	ACADEMIC
Loop-Diffusion	University of Washington	ACADEMIC
GenSLM	California Institute of Technology	ACADEMIC
PRISM (str. enc. only)	Carnegie Mellon University	ACADEMIC
PRISM	Carnegie Mellon University	ACADEMIC
TENet	Stanford University	ACADEMIC
β-GAN	Georgia Institute of Technology	ACADEMIC
Stable FT-AF2	University of California	ACADEMIC
Stable Structure-Split FT-AF2	University of California	ACADEMIC
SuperWater	Vanderbilt University	ACADEMIC
TrMRF	MIT	ACADEMIC
structure-conditioned sequence generative model	Generate Biomedicines	CORPORATE

Twin Peaks	Iowa State University	ACADEMIC

GOVERNANCE INSTRUMENTS 14

INSTRUMENT	TYPE	COVERAGE
~~EO 14110 — Safe, Secure & Trustworthy AI~~ Directed OSTP/NIH to assess AI biosecurity risks. Revoked Jan 2025 by EO 14179. No biosecurity replacement issued. ↗ Primary source	Executive Order	~~REVOKED~~
EO 14179 — Removing Barriers to American Leadership in AI Current operative EO. Reorients US AI policy toward removing regulatory barriers. No biosecurity provisions. ↗ Primary source	Executive Order	NONE
America's AI Action Plan Tied to EO 14179. Competitiveness-focused. No protein engineering governance provisions. ↗ Primary source	Official Policy	NONE
OSTP Nucleic Acid Synthesis Screening Framework (Sep 2024) Most proximate instrument. Regulates synthesis providers downstream of AI model output — does not address AI-generated protein design upstream. ↗ Primary source	Official Policy	PARTIAL
HHS Screening Framework Guidance for Synthetic Nucleic Acids (Oct 2023) Applies at synthesis transaction, not at model output. ↗ Primary source	Federal Guidance	PARTIAL
NIH Notice NOT-OD-25-012 — OSTP Framework Adherence		

NIH Notice NOT-OD-25-012 — OSTP Framework Adherence

Applies OSTP screening to NIH-funded procurements only.

↗ Primary source

NIH Notice | PARTIAL

NIH Guidelines — Recombinant or Synthetic Nucleic Acids (Apr 2024)

Governs recombinant nucleic acid research. AI-assisted protein design not explicitly within scope.

↗ Primary source

NIH Guidelines | PARTIAL

USG Policy for Oversight of DURC and PEPP (May 2024)

Anchored to listed agents and pre-AI research types. Does not create a trigger for AI-generated outputs.

↗ Primary source

Official Policy | LIMITED

DURC and PEPP Implementation Guidance (May 2024)

Implements USG DURC/PEPP policy. Same scope limitations.

↗ Primary source

Official Policy | LIMITED

15 CFR Part 774 — Commerce Control List (EAR)

AI model weights and protein design software not listed as controlled items.

↗ Primary source

CFR Regulation | LIMITED

NIST AI RMF 1.0 (2023)

General AI risk framework. No biosecurity-specific controls. Bioscience sector playbook not published.

↗ Primary source

NIST Framework | LIMITED

NIST RFI — Safety Considerations for Chemical/Biological AI Models (Oct 2024)

First formal US acknowledgment that biological AI models need dedicated safety consideration. Still information-gathering.

Federal RFI | LIMITED

NIST RFI — Safety Considerations for Chemical/Biological AI Models (Oct 2024)

First formal US acknowledgment that biological AI models need dedicated safety consideration. Still information-gathering.

↗ Primary source

Federal RFI | LIMITED

FDA Draft Guidance — AI-Enabled Device Software Functions

Applies to medical AI devices only. Not applicable to research protein engineering tools.

↗ Primary source

FDA Guidance | LIMITED

California SB 53 — Large AI Developer Obligations

State-level large developer obligations. Not biosecurity-specific.

↗ Primary source

State Legislation | LIMITED

Research note. Coverage ratings are preliminary assessments based on English-language public sources only. Not legal advice. Tool set cross-referenced from two EpochAI databases. v1 scope: protein engineering only.

