Responsible Biodesign Workshop: AI, Protein Design, and the Biosecurity Landscape – Recommended Actions

Sebastian Rivera [1], Moritz S. Hanke [2], Samuel Curtis [3], Neil Cherian [4], Anthony Gitter [5], Jeffrey J. Gray [6], Andrew Hebbeler [4], Stephen McCarthy [7], David Nannemann [8], Claire Qureshi [9], Brian D. Weitzner [10], Ian C. Haydon [11,*]

[1] Engineering Biology Research Consortium, USA
[2] Johns Hopkins Center for Health Security, USA
[3] Rosetta Commons, USA *(formerly, until January 2025)*
[4] Coalition for Epidemic Preparedness Innovations, NO
[5] University of Wisconsin-Madison, USA
[6] Johns Hopkins University, USA; Rosetta Commons, USA
[7] University of California, Irvine, USA
[8] Rosetta Design Group, USA
[9] Sentinel Bio, USA
[10] Outpace Bio, USA
[11] Institute for Protein Design, University of Washington, USA
*Correspondence: ichaydon@uw.edu

Abstract: This report presents Recommended Actions from the January 2025 Responsible Biodesign Workshop, which convened leading experts across AI-enabled biomolecular design and biosecurity policy. Building on existing community commitments for the Responsible Development of AI for Protein Design, the Recommended Actions aim to guide scientists, policy practitioners, and funding bodies in ensuring safe and beneficial development of AI-enabled biomolecular design tools. The Recommended Actions focus on advancing AI-Resilient nucleic acid synthesis security screening, assessing the risk-benefit landscape of biomolecular design capabilities, and building fora for sustained engagement between scientists and policy practitioners.

Keywords: Biosecurity; Biodesign; Protein Design; Biomolecular Design; Artificial Intelligence; Responsible Innovation; Nucleic Acid Synthesis Screening; Technology Policy

Background

The Responsible Biodesign Workshop, held on January 10th, 2025, in Washington, D.C., convened key stakeholders from biomolecular software development, biosecurity policy, and adjacent expert communities to advance implementation of the commitments made in the Community Values, Guiding Principles, and Commitments for the Responsible Development of AI for Protein Design statement [1] (hereafter "Community Statement"). Published on March 8, 2024, the Community Statement highlighted the great potential of biomolecular design for solving some of humanity's most pressing health, environmental, and energy challenges; it demonstrated a commitment from scientists who develop or use AI tools for biomolecular research to ensure these technologies are trustworthy and advance responsibly. The Community Statement was signed by over 170 leading scientists from around the world.

Below are proposed Recommended Actions—concrete next steps that aim to advance implementation of the commitments made in the Community Statement via collaborative efforts. They reflect areas of broad agreement reached during the Workshop's four sessions, which featured presentations and open discussions around biomolecular AI tools and the biosecurity

landscape, their applications for enhancing public health resilience, nucleic acid synthesis screening, and tangible next steps to advance responsible biodesign.

Participants agreed that AI-enabled biomolecular modeling tools have the transformative potential to accelerate and expand previous biomolecular design capabilities to help solve some of the world's most critical challenges. However, these accelerating AI-enabled biomolecular design capabilities may also heighten risks from harmful applications of these tools. Therefore, the Recommended Actions seek to promote the development of beneficial technologies while mitigating their potential to be misused, intentionally or otherwise, to cause harm.

Because the Community Statement is an expression of commitments by scientists who develop and use biomolecular design tools, the Recommended Actions primarily focus on steps ***scientists*** should take (herein referring to developers and users of biomolecular design tools and not to scientists from other disciplines). However, meaningfully advancing these commitments would greatly benefit from action by and coordination with organizations outside of these research communities, such as ***policy practitioners*** (policymakers, standard-setting bodies, civil society groups, etc.) and ***funding bodies*** (governmental and nongovernmental grantmaking bodies, research institutions, philanthropic groups, etc.). Thus, the Recommended Actions also put forward proposals for these actors.

Recommended Actions are divided into three categories:

1. Advance AI-Resilient Nucleic Acid Synthesis Security Screening
2. Assess the Risk-Benefit Landscape of Biomolecular Design Capabilities
3. Build Fora for Sustained Engagement between Scientists and Policy Practitioners

Workshop participants were given the opportunity to review and provide feedback on the Recommended Actions before publication. The content expressed herein does not necessarily reflect the consensus positions of, or endorsement by, the authors, Workshop contributors, and participants, or their respective institutions. The authors of this document aim to ensure a continuous effort to approach well-positioned, respective individuals and organizations from the scientist, policy practitioner, and funding body communities to operationalize and carry out the Recommended Actions below.

