

Testimony of

Jassi Pannu, MD

**Senior Scholar, Johns Hopkins Center for Health Security
Assistant Professor, Johns Hopkins Bloomberg School of Public Health**

Before the

**US House of Representatives Committee on Energy and Commerce
Subcommittee on Oversight and Investigations**

“Examining Biosecurity at the Intersection of AI and Biology”

December 17, 2025

Introduction

Chairman Joyce, Ranking Member Clarke, Chairman Guthrie, Ranking Member Pallone, and distinguished members of the Subcommittee, thank you for the opportunity to appear before you today to discuss biosecurity issues at the intersection of artificial intelligence (AI) and biology. The application of AI to solving challenging health and biomedical problems could be one of the most beneficial uses of this technology for American society and the world. In order to harness these benefits, we should seek to understand, anticipate, and prevent, the narrow set of biosafety and biosecurity risks that could significantly impact society. I commend the subcommittee for their attention to these issues.

I am currently an Assistant Professor at the Johns Hopkins Bloomberg School of Public Health, and a Senior Scholar at the Johns Hopkins Center for Health Security. The opinions expressed herein are my own and do not necessarily reflect the views of Johns Hopkins University.

I am a medical doctor by training, and completed my medical degree and internal medicine residency at Stanford University. During the peak of COVID-19 pandemic, I worked full time on the frontlines caring for patients, and I have worked as a doctor internationally in Uganda, a country that frequently faces both endemic and emerging infectious disease outbreaks. I also previously worked at Google AI, where I developed AI-enabled diagnostics. These experiences have shaped my view that technology advances can provide immense public health benefits; but also, that the prevention of biothreats with societal impact is essential. Prevention saves lives. It is also cost-effective and protects the American economy.

My work now focuses on tracking AI advancements for biology in order to develop technology and policy approaches for mitigating risks. I also investigate ways in which AI could improve our public health and biodefense readiness.

Today, I was asked to provide comments on how we can guard against potential harms of AI while at the same time working to ensure that AI will improve the nation's biosecurity and public health outcomes. Prior to offering these comments, I want to share my top-line recommendations as to the actions Congress should consider to address the biosafety and biosecurity risks of AI.

To that end, I recommend that Congress:

- 1) Provide CAISI with the authority to define the narrow subset of biological AI models that pose societal risks, to enable targeted risk assessment and mitigation;
- 2) Resource CAISI to closely track the benefits and risks of LLMs, AI agents, and autonomous robotics for biology, using benchmarks and evaluations;
- 3) Task the Office of Science and Technology Policy (OSTP) to formalize a thoughtful and targeted framework for biological data controls; and
- 4) Ensure that the currently ongoing NIH Biosafety Oversight modernization initiative addresses the use of AI in the design of biological agents and toxins;
- 5) Pass legislation to govern de novo gene synthesis.


These measures will protect the public from risks, engender trust in science, and enable society to fully harness the benefits of important research at the intersection of AI and biology.

The Current Landscape of AI for Biology in Brief

There are numerous different ways artificial intelligence is being used to accelerate and scale biological research.