AI BIOSECURITY RESEARCH · PROTEIN ENGINEERING SCOPE

AI-Enabled Biological Tool Policy Dashboard

Governance coverage assessment — United States & United Kingdom

v6 · Apr 2026

Scope: Protein engineering only

UNITED STATES | UNITED KINGDOM | COMPARE US / UK

COVERAGE ● Strong ● Partial ● Limited ● None

⚠ STRUCTURAL GAP — SHARED BY BOTH JURISDICTIONS

No governance instrument in either jurisdiction creates an enforceable trigger point for AI-generated protein design outputs that precede physical material production. Existing controls are calibrated to physical agents, synthesis transactions, and pre-AI research categories. The most capable protein engineering tools operate entirely upstream of these trigger points.

MODELS	ACADEMIC	CORPORATE	INSTRUMENTS	ASSESSMENT
11	3	8	6	Third-ranked contributor. Stronger strategic framing than the US — UK Biological Security Strategy explicitly names AI — but no binding regulatory instrument specific to protein engineering exists.

WEIGHTED COVERAGE SCORE 30%

PROTEIN ENGINEERING MODELS 11

MODEL	INSTITUTION / ORGANISATION	SECTOR
DynamicMPNN	University of Cambridge	ACADEMIC

DynamicMPNN	University of Cambridge	ACADEMIC
AlphaProteo	Google DeepMind	CORPORATE
CPDiffusion	Cambridge University / Shanghai Jiao Tong University	ACADEMIC
AlphaProteo v1	Google DeepMind	CORPORATE
BAGEL	Imperial College London	ACADEMIC
AbBFN	InstaDeep	CORPORATE
AbBFN+	InstaDeep	CORPORATE
BaseFold	Basecamp Research	CORPORATE
Metalic-AuxIF	InstaDeep	CORPORATE
ProtBFN	InstaDeep	CORPORATE
Zero-shot MSA Transformer	Cambridge Consultants	CORPORATE

GOVERNANCE INSTRUMENTS 7

INSTRUMENT	TYPE	COVERAGE
UK Biological Security Strategy (2023) Explicitly names AI as a biosecurity consideration. Calls for assessment of AI tools enabling pathogen design. Binding rules still pending.	Official Policy	PARTIAL

UK Biological Security Strategy Implementation Report (Jun 2023–Jun 2025)

Confirms ongoing work. No binding regulatory outcomes for protein engineering specifically reported.

↗ Primary source

Official Policy | PARTIAL

UK AI Safety Institute — Frontier AI Safety Framework

AISI mandate focuses on frontier AI systems broadly. No biosecurity-specific protein engineering assessment published.

↗ Primary source

Official Policy | LIMITED

UK National AI Strategy

Broad AI governance. No biosecurity-specific provisions or protein engineering risk categories.

↗ Primary source

Official Policy | LIMITED

Code of Practice for Cyber Security of AI

Addresses security of AI systems, not biosecurity risks from AI outputs.

↗ Primary source

Official Policy | NONE

Seoul Declaration on AI Safety (2024)

Non-binding commitment to CBRN risk cooperation. No implementation mechanism for protein engineering.

↗ Primary source

International Declaration | LIMITED

ISO 42001

General AI management system standard. No biosecurity-specific controls.

Technical Standard | LIMITED

Research note. Coverage ratings are preliminary assessments based on English-language public sources only. Not legal advice. Tool set cross-referenced from two EpochAI databases. v1 scope: protein engineering only.

AI BIOSECURITY RESEARCH · PROTEIN ENGINEERING SCOPE

AI-Enabled Biological Tool Policy Dashboard

Governance coverage assessment — United States & United Kingdom

v6 · Apr 2026

Scope: Protein engineering only

UNITED STATES | UNITED KINGDOM | COMPARE US / UK

COVERAGE ● Strong ● Partial ● Limited ● None

⚠ STRUCTURAL GAP — SHARED BY BOTH JURISDICTIONS

No governance instrument in either jurisdiction creates an enforceable trigger point for AI-generated protein design outputs that precede physical material production. Existing controls are calibrated to physical agents, synthesis transactions, and pre-AI research categories. The most capable protein engineering tools operate entirely upstream of these trigger points.

UNITED STATES

MODELS	ACADEMIC	CORPORATE	INSTRUMENTS	ASSESSMENT
42	29	13	14	Leading contributor to protein engineering AI development. Governance posture regressed following EO 14110 revocation (Jan 2025) with no biosecurity replacement issued.