Recommended Actions

Advance AI-Resilient Nucleic Acid Synthesis Security Screening

(Community Statement Commitments 3, 4)

Nucleic acid synthesis is critical to modern life sciences research, and the plummeting cost of custom DNA molecules is enabling breakthroughs in medicine and engineering. However, as technological capabilities advance and synthesized DNA becomes more widely accessible, securing the digital-physical interface against unchecked availability of DNA encoding "sequences of concern" (SOCs) will require advanced security systems and international cooperation to implement universally enforced security standards. In this context, sequences of concern are those that could confer toxicity or enhanced transmissibility, pathogenicity, or virulence to an organism or directly encode a biotoxin. As the key process required to materialize most computationally designed biomolecules, securing nucleic acid synthesis capabilities is critical for

biosecurity.

Currently, most nucleic acid synthesis screening tools rely primarily on sequence homology to known sequences. As biomolecular design tools continue to rapidly advance, particularly those that can re-design sequences and create *de novo* sequences, the prevalence of SOCs that have little to no similarity to known sequences will rise. New screening tools are needed for assessing the risks associated with these novel sequences. Such tools would have broad biosecurity applications, particularly in critical industries like commercial nucleic acid synthesis, but also in biosurveillance. As demonstrated by a recently released report [2], screening systems can be strengthened to detect AI-generated SOCs, such as novel sequences and functional homologs that have little similarity to known sequences. Updated synthesis security methods, however, should also take into account that the vast majority of AI-generated proteins will have legitimate, beneficial purposes and should not unduly create barriers to researchers obtaining these sequences.

A. Develop and implement next-generation screening algorithms for detecting AI-generated sequences of concern

Scientists should leverage their unique technical expertise to contribute to the development of next-generation screening software (e.g., sequence-to-function prediction tools) and closely work with respective nucleic acid synthesis providers, especially toward the analysis and detection of AI-designed SOCs. Such defensive tools could also be beneficial for up- and downstream biosecurity measures, for instance, in the context of pathogen surveillance. Given that these tools will be used to predict harmful functions from sequence, these tools should be developed with appropriate security measures as well.

Policy practitioners should pursue policy options to advance the security of next-generation screening software (regarding SOCs and customers) and hardware (particularly benchtop synthesis devices) and strive towards harmonized, universal, international adoption. These policy options should include means for identifying and disclosing vulnerabilities in screening tools to tool developers (particularly relevant for ***Scientists*** due to their technical expertise in biomolecular design tool usage) so that they can rapidly address them.

Funding bodies should sponsor efforts, such as pilot projects and competitions, to develop next-generation security solutions, especially regarding vulnerabilities introduced to current safeguards by biomolecular design tools.

B. Advance nucleic acid synthesis security as an international policy priority

Scientists should publish a statement featuring a broad international coalition in support of harmonized international nucleic acid synthesis security. It should include support for the Office of Science and Technology Policy's Framework for Nucleic Acid Synthesis Screening and the Administration for Strategic Preparedness and Response's Screening Framework Guidance for Providers and Users of Synthetic Nucleic Acids (or similar future guidance) and advocate for regulatory requirements for nucleic acid synthesis screening. This statement should also endorse the biosecurity recommendations in the International Gene Synthesis Consortium's standard ISO 202688-2:2024, acknowledge and support the United Kingdom's screening guidance on synthetic nucleic acids and New Zealand's Gene

Technology bill, and call on other governments to adopt similar frameworks and guidance.

Policy practitioners should organize efforts to coordinate international agreement on the implementation of nucleic acid synthesis screening best practices. This should include convenings at various strategic geographic locations to ensure participation from a broad international coalition.

Funding bodies should sponsor efforts by ***policy practitioners*** to support international agreement and implementation of nucleic acid synthesis screening best practices. They should also develop commitments integrating nucleic acid synthesis security practices, like a requirement to only acquire nucleic acids from providers who screen, in their funding requirements.