Large Language Models and Reasoning Models

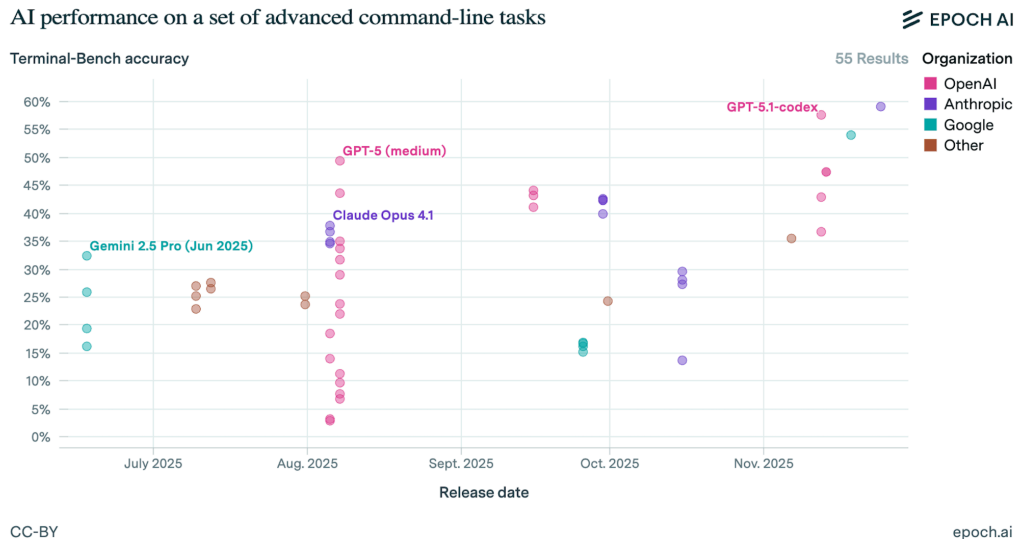
Most familiar to the public are large-language models (LLMs). LLMs currently surpass any individual human in terms of the breadth of their biology knowledge, and researchers can query LLMs to quickly access biology related information and expedite their work. Benchmarks such as GPQA and GPQA Diamond (1,2) that test LLMs on their ability to answer extremely difficult, Google-Proof questions show the latest LLMs exceed human expert performance in academic subjects including biology. A benchmark that specifically assesses virology knowledge (the Virology Capabilities Test), including fundamental, tacit, and visual knowledge, shows that some LLMs now outperform expert virologists (3).



To further enable AI-assisted biomedical research, models will need to move beyond simple recall and towards complex reasoning. Specialized forms of LLMs called reasoning models are built to solve complex, multi-step reasoning problems. A benchmark known as Humanity's Last Exam (4) that tests LLMs on expert-written academic questions across 100+ disciplines including biology, and that requires multi-step reasoning to answer correctly, shows that as of today reasoning models do not yet outperform expert humans. However, results do demonstrate that reasoning models over the last 2 years have significantly improved. Some benchmarks, such as LAB-Bench, seek to measure how LLMs can assist with practical biology research tasks, such as interpretation of figures, navigation of databases, and comprehension of DNA and protein sequences (5). Performance on this benchmark has been steadily increasing due to improvements such as increasing context window size, enabling models to ingest entire corpora of research papers.

Coding Capabilities

LLMs are increasingly proficient in writing code, including code for computational biology applications, as measured by benchmarks like Terminal-Bench (6). BioLLMBench is a bioinformatics-focused benchmark that assesses bioinformatics knowledge and bioinformatics coding capabilities (7); frontier models are saturating the knowledge components of this benchmark but are still challenged by the agentic and synthesis components.





AI Agents

Researchers now seek to scale biomedical research through the use of LLM-based agents to complete real-world tasks. AI agents can act as intelligent collaborators – planning workflows, processing data, and planning experiments. BixBench, a comprehensive benchmark for AI agents in computational biology, demonstrates that current frontier models still need significant development to achieve the goal of a completely autonomous bioinformatician “colleague” (8). On the other hand, another benchmark, ABC-Bench, found that top-performing LLMs could outperform the majority of human experts on tasks such as managing robotic lab equipment and designing DNA fragments (9). The team behind ABC-Bench found that model-generated code was executable in a real laboratory setting.

Biological AI Models

Separately from LLMs, there is a broad landscape of biological AI models – models trained on biological sequence, structure, function or imaging data (as compared to human-created text or images). Combining AI approaches with large biological datasets could enable the understanding and design of functional biological constructs, such as protein complexes. They could in the future enable a detailed understanding of cellular states, or the simulation of cellular responses entirely *in silico*. There are many categories of biological AI models, including: genomic language models, transcriptomics models, protein structure prediction models, protein language models, cell-state models, perturbation models, etc. The resources being invested in computation and data generation for biological AI models are increasing.



One increasing application of biological AI models is to generate biological designs that are constructed in the wet-lab. Because biological AI models currently have varied performance, constructing these designs in the wet-lab is often done in order to test the performance of the AI model.

Autonomous Robotics and Laboratory Equipment

To fully realize autonomous, AI-enabled biological research, the above described AI models and tools will need to be integrated into end-to-end research pipelines that leverage robotic laboratory equipment. Lab equipment is increasingly robotic and controlled by computer code in order to further scale research and reduce the expertise and human labor required at each step. Google DeepMind recently announced a partnership with the UK government that will establish a first-of-its-kind automated laboratory (10). Ginkgo Bioworks this month announced a \$47M contract from the USG to build a fully autonomous microbial laboratory at Pacific Northwest National Labs (11).

