

EDWARD H. YOU

Biosecurity for the Age of AI

There is much discussion on how artificial intelligence poses threats to biosecurity; there should be at least as much focus on how AI can expand the solutions.

Today's biosecurity approach is still grounded in a foundation laid nearly two decades ago. In 2008, when I reported to Federal Bureau of Investigation Headquarters as a newly minted supervisory special agent to support biosecurity policy, one important, emerging question was how providers of synthetic DNA should screen orders from customers. As commercial gene synthesis became increasingly accessible, many people worried that individuals could order genetic sequences coding for toxins and pathogens and then use them to cause harm.

I was fresh from the field, where I'd been conducting counterterrorism investigations and serving on the Hazardous Materials Response Team, a specialized unit trained to manage crime scenes and collect forensic evidence involving chemical, biological, radiological, and nuclear materials, plus weapons of mass destruction. My new mission was to inform an interagency process that went on to produce the 2010 federal guidance developed specifically for companies that manufacture and supply synthetic DNA. The goal was to preserve legitimate access for research, biotechnology, and other beneficial applications while mitigating risks. This voluntary framework established an approach that still stands today: screening both customers and sequence orders before DNA synthesis. I participated as the FBI law enforcement representative supporting the US Department of Health

and Human Services on the guidance and in related efforts, including work by the National Science Advisory Board for Biosecurity. We convened many roundtables on how to support legitimate research while preventing bioterror.

The questions we needed to answer were straightforward, if not always simple to answer: What are people ordering? And with what intent?

By pinpointing sequences of concern, we developed screens and lists for the first question, where our expertise was stronger. We did much less with the second, which involved understanding human behavior.

I've given a lot of thought to how, in hindsight, I would change my approach to this process. The participants who crafted the guidance brought deep expertise in biology, public health, policy, national security, and the emerging synthetic DNA industry. Knowing what I know now, I would add more seats to the table. I would invite people who understand logistics, customer and vendor relationships, enterprise procurement, fraud detection, payments, and how goods move through real markets—people from companies like Amazon, PayPal, FedEx, and Oracle, and institutions such as the Financial Crimes Enforcement Network, the bureau at the US Treasury Department that tracks financial crimes. They would know less about synthetic biology than the scientists, but they would understand how to establish trust across millions of

routine transactions, how to identify unusual behavior, and how to decide when an anomaly deserves a closer look.

That understanding is especially needed now. Since the process to place a synthetic DNA order increasingly moves through standardized digital systems, one would expect those systems to include automated analysis that detects suspicious activity. Yet the biosecurity community has done little to build this capacity. Meanwhile, generative artificial intelligence increasingly makes it possible to produce biological designs of concern that are not detected with the conventional approach, which issues alerts based on sequence alone.

The biosecurity community must adapt to answer the two essential questions—“What?” and “With what intent?”—in today’s technological context. That requires more sophisticated ways of assembling and cross-referencing information on both questions and improving the community’s ability to separate meaningful signals from ordinary noise.

Lists are a day late and a dollar short

The federal guidance released in 2010 succeeded in getting many DNA synthesis companies to screen their orders. Even today, requested sequences are compared against lists of viruses, bacteria, and toxins known to pose significant risk to human, animal, and plant health. By 2010, these lists had already been defined through the federal government’s select agents regulations, export control classifications, the Commerce Control List, and international arrangements, such as the Australia Group Common Control Lists. The logic was simple: Compare what DNA is ordered with known threats. In practice, this meant matching sequences ordered to sequences on lists. This is technically appealing, since sequences can be compared, scored, and flagged.

But lists are retrospective. New pathogens cannot be placed on lists until they are recognized. Viruses such as MERS-CoV, SARS-CoV, the highly pathogenic H5N1 avian influenza, and even SARS-CoV-2 were already affecting the world before regulatory processes placed them on lists. In other words, lists follow the biology and so can never be comprehensive. Generative AI compounds this problem, as it could produce off-list sequences that encode a potentially harmful structure, function, phenotype, or capability without closely resembling sequences that are already known.

AI also broadens what might constitute a threat. One important dimension of biosecurity appropriately worries about terrorists, rogue states, and deliberately engineered pandemic threats. However, as engineering biology becomes increasingly accessible as a design platform, criminal markets may attempt to manufacture illicit drugs, performance-enhancing compounds, counterfeit or unlicensed biologics, and similar products that could cause harm but are not covered by lists.

So although sequence screening is important, it will always provide inadequate and incomplete answers to biosecurity’s two essential questions.

Behavior can give a stronger signal than sequence

What’s more, the DNA molecule itself cannot carry out an attack. For a sequence to cause harm, someone must generate and distribute the encoded product, for example, by expressing it in cells in a laboratory with specialized equipment.