Assess the Risk-Benefit Landscape of Biomolecular Design Capabilities
(Community Statement Commitments 1, 2, 5, 6, 7, 8, 9)

Scientists, policy practitioners, and funding bodies share a common interest in maximizing the benefits and mitigating potential risks of AI-enabled biomolecular design. However, there is little shared, detailed understanding of the potential impacts of current and future research on the risk-benefit landscape and how proposed or adopted policy frameworks may affect scientists. As is inherent to dual-use technologies, capabilities that can provide benefits to society through legitimate research and use also carry risks of causing societal harm through deliberate or accidental misuse. The risk assessment of dual-use technology is challenging, given the overlap of beneficial and harmful capabilities. Over the past decade, researchers and policymakers have thoroughly engaged with nuanced biosecurity risk-benefit tradeoffs, leading to multiple guidelines and policies, most recently the "United States Government Policy for Oversight of Dual Use Research of Concern and Pathogens with Enhanced Pandemic Potential" (DURC/PEPP). While this policy focuses primarily on physical life sciences research, it includes a section encouraging voluntary institutional oversight of *in silico* DURC. To inform potential future policies, it will be essential to incorporate biomolecular design scientists' perspectives and ensure thorough analyses of risk-benefit tradeoffs.

Frameworks for rigorous risk-benefit analyses of specific biomolecular design capabilities and their applications could provide a foundation for productive dialogue on advancing benefits while minimizing risk. Additionally, careful threat modeling is needed, for instance, to consider if and how biomolecular design tools add real-world misuse potential over already available methods. Through these analyses, scientists, in collaboration with biosecurity practitioners, could identify specific 'capabilities of concern' that pose the greatest risk of harm under specific threat models, as well as the research products (e.g., datasets, source code, model weights) that could enable them. In accompanying efforts to identify such harmful applications, it is pertinent to prioritize realistic large-scale societal harms, for instance, to health, national security, the economy, or the environment through pathogens, toxins, or other biological agents. Stakeholders need to be conscious of weighing the misuse risk of such capabilities against the plethora of important beneficial applications to society, their potential to reduce large-scale harm, and the risks of slowing down and limiting access to beneficial research applications. Such a shared understanding would support the development of pre- and post-development risk-benefit review processes and evaluation methods akin to existing DURC/PEPP policies, as well as tools that monitor and forecast

future capabilities to help scientists ensure their research is conducted responsibly. Additionally, such efforts would inform and contribute to existing efforts spelling out potential risk mitigation strategies without unduly impeding beneficial research.

C. **Assess the potential health and economic benefits and risks of AI-enabled biomolecular design research.**

   Funding bodies should commission studies based on existing or original research involving qualitative and quantitative risk-benefit analyses and bioethical considerations of current and anticipated technological capabilities enabled by advances in biomolecular design research (involving input from ***scientists*** and ***policy practitioners***, among others). This effort should also identify potential priority research areas that could yield significant benefits with minimal risks.

D. **Identify "capabilities of concern" and corresponding research products that warrant pre- and post-development risk-benefit review.**

   Scientists and ***policy practitioners*** should collaborate to precisely and continuously define and identify "capabilities of concern"—specific technological capabilities enabled by advances in biomolecular design research that would meaningfully increase the risk of large-scale biological harm if misused—and the research products that could give rise to such capabilities (datasets, etc.). This should be accompanied by an effort to clarify what exact large-scale biological harm warrants concern and, conversely, an effort to define "capabilities without applications of concern" where research would not require pre- and post-development risk-benefit review. A key aspect of this work should also involve the development of shared language and terminology.

E. **Coordinate an interdisciplinary effort to develop pre- and post-development review processes to identify and evaluate capabilities of concern in relevant tools and explore developing and recommending proportional risk mitigation measures.**

   Scientists and ***policy practitioners*** should conduct an interdisciplinary effort to develop and test streamlined pre- and post-development risk-benefit review processes that support scientists in determining whether their research enables capabilities of concern and support scientists in, if needed, taking proportional measures to mitigate risks and enable benefits for such tools. As part of this effort, scientists should develop rigorous evaluation methods for quantitatively assessing capabilities of concern and explore the implementation and evaluate the feasibility and efficacy of potential risk mitigation measures that do not limit benefits or stall innovation and are compatible with the needs of scientists.