WEIGHTED COVERAGE SCORE 30%

MODELS 42

UNITED KINGDOM

MODELS	ACADEMIC	CORPORATE	INSTRUMENTS	ASSESSMENT
11	3	8	6	Third-ranked contributor. Stronger strategic framing than the US — UK Biological Security Strategy explicitly names AI — but no binding regulatory instrument specific to protein engineering exists.

WEIGHTED COVERAGE SCORE 30%

MODEL	INSTITUTION / ORGANISATION	SECTOR
ProteinGenerator	University of Washington	ACADEMIC
JAM-2	Nabla Bio	CORPORATE
Proteina 200M (+15M triangle layers)	NVIDIA	CORPORATE
ProteinZero	University of Illinois	ACADEMIC
Proteina 400M	NVIDIA	CORPORATE
Proteina 200M	NVIDIA	CORPORATE
RFAntibody	University of Washington	ACADEMIC
Chai-2	Chai Discovery	CORPORATE
GPT-4b micro	OpenAI	CORPORATE
Germinal	Stanford University	ACADEMIC
RFDiffusion2	University of Washington	ACADEMIC
ProGen3 46B	Profluent Bio	CORPORATE
AF2Ring	University of Missouri	ACADEMIC

MODELS 11

MODEL	INSTITUTION / ORGANISATION	SECTOR
DynamicMPNN	University of Cambridge	ACADEMIC
AlphaProteo	Google DeepMind	CORPORATE
CPDiffusion	Cambridge University / Shanghai Jiao Tong University	ACADEMIC
AlphaProteo v1	Google DeepMind	CORPORATE
BAGEL	Imperial College London	ACADEMIC
AbBFN	InstaDeep	CORPORATE
AbBFN+	InstaDeep	CORPORATE
BaseFold	Basecamp Research	CORPORATE
Metalic-AuxIF	InstaDeep	CORPORATE
ProtBFN	InstaDeep	CORPORATE
Zero-shot MSA Transformer	Cambridge Consultants	CORPORATE

INSTRUMENTS 7

CSDesign	Brigham Young University	ACADEMIC
aSAMc	Michigan State University	ACADEMIC
Deep Kernel Inversion (DKI)	Nosis Bio	CORPORATE
ClusPro Team protein–protein docking protocol	Stony Brook University	ACADEMIC
ASR-Max	Michigan State University	ACADEMIC
BindCraft	University of Washington	ACADEMIC
Caliby	Stanford University	ACADEMIC
AF3Score	Stony Brook University / Boston University	ACADEMIC
BayesDesign	Brigham Young University	ACADEMIC
AIDO.Protein-16B	GenBio AI	CORPORATE
AIDO.Protein-IFdiff	GenBio AI	CORPORATE
DFMDock	MIT / Broad Institute	ACADEMIC
BaseFold	Basecamp Research	CORPORATE
ASR-Dist	Michigan State University	ACADEMIC

INSTRUMENT	TYPE	COVERAGE
UK Biological Security Strategy (2023) Explicitly names AI as a biosecurity consideration. Calls for assessment of AI tools enabling pathogen design. Binding rules still pending. ↗ Primary source	Official Policy	PARTIAL
UK Biological Security Strategy Implementation Report (Jun 2023–Jun 2025) Confirms ongoing work. No binding regulatory outcomes for protein engineering specifically reported. ↗ Primary source	Official Policy	PARTIAL
UK AI Safety Institute — Frontier AI Safety Framework AISI mandate focuses on frontier AI systems broadly. No biosecurity-specific protein engineering assessment published. ↗ Primary source	Official Policy	LIMITED
UK National AI Strategy Broad AI governance. No biosecurity-specific provisions or protein engineering risk categories. ↗ Primary source	Official Policy	LIMITED
Code of Practice for Cyber Security of AI Addresses security of AI systems, not biosecurity risks from AI outputs. ↗ Primary source	Official Policy	NONE

