Present and future of safe biological AI

By Brian Hie · August 15, 2026

Given the public conversation around the generative phage design paper and in response to several questions I have been receiving, I wanted to provide my personal thoughts on future biosafety and biosecurity implications, not just of the phage work, but of generative genomics and generative biology more broadly. We as researchers do need to be extremely careful and proactive, but I am optimistic that we can steer this technology to maximize human good.

First, our specific work in designing phages (viruses that infect bacteria) is aimed ultimately at preventing the spread of dangerous and evolving bacterial pathogens. AI-guided phage engineering is an active area of research pursued by many academic labs and several companies. This line of work is extremely important given the rise of bacterial infections that are resistant to most or all of our antibiotics. The experiments we performed in the paper were also safe and controlled, and while the generated phages had biologically interesting new sequences, structures, and functions (including greater resilience against resistant bacteria), all of the generated phages remain close in sequence space to the design template, wildtype ΦX174 (we did not attempt the design of completely *de novo* phages). Our paper further lays out detailed thoughts on the biosafety and biosecurity implications of our phage design work in a supplementary discussion. Moreover, obtaining these phages also required substantial time, resources, and a team of multidisciplinary scientific experts spanning machine learning, bioinformatics, and microbiology. Given what we know from this paper, I would assess the current threat of transferring these techniques from phages to the synthesis of pathogenic, highly novel, AI-generated viruses that infect humans to be very low.

However, the work has generated substantial interest in the press and among the public, not for the current state of capabilities, but for what future capabilities may lie ahead. Our paper represents a real technological milestone in producing the first AI-generated genomes, and genome-scale design unlocks new functions beyond the design of individual genes or molecules. AI for protein and molecular design will continue to revolutionize medicine, sustainability, and many other fields. At the same time, AI-generated biomolecules or AI-guided mutations could aid in causing significant harm to humans. As generative biology advances to the genome-scale and beyond, it is a near certainty that generative models will achieve much more sophisticated applications with the potential for both positive and negative consequences.

As someone on the frontlines of biological AI development, my view on safety is therefore informed by both current capabilities and the trajectory of progress. While there is no imminent danger from AI-

generated human viruses, I also agree that we must undergo urgent preparations for a future in which (alongside any existing or future computational safeguards) we will need to rapidly identify and defend against threats with pandemic potential generated by nature, human researchers, or an AI system. Doing so will include improved DNA synthesis screening, sequencing-based pathogen surveillance, intelligent containment strategies, broad-spectrum interventions, and targeted interventions, where AI tools will play a major role across all of these.

To ensure a beneficial future for biological AI, one of the most useful things I can do as a technology developer is to build advanced new tools that can fight and prevent threats with pandemic potential. A core aim of my research lab from the beginning has been to develop safe biological AI for the good of humanity. A substantial portion of my lab's ongoing efforts have immediate applications in improving biosafety and biosecurity, including research on models for better responding to both current and future forms of pathogen evolution. This includes better methods for predicting how viral diseases can change and designing therapeutics or vaccines that respond to these changes, for which we have new research that I am excited to share soon.

As a professional scientist, I do not claim to know the best forms of governance and policy for AI for biology, although there are a few areas in which I do have stronger convictions. First, the current biosafety level (BSL) system, a globally recognized framework for safety and containment practices, does provide a useful existing foundation for future work on AI-generated biology. The BSL system already covers techniques like introducing random mutations or environmental sampling that could yield unpredictable functions, and this system could be adapted to consider both the safety and societal consequences of AI-generated sequence changes as well. Moreover, AI models that predict function from sequence could form the basis of rigorous, AI-assisted biosafety assessments that improve the safety profiles of AI-designed sequences compared to previous experimental design methods. Second, I hope the majority of models can continue to be released openly to advance broadly useful and reproducible science. AI models will play a critical role in flagging dangerous new pathogens and in developing new therapeutics and vaccines, and the scientific community should have wide access to these tools. Conversations around limited or closed models in biology for biosecurity reasons should involve many different perspectives, including the model developers themselves. Third, it is essential that regulatory frameworks support rather than inadvertently hinder safety-oriented researchers. A highly restricted ecosystem risks disadvantaging academic labs working on defensive tools, while potentially leaving non-compliant actors unaffected by those same limitations.

Finally, I want to end on a note of optimism. Unfortunately, it is already relatively easy to design dangerous biological systems without any use of AI (and nature continuously produces new diseases without any human engineering). What humanity cannot do is easily identify and respond to new pandemic-level threats, or to solve its most intractable diseases. I thus continue to maintain that bet-

ter biological AI should have much greater potential benefits for humanity than risks, and it is why I have personally dedicated myself to the mission of realizing these benefits. We as humans are almost completely helpless in the face of drug-resistant pathogens or metastatic cancers or age-related degeneration. If developed thoughtfully, correctly, and urgently, AI tools for biology will be indispensable to resolving our longstanding fight against human suffering and disease.

I would like to thank Trevor Bedford, Dave Burke, Patrick Collison, Hani Goodarzi, Patrick Hsu, Julia Kazaks, Samuel King, Silvana Konermann, David Li, Usman Muzaffar, Jeremy Ratcliff, and Alden Woodrow for feedback on this document, though the views here are my own.