Build Fora for Sustained Engagement between Scientists and Policy Practitioners
(Community Statements Commitments 5, 7, 9)

Sustained engagement between experts is needed to enable more effective translation of technical insights into policy frameworks and help scientists better understand and engage with policy processes. While the Workshop successfully kick-started discussions, it also highlighted areas that need thorough downstream engagement. These include: 1) increasing understanding of biosecurity threat models and existing DURC, security, and AI governance frameworks and discussions for scientists 2) informing policy practitioners about beneficial applications, realities of

biomolecular design research, current tool capabilities and future trajectories, and 3) discussing digital and physical biosecurity vulnerabilities, detailed risk-benefit analyses and concrete risk mitigation measures. Creating structured opportunities for such meaningful information exchange and collaboration between scientists and policy practitioners will aid in improving our collective understanding of risk and in developing technical solutions and well-informed policies and processes that support the pursuit of beneficial research and development while proportionally guarding against potential risks from accidental or deliberate misuse.

F. Establish a scientists' "Responsible AIxBiodesign" working group that engages in policy development, technical advice, and analysis.

Scientists should develop and launch a working group (potentially akin to the Frontier Model Forum) that engages in policy development processes around responsible biomolecular design and serves as a point of contact for policy practitioners and researchers engaged in relevant work. Given their deep technical expertise, such a working group should advise on technical questions that inform policies and explore risk mitigation strategies that are compatible with ***scientists'*** priorities around accessibility, openness, and reproducibility.

G. Establish a research advisory committee to assist scientists with questions or concerns about the safety, security, or ethical implications of their work.

Scientists and ***policy practitioners*** should create an interdisciplinary committee comprising biomolecular design scientists, virologists, immunologists, security practitioners, and ELSI (ethical, legal, and social implications) experts to support scientists in navigating decisions related to their research conceptualization, development, publication, and release (particularly pre- and post-development risk-benefits assessment and proportional risk mitigation, see **E.**). Over time, this committee can develop private and public guidance on resulting best practices and support educational efforts and coursework for scientists around biorisk and security, DURC, and responsible conduct.

H. Facilitate policy professionals' participation in scientific conferences, scientists' involvement in policy fora, as well as conducting joint events.

Scientists should invite policy professionals to attend and participate in relevant scientific conferences and ensure efforts around responsible biomolecular design research are represented on the agenda.

Policy practitioners should invite scientists to attend and participate in pertinent fora for biomolecular design.

Scientists and ***Policy practitioners*** should regularly convene at dedicated, multi-stakeholder events (like this Workshop) to advance responsible biomolecular design practices.

I. Create programs to advance Recommended Actions F – H.

Funding bodies should support programs and provide financial, operational, and logistical support for implementing the respective recommendations **F-H** for ***scientists*** and ***policy practitioners***.

Acknowledgments: This Workshop was made possible by generous support from the Center for Epidemic Preparedness Innovations (CEPI), Rosetta Commons, Open Molecular Software Foundation (OMSF), and the Rockefeller Foundation.

Workshop Contributors: This list includes individuals who served as presenters and/or moderators at the Responsible Biodesign Workshop. Although these experts, as well as other Workshop participants, heavily informed this report's contents, the synthesis of information by the authors and the recommendations do not reflect a consensus or endorsement of this group or their institutions.

Polina Brangel, Coalition for Epidemic Preparedness Innovations
Erika DeBenedictis, Pioneer Labs / Align to Innovate
James Diggans, Twist Bioscience / International Gene Synthesis Consortium
Ian Haydon, Institute for Protein Design, University of Washington
Neil King, Institute for Protein Design, University of Washington
Sheng Lin-Gibson, US National Institute of Standards and Technology
Jassi Pannu, Stanford University / Johns Hopkins University
Gopal Sarma, White House Office of Science and Technology Policy (*formerly*)
Lynda Stuart, Institute for Protein Design, University of Washington *(formerly)*
Matthew Walsh, Johns Hopkins University
Jaime Yassif, NTI|bio

Authorship Disclaimer: The content expressed herein does not necessarily reflect the consensus positions of, or endorsement by, the authors or their respective institutions.

References

[1] Community Values, Guiding Principles, and Commitments for the Responsible Development of AI for Protein Design. Responsible AI X Biodesign n.d. https://responsiblebiodesign.ai/ (accessed May 23, 2025).
[2] Wittmann BJ, Alexanian T, Bartling C, Beal J, Clore A, Diggans J, et al. Toward AI-resilient screening of nucleic acid synthesis orders: Process, results, and recommendations. bioRxiv 2024:2024.12.02.626439. https://doi.org/10.1101/2024.12.02.626439.