Potential Benefits for Biosecurity and Biodefense

Within the biological sciences, AI systems will likely provide immense benefit. They are likely to be employed to improve the discovery, diagnosis, and treatment of diseases, to boost agricultural yields, and to optimize the biosynthesis of useful products, among many other uses now being explored. AI also holds great promise in areas of biosecurity and biodefense. However, these areas are underincentivized by the private market, and companies may not seek to apply AI for the improvement of biosecurity and biodefense by default. The government should consider investing in defense-forward applications of AI in order to strengthen our resilience to biological threats.

Biosurveillance and Diagnostics

AI should be employed to improve our ability to detect outbreaks early, by advancing pathogen biosurveillance and scalable diagnostics. For example, AI could be used to synthesize multiple data streams, such as symptom search queries, wastewater sequence data, and vital signs from wearable devices.

Data Generation

A current barrier to biological AI model capabilities is high quality data. A growing proportion of wet-lab work for data generation, such as data visualization, data analysis, and sample creation, can be conducted by autonomous machines, including machines that researchers pay to access remotely, known as cloud labs; this work could be further scaled using AI.

Medical Countermeasure Development, Production and Distribution

Within medical countermeasure production and distribution, AI can potentially: (1) increase supply chain optimization through demand forecasting, inventory management, production capacity planning, and distribution network optimization; (2) enhance strategic national stockpile management through real-time tracking, predictive maintenance, and automated procurement; and (3) improve manufacturing coordination through production line scheduling, cross-manufacturer standardization, and quality assurance.

Materials and Hardware Development and Use

Current bioterror readiness is limited by the types of available personal protective equipment and built environment pathogen disinfection technologies. AI could be used to advance the materials science behind filters for reusable respirators, and to help users correctly fit test respirators themselves.

Strengthening Health Systems Resilience

Within healthcare personnel and resource deployment, AI can strengthen healthcare workforce management through: (1) personnel availability tracking, skills-based deployment matching, training coordination, and burnout prediction; (2) optimizing healthcare facility coordination through bed capacity management, equipment allocation, and surge capacity planning; and (3) enhance emergency medical services coordination through ambulance dispatch optimization, emergency department capacity management, and inter-facility patient transport.

Response Coordination & Treaty Monitoring

AI-agents could be leveraged for rapid outbreak response coordination, and LLMs could be used to disseminate reliable public health information. AI tools for open source intelligence could be leveraged to, for example, ensure compliance with treaties such as the Biological Weapons Convention.

Emerging Biosafety and Biosecurity Risks

Taken together, trends in AI progress across multiple domains suggest that AI model capabilities may play an increasing role in enabling high-consequence biosafety and biosecurity risks in the coming years.

Anticipating and mitigating these risks is necessary in order to secure these tools against both accident and misuse, and avoid impeding immense potential benefits.

LLM and Agent-Related Risks

As described above, LLMs and reasoning models can provide accurate biology information; a subset of this information could potentially enable a bad actor (e.g. a terrorist group) to design and conduct a bioweapons attack. In the field this is referred to as AI “uplift” (12). In the biosecurity context, AI uplift increases the number of users capable of carrying out a bioweapons attack. Frontier AI companies conduct safety evaluations to determine how much uplift their models provide, and implement mitigations such as prompt refusals in order to prevent users from querying models for bioweapons related knowledge. Frontier AI companies can voluntarily report the results of their safety evaluations to the US Center for AI Standards and Innovation (CAISI), however there is no federal legal requirement to do so. Importantly, bioweapons-related knowledge does not encompass all of biology. The latest LLMs answer many biology-related questions, but refuse queries related to biological weapons ideation and design, methods for obtaining physical materials (such as synthetic DNA), protocols validating bioweapons candidates in the wet lab, and instructions for the dispersal of a bioweapon. Similarly, AI agents and the broader set of autonomous scientific workflows described previously, may contribute further bioweapons uplift in the absence of safeguards.

Biological AI Model Risks

Biological AI models pose a different set of risks; consideration of these risks is relatively neglected compared to LLM-related risks. Biological AI models are often research tools, intended to be used by expert scientists to tackle specific biology-related problems, though advancements in AI agents and AI-enabled coding are also making these biology models more accessible to non-experts.