LASErMPNN	Harvard University	ACADEMIC
Metalic-AuxIF	InstaDeep	CORPORATE
EvoSeq-ML	Michigan State University	ACADEMIC
Loop-Diffusion	University of Washington	ACADEMIC
GenSLM	California Institute of Technology	ACADEMIC
PRISM (str. enc. only)	Carnegie Mellon University	ACADEMIC
PRISM	Carnegie Mellon University	ACADEMIC
TENet	Stanford University	ACADEMIC
β-GAN	Georgia Institute of Technology	ACADEMIC
Stable FT-AF2	University of California	ACADEMIC
Stable Structure-Split FT-AF2	University of California	ACADEMIC
SuperWater	Vanderbilt University	ACADEMIC
TrMRF	MIT	ACADEMIC
structure-conditioned sequence generative model	Generate Biomedicines	CORPORATE

Seoul Declaration on AI Safety (2024)

International Declaration LIMITED

Non-binding commitment to CBRN risk cooperation. No implementation mechanism for protein engineering.

↗ Primary source

ISO 42001

Technical Standard LIMITED

General AI management system standard. No biosecurity-specific controls.

Twin Peaks	Iowa State University	ACADEMIC

INSTRUMENTS 14

INSTRUMENT	TYPE	COVERAGE
~~**EO 14110 — Safe, Secure & Trustworthy AI**~~ Directed OSTP/NIH to assess AI biosecurity risks. Revoked Jan 2025 by EO 14179. No biosecurity replacement issued. ↗ Primary source	Executive Order	~~REVOKED~~
EO 14179 — Removing Barriers to American Leadership in AI Current operative EO. Reorients US AI policy toward removing regulatory barriers. No biosecurity provisions. ↗ Primary source	Executive Order	NONE
America's AI Action Plan Tied to EO 14179. Competitiveness-focused. No protein engineering governance provisions. ↗ Primary source	Official Policy	NONE
OSTP Nucleic Acid Synthesis Screening Framework (Sep 2024) Most proximate instrument. Regulates synthesis providers downstream of AI model output — does not address AI-generated protein design upstream. ↗ Primary source	Official Policy	PARTIAL

OSTP Nucleic Acid Synthesis Screening Framework (Sep 2024)

Most proximate instrument. Regulates synthesis providers downstream of AI model output — does not address AI-generated protein design upstream.

↗ Primary source

Official Policy

PARTIAL

HHS Screening Framework Guidance for Synthetic Nucleic Acids (Oct 2023)

Applies at synthesis transaction, not at model output.

↗ Primary source

Federal Guidance

PARTIAL

NIH Notice NOT-OD-25-012 — OSTP Framework Adherence

Applies OSTP screening to NIH-funded procurements only.

↗ Primary source

NIH Notice

PARTIAL

NIH Guidelines — Recombinant or Synthetic Nucleic Acids (Apr 2024)

Governs recombinant nucleic acid research. AI-assisted protein design not explicitly within scope.

↗ Primary source

NIH Guidelines

PARTIAL

USG Policy for Oversight of DURC and PEPP (May 2024)

Anchored to listed agents and pre-AI research types. Does not create a trigger for AI-generated outputs.

↗ Primary source

Official Policy

LIMITED

DURC and PEPP Implementation Guidance (May 2024)
Implements USG DURC/PEPP policy. Same scope limitations.
↗ Primary source

Official Policy | LIMITED

15 CFR Part 774 — Commerce Control List (EAR)
AI model weights and protein design software not listed as controlled items.
↗ Primary source

CFR Regulation | LIMITED

NIST AI RMF 1.0 (2023)
General AI risk framework. No biosecurity-specific controls. Bioscience sector playbook not published.
↗ Primary source

NIST Framework | LIMITED

NIST RFI — Safety Considerations for Chemical/Biological AI Models (Oct 2024)
First formal US acknowledgment that biological AI models need dedicated safety consideration. Still information-gathering.
↗ Primary source

Federal RFI | LIMITED

FDA Draft Guidance — AI-Enabled Device Software Functions
Applies to medical AI devices only. Not applicable to research protein engineering tools.
↗ Primary source

FDA Guidance | LIMITED

California SB 53 — Large AI Developer Obligations

State Legislation | LIMITED

California SB 53 — Large AI Developer Obligations State Legislation LIMITED

State-level large developer obligations. Not biosecurity-specific.

↗ Primary source

Research note. Coverage ratings are preliminary assessments based on English-language public sources only. Not legal advice. Tool set cross-referenced from two EpochAI databases. v1 scope: protein engineering only.