There is no comprehensive shared understanding of what kinds of capabilities pose consequential biosecurity risks – though our team and other groups have made initial progress on this question (13,14) – nor what methods should be used to evaluate them, or the risk mitigation measures that should be undertaken.

Current biological AI models possess some capabilities that might be misused: they can be used to design novel viral capsids, forecast pathogen evolution, craft nucleic acid sequences to evade safety screening software, and create novel bacteriophage genomes with augmented fitness when synthesized in the wet lab, per claims made by their developers and outside researchers (15–19). Developers are releasing new, more proficient biological models, often without conducting basic safety assessments, a practice that would be unacceptable in other domains of life-science research (20). Many biological AI models are general-purpose, meaning the same biological models able to design a benign viral vector to deliver gene therapy could be used to design a more pathogenic virus capable of evading vaccine-induced immunity. This makes it increasingly important to educate developers on potential risks, and to develop objective risk assessment methods.

While current models have significant performance limitations, both academic and industrial incentives push researchers to develop future models with reliable scientific abilities; as described above, these may be integrated with AI agents or wet-lab experimentation. As model capabilities become more powerful, systems made from combinations of these tools may enable more concerning pathogen-related applications. These include the design of novel pandemic pathogens, such as a novel variant of pandemic influenza, or viruses engineered to evade vaccine-induced immunity. Because of the immense harm models possessing these capabilities could have if misused, biological AI models are often referred to as “raising the ceiling of harm”.

Oversight and Risk Mitigation Approaches

Proactive oversight and risk mitigation is needed given the speed at which AI is advancing; such oversight would protect the public and engender trust in science. Both the frontier AI and biological AI communities have taken steps towards this proactive oversight. In 2023, several Frontier AI companies committed voluntarily to conducting safety evaluations for biosecurity risks (21). In 2024, a global consortium of biological AI model developers signed the Responsible AI x Biodesign statement of community values and commitments, which also included safety evaluations and risk mitigation (22).

Define Biological AI Model Capabilities of Concern

A National Academies Workshop held to discuss the risks of in silico and artificial intelligence approaches to biological research (23) highlighted that an authoritative working group should be established to determine a tiered oversight approach, where high risk research would undergo further risk assessment while allowing lower risk research to proceed and be shared openly. Recommendations also included developing a “Capability of Concern” list, akin to the Dual-Use Research of Concern categories, to help researchers evaluate the risks of their work. Establishing an understanding of these capabilities would improve upon the current ad hoc approach. Congress should give CAISI the authority to define capabilities of concern for the narrow subset of biological AI models that pose societal risks, and should specify in its tasking sufficient resources and staffing to re-assess this taxonomy at least every 6 months or more frequently as the rate of technological progress warrants.

Identifying which biological AI capabilities pose the greatest biosecurity and biosafety concerns is necessary in order to establish targeted oversight and avoid impeding low-risk research. Within the broad range of AI models being developed, only a narrow set of advanced biological AI models have characteristics that currently warrant oversight. By focusing on these classes of models, officials are more likely to target governance on only those models that pose the greatest risks without unduly hampering innovation and the benefits of AI. Such governance also enhances innovation because clear safety standards provide researchers with the confidence to pursue high-impact applications without fearing unintentional harms.

Develop Biological AI Model Evaluation and Risk Mitigation Methods

It is encouraging to see CAISI partner with frontier AI labs to test LLMs for uplift capabilities related to biothreats. Autonomous science systems and biological AI models have not received the same treatment. Congress should empower CAISI with the authority and resources to develop methods to evaluate biological AI models for capabilities of concern, including prior to development and prior to release. Results of these safety evaluations should be linked to risk mitigation standards. For example, generalizable AI models for functional virus design, if created in the future, may warrant standards and requirements such as not disclosing model code or model weights, in order to prevent accidents and misuse.

Modernize Biosafety Oversight for AI-Enabled Pathogen and Toxin Research

The NIH is currently conducting an important initiative to modernize and strengthen biosafety oversight (24). Biosafety oversight has long played an important role in reducing risks to individual laboratory staff when handling infectious agents and toxins. However, consideration of the risks infectious agents pose on a societal scale have been insufficiently emphasized when defining risk groups and biosafety levels. This has resulted in pathogens capable of causing pandemics, widespread loss of life and societal disruption, as being tiered as lower risk groups or biosafety levels.

Future biosafety oversight should aim to harmonize the NIH's recombinant/synthetic policies with the CDC's Biosafety in Microbiological and Biomedical Laboratories (BMBL) best practices, to create a single, unified, and easy-to-implement policy for researchers. This policy should ensure that known and potential pandemic pathogens are subject to stronger biosafety oversight, given their associated societal risks. Any harmonized policy should also move away from mere best practices, towards explicit requirements to receive federal research funding.

As described above, AI and new synthetic biology approaches will increasingly enable the design and engineering of agents and toxins. Researchers recently synthesized viable, complete AI-generated genomes of bacteriophages in the wet-lab (19), demonstrating an important proof-of-concept for how AI will be used for pathogen design. It is therefore necessary and prudent for us to consider how biosafety oversight handles pathogens and toxins that are designed using AI.

The current NIH Guidelines assume a researcher intentionally chooses the genetic alterations they intend to make. With this information, one can determine the Risk Group and Experimental Category of the work, and therefore make a biosafety assessment. Generative AI (and some other computational approaches) present novel challenges. Rather than human intention, an AI model is used to generate novel designs, and may draw on any information present in its training data. The NIH must be clear at what stage these AI-generated designs are reviewed by a human biosafety expert, and how determinations of the appropriate biosafety procedures are made. Congress should direct the NIH to include these considerations as part of their ongoing policy modernization effort (24).

Established Targeted Biological Data Controls

Progress in AI for biomedicine will require large biological datasets. The National Security Commission on Emerging Biotech (NSCEB) describes biological data as a strategic resource that should be invested in and standardized to ensure its ready use in AI models (25). The recently announced Genesis Mission, described in EO 14363 (26), will coordinate a national effort to integrate Federal scientific datasets to enable AI, including AI for biology. The previous administration's Frontiers in Artificial Intelligence for Science, Security, and Technology (FASST) initiative also supported the creation of AI-ready datasets. Private

research organizations are also heavily investing in biological data generation. These efforts are very important to successfully advance AI for biology.

Biological data is a strategic resource, and a narrow set of biological data will require security to prevent misuse. The NSCEB, as well as over 100 researchers at the 50th anniversary Asilomar conference held earlier this year (27), have endorsed targeted security controls on biological data that is particularly susceptible to enabling AI misuse.

Data-access limitations are a policy approach already widely accepted in the domain of privacy protection. Scientists broadly accept data-access limitations to protect privacy. Government programs in the United States, Europe, and elsewhere mediate access to personally identifiable and sensitive data—such as genetic data, health records, and financial reports—to permit rigorous research without compromising privacy. Though imperfect, these frameworks have shown that responsible data governance and scientific progress are not contradictions.

The NSCEB’s recommendation that Congress establish a central office within the Department of Energy (DOE) to manage and store biological data, would both spur innovation and ensure data security. Congress is currently considering legislation that would, following the Commission’s recommendation, establish a “Web of Biological Data,” including tiers of access controls to prevent misuse, and I urge Congress to advance it.

The types of data that warrant security and access limitations are likely a narrow band of novel and sensitive pathogen data—such as comprehensive genotype-to-phenotype mappings for pandemic pathogens. Data access can be carefully scoped, allowing governments to impose access restrictions on a small set of data, successfully reducing risk, without affecting the vast majority of biological data. Without such controls, data can be anonymously and irrevocably accessed by bad actors and used to train AI models.

To make controlled data broadly available to responsible researchers, Congress should pair controls with support for the standardization and development of trusted research environments (TREs) for use with the highest tiers of control. My team has enumerated several promising cybersecurity controls that could be implemented via TREs (28), work recently presented at the leading AI conference *NeurIPS*. The DOE’s Scalable Protected Data cloud research environment, which meets National Institute of Standards and Technology (NIST) control standards, could serve as one template for such TREs. Whether or not the US government runs the platforms themselves, TREs should be federally supported and required to meet stringent security standards.

Congress should task the Office of Science and Technology Policy (OSTP) to formalize and standardize a thoughtful framework for biological data access limitations and security controls. Some prominent biological model developers have voluntarily removed data related to viruses and toxins from their training data, but no government-backed expert panel has given guidance on specific data of concern. My colleagues and I have created a preliminary framework consisting of five Biosecurity Data Levels (BDLs), ranging from BDL-0 to BDL-4 (28). Each level governs data that present progressively greater risk compared to reward, and imposes layered controls based on past approaches to guard sensitive biological data. In designing the BDL tiers, we have drawn on the work of expert panels and government policies identifying experimental outcomes that pose the greatest biosecurity threats, such as those that increase pathogen transmissibility, virulence, stability, host range, and immune evasion.

Governance of *de novo* Gene Synthesis

Congress should also govern the part of the supply chain that would potentially enable a bad actor to acquire the genetic materials necessary to build a bioweapon, or for a researcher to conduct risky research that results in an accidental outbreak. This part of the supply chain is often referred to as the digital-to-physical barrier.

Gene synthesis is the process of constructing physical genetic material from digital genetic sequence information. This technology is not new; the first synthesized gene was created in 1972. However, advancements since then have allowed companies (known as gene synthesis providers) to routinely create these physical genetic material at scale, contributing to an affordable and efficient gene synthesis industry that supports life sciences research and development broadly. These benefits mean targeted governance mechanisms that do not impede the benefits of this technology are needed.

Only a small subset of genetic sequences contribute to misuse risk. To create targeted safeguards, researchers and industry have coalesced around a set of protocols known as gene synthesis screening (GSS). GSS involves implementing a know-your-customer (KYC) protocol for the customers ordering genetic material to ensure they are using it for legitimate research purposes, and screening for sequences of concern (genetic sequence known to encode pathogens or toxins). This screening limits access to potentially dangerous genetic material to legitimate researchers.

Many biotechnology companies that are gene synthesis providers already implement their own screening standards, with more than 30 leading providers (mostly US companies) committing to voluntary nucleic acid synthesis screening (29). While these industry standards are a positive step, voluntary standards are not a replacement for federal guidance. Federal rules can also level the playing field, ensuring that companies that conduct screening and are responsible actors who invest in safety, are not undercut by those that do not.

The research community is also supportive of universal GSS. Last year, over 150 signatories of the Responsible AI x Biodesign statement, committed to only purchase synthesized genes from providers that screen their orders and customers (22).

Both the previous and current administrations have taken positive steps to put forward safeguards and guidance for the community. The previous administration released the “The US Framework for Nucleic Acid Synthesis Screening” (the Framework), that established procedures for all federally funded entities to purchase genetic material only from providers that comply with the Framework (30). My colleagues at the Center for Health Security created a tracker for providers to self-attest compliance with the Framework (31). By establishing clear federal requirements tied to US funding, a framework of this kind would create a competitive advantage for compliant providers and pressure international competitors to meet these higher biosecurity standards to access the lucrative US market. It would also demonstrate US leadership in biosecurity and help export American safety standards globally, as many US companies are already voluntarily participating and setting the global standard for responsible nucleic acid synthesis screening.

Executive Order (EO) 14292 is intended to replace the Framework, and directs OSTP to release new guidance and regulations as well as an enforcement mechanism. EO 14292 also directs the creation of a legislative proposal for universal GSS, in order to require non-federally funded entities to conduct screening as well (32). The new enforcement mechanism has not yet been released, leaving industry without regulatory clarity. This subcommittee can call for the release of the new framework, which is delayed well past the 90-day timeline set in the EO.

Congress should also act on EO 14292's request for a legislative proposal for universal GSS that requires providers of synthetic genes and manufacturers of synthetic gene equipment to screen orders and customers for potential threats by drafting legislation tasking a single agency, such as the Department of Commerce or the Department of Health and Human Services, with this oversight and governance work.

Conclusion

The convergence of artificial intelligence and biology represents one of the most consequential scientific developments of our time. The oversight approach outlined above is intended to be proportionate, targeted, and adaptive as capabilities evolve. Congress should act now, while we still have the opportunity to shape this rapidly advancing field's trajectory thoughtfully and deliberately. I thank the Subcommittee for its attention to these critical issues and stand ready to support your efforts to advance sensible biosecurity policy.

References

1. [2311.12022] GPQA: A Graduate-Level Google-Proof Q&A Benchmark [Internet]. [cited 2025 Dec 14]. Available from: <https://arxiv.org/abs/2311.12022>
2. Epoch AI [Internet]. [cited 2025 Dec 14]. GPQA Diamond. Available from: <https://epoch.ai/benchmarks/gpqa-diamond/>
3. Virology Capabilities Test [Internet]. [cited 2025 Dec 14]. Available from: <https://www.virologytest.ai/>
4. Phan L, Gatti A, Han Z, Li N, Hu J, Zhang H, et al. Humanity's Last Exam [Internet]. arXiv; 2025 [cited 2025 Dec 14]. Available from: <http://arxiv.org/abs/2501.14249>
5. Laurent JM, Janizek JD, Ruzo M, Hinks MM, Hammerling MJ, Narayanan S, et al. LAB-Bench: Measuring Capabilities of Language Models for Biology Research [Internet]. arXiv; 2024 [cited 2025 Dec 14]. Available from: <http://arxiv.org/abs/2407.10362>
6. Terminal-Bench [Internet]. [cited 2025 Dec 14]. Terminal-Bench. Available from: <https://www.tbench.ai>
7. Sarwal V, Munteanu V, Suhodolschi T, Ciorba D, Eskin E, Wang W, et al. BioLLMBench: A Comprehensive Benchmarking of Large Language Models in Bioinformatics [Internet]. bioRxiv; 2023 [cited 2025 Dec 14]. p. 2023.12.19.572483. Available from: <https://www.biorxiv.org/content/10.1101/2023.12.19.572483v1>
8. Mitchener L, Laurent JM, Andonian A, Tenmann B, Narayanan S, Wellawatte GP, et al. BixBench: a Comprehensive Benchmark for LLM-based Agents in Computational Biology [Internet]. arXiv; 2025 [cited 2025 Dec 14]. Available from: <http://arxiv.org/abs/2503.00096>
9. Liu AB, Nedungadi S, Cai B, Kleinman A, Bhasin H, Donoughe S. ABC-Bench: An Agentic Bio-Capabilities Benchmark for Biosecurity. In 2025 [cited 2025 Dec 14]. Available from: [https://openreview.net/forum?id=mo5H9VAr6r&referrer=%5Bthe%20profile%20of%20Samira%20Nedungadi%5D\(%2Fprofile%3Fid%3D~Samira_Nedungadi1\)](https://openreview.net/forum?id=mo5H9VAr6r&referrer=%5Bthe%20profile%20of%20Samira%20Nedungadi%5D(%2Fprofile%3Fid%3D~Samira_Nedungadi1))
10. Google DeepMind [Internet]. 2025 [cited 2025 Dec 14]. Our partnership with the UK government. Available from: <https://deepmind.google/blog/strengthening-our-partnership-with-the-uk-government-to-support-prosperity-and-security-in-the-ai-era/>

11. Energy.gov [Internet]. 2025 [cited 2025 Dec 14]. Energy Department Launches Breakthrough AI-Driven Biotechnology Platform at PNNL. Available from: <https://www.energy.gov/articles/energy-department-launches-breakthrough-ai-driven-biotechnology-platform-pnnl>
12. Dual-Use AI Capabilities and the Risk of Bioterrorism | GovAI [Internet]. [cited 2025 Dec 15]. Available from: <https://www.governance.ai/research-paper/dual-use-ai-capabilities-and-the-risk-of-bioterrorism-converting-capability-evaluations-to-risk-assessments>
13. Pannu J, Bloomfield D, MacKnight R, Hanke MS, Zhu A, Gomes G, et al. Dual-use capabilities of concern of biological AI models. *PLoS Comput Biol*. 2025 May 8;21(5):e1012975.
14. Pannu J, Gebauer SL, Bradley HA, Woods D, Bloomfield D, Berke A, et al. Defining Hazardous Capabilities of Biological AI Models: Expert Convening to Inform Future Risk Assessment [Internet]. 2025 Aug [cited 2025 Dec 14]. Available from: https://www.rand.org/pubs/conf_proceedings/CFA3649-1.html
15. Hayes T, Rao R, Akin H, Sofroniew NJ, Oktay D, Lin Z, et al. Simulating 500 million years of evolution with a language model. *Science*. 2025 Jan 16;0(0):eads0018.
16. Thadani NN, Gurev S, Notin P, Youssef N, Rollins NJ, Ritter D, et al. Learning from prepandemic data to forecast viral escape. *Nature*. 2023 Oct;622(7984):818–25.
17. Nisanov AM, Rivera de Jesús JA, Schaffer DV. Advances in AAV capsid engineering: Integrating rational design, directed evolution and machine learning. *Mol Ther J Am Soc Gene Ther*. 2025 May 7;33(5):1937–45.
18. Wittmann BJ, Alexanian T, Bartling C, Beal J, Clore A, Diggans J, et al. Strengthening nucleic acid biosecurity screening against generative protein design tools. *Science*. 2025 Oct 2;390(6768):82–7.
19. King SH, Driscoll CL, Li DB, Guo D, Merchant AT, Brixi G, et al. Generative design of novel bacteriophages with genome language models [Internet]. bioRxiv; 2025 [cited 2025 Dec 14]. p. 2025.09.12.675911. Available from: <https://www.biorxiv.org/content/10.1101/2025.09.12.675911v1>
20. Villalobos P. Epoch AI. 2025 [cited 2025 Mar 9]. Announcing our Expanded Biology AI Coverage. Available from: <https://epoch.ai/blog/announcing-expanded-biology-ai-coverage>
21. House TW. FACT SHEET: Biden-Harris Administration Secures Voluntary Commitments from Eight Additional Artificial Intelligence Companies to Manage the Risks Posed by AI [Internet]. The White House. 2023 [cited 2025 Dec 14]. Available from: <https://bidenwhitehouse.archives.gov/briefing-room/statements-releases/2023/09/12/fact-sheet-biden-harris-administration-secures-voluntary-commitments-from-eight-additional-artificial-intelligence-companies-to-manage-the-risks-posed-by-ai/>
22. Responsible AI x Biodesign [Internet]. [cited 2024 Mar 27]. Responsible AI x Biodesign. Available from: <https://responsiblebiodesign.ai/>
23. Navigating the Benefits and Risks of Publishing Studies of In Silico Modeling and Computational Approaches of Biological Agents and Organisms – A Workshop [Internet]. [cited 2025 Dec 14]. Available from: <https://www.nationalacademies.org/projects/DELS-BLS-24-06/event/44297>
24. NIH Launches Initiative to Modernize and Strengthen Biosafety Oversight | National Institutes of Health (NIH) [Internet]. [cited 2025 Dec 15]. Available from: <https://www.nih.gov/about-nih/nih-director/statements/nih-launches-initiative-modernize-strengthen-biosafety-oversight>
25. Treat Biological Data as a Strategic Resource [Internet]. Biotech. [cited 2025 Dec 15]. Available from: <https://www.biotech.senate.gov/final-report/chapters/chapter-4/section-1/>

26. Orders E. The White House. 2025 [cited 2025 Dec 15]. Launching the Genesis Mission. Available from: <https://www.whitehouse.gov/presidential-actions/2025/11/launching-the-genesis-mission/>
27. Bromberg Y, Altman R, Imperiale M, Horvitz E, Dus M, Townshend R, et al. 3.1 Artificial Intelligence and the Future of Biotechnology. 2025 [cited 2025 Dec 15]; Available from: <https://hdl.handle.net/1911/118555>
28. Bloomfield D, Hanke MS, Maiwald A, Black JRM, Webster T, Hernandez-Boussard T, et al. Securing Dual-Use Pathogen Data of Concern. In 2025 [cited 2025 Dec 15]. Available from: <https://openreview.net/forum?id=ZgD951FVe7¬eId=ZgD951FVe7>
29. Home | International Gene Synthesis Consortium [Internet]. International Gene Synthesis Consortium | The Promotion of Biosecurity. 2017 [cited 2025 Dec 15]. Available from: <https://genesynthesisconsortium.org/>
30. Framework for Nucleic Acid Synthesis Screening | OSTP [Internet]. The White House. 2024 [cited 2025 Dec 15]. Available from: <https://bidenwhitehouse.archives.gov/ostp/news-updates/2024/04/29/framework-for-nucleic-acid-synthesis-screening/>
31. Gene Synthesis Screening Information Hub [Internet]. [cited 2025 Dec 15]. Available from: <https://genesynthesiscreening.centerforhealthsecurity.org/>
32. Orders E. The White House. 2025 [cited 2025 Dec 15]. Improving the Safety and Security of Biological Research. Available from: <https://www.whitehouse.gov/presidential-actions/2025/05/improving-the-safety-and-security-of-biological-research/>