One strategy to assess ill intent is to look for signals beyond the sequence. The early DNA screening efforts already recognized that customer legitimacy mattered. A concerning order from an established research institution is different from the same order placed through a fictitious identity or delivered to an unrelated location.


The primary reason the biotechnology experts who assembled two decades ago neglected the question of human intent is that it is much harder to assess than a DNA order. Simply gaining possession of a sequence that codes for something dangerous says nothing about whether the recipient intends to do something dangerous. After all, legitimate researchers also need such sequences to detect disease and develop preventions and cures.

Having spent much of my career in weapons of mass destruction investigations, I am uncomfortable when conversations about sequence screening gloss over everything that comes after someone acquires a concerning pellet of DNA. Depending on the objective, a person may still need to synthesize or assemble material, characterize it, produce it at useful scale, stabilize or formulate it, and ultimately deliver or disseminate it. Each step presents a logistical hurdle. And each step can generate observable activity.

Orders for biological materials do not occur in a vacuum. Laboratories buy equipment, reagents, growth media, cell lines, consumables, and services. That generates purchase orders, grant charges and cost-center records, shipping information, delivery notices, laboratory information-management data, and all sorts of other administrative traces that can be examined. The further someone moves toward operational capability, the more opportunities there are for detection and intervention.

A single unusual purchase may mean nothing. So may an unexpected shipping address, an unfamiliar vendor, or a change in purchasing frequency. But when otherwise ambiguous indicators converge, the signal can become stronger. The pattern is not proof; it is a lead—and the biosecurity community needs a better way to follow these leads.

Financial institutions and law enforcement provide a useful model for working with fragments of information. Suspicious-activity reporting and fraud-detection systems do not assume that an unusual transaction proves criminal intent. Instead, they uncover patterns, relationships, and aberrant behavior that may warrant further review. Similarly, biosecurity needs to be able to use patterns of behavior, which can reveal more than a one-time transaction, such as a commercial DNA order.



Mapping the biological terrain

The biological content of an order still matters, and again AI has the potential to be part of the solution. In the same way that AI is making it possible to design biology to evade sequence-based comparison, it is also improving researchers' ability to navigate relationships among sequence, structure, molecular interaction, post-translational modification, and biological function. Much of the current AI-biosecurity debate focuses on what AI might enable someone to design, but the biosecurity community should be equally interested in what AI might enable us to detect.

I am not suggesting that today's AI can be used to look at a novel sequence, predict its structure, and definitively determine what it will do. Biology is far too complex for that. But AI can be used to do what I call *biological terrain mapping* to extract more meaning from a sequence than simply determining whether it matches a list.

One way to think about this is as a key and a lock. The key is the biological design or sequence. The lock is the biological target, structure, pathway, or system with which that design may interact to produce an effect. Traditional sequence screening asks whether a new key looks sufficiently similar

and may undergo post-translational modifications that alter its behavior. Understanding what it may actually do requires knowing expression levels, localization, molecular interactions, cellular context, host biology, dose, and system-level effects. And use context matters, too: Botulinum toxin, for example, is among the most poisonous biological substances known, yet it is widely used for medical and cosmetic purposes.

Recent work using genome language models to generate functional bacteriophages—viruses that infect bacteria—illustrates why a broader perspective matters. Generative AI proved able to create over a dozen highly diverse viruses that together could infect bacteria resistant to a natural bacteriophage. The AI-generated bacteriophages also had novel sequences that did not match that of known bacteriophages. Bacteriophages don't infect human cells, but they could affect human health by, for instance, affecting the microbiome. This work clearly demonstrates how insufficient it is now to ask simply whether an ordered DNA sequence matches a listed sequence. The more relevant security question becomes: What biological system might this design affect, and what could happen next?

Much of the current AI-biosecurity debate focuses on what AI might enable someone to design, but the biosecurity community should be equally interested in what AI might enable us to detect.

to keys that have already been identified as concerning. But generative AI raises the possibility of creating a very different-looking key that may nevertheless fit the same lock. Biological terrain mapping entails examining the lock—and the interaction between key and lock—rather than relying principally on how much one key resembles another.

The analogy is useful because biology is not unconstrained. Life uses an immense variety of proteins, but nature relies on a comparatively limited repertoire of recurring three-dimensional folding architectures, numbering in the low thousands by common classification schemes. The numbers of possible amino-acid sequences are astronomical. Nonetheless, physical constraints shape how protein chains fold into stable, functional machinery, including binding sites and molecular interfaces. The universe of possible keys may be enormous, but the biological terrain in which they must operate has structure and constraints. In other words, the locks are finite. Protein structure is an obvious place to start for biological terrain mapping, but it is not the endpoint. A protein does not emerge from translation as a static object whose function can always be read directly from its amino-acid sequence. It folds, interacts with other molecules, moves within a cell,

At the same time that generative AI is expanding the universe of potentially concerning biological designs, AI-enabled structural and systems biology is making it easier to understand the constraints those designs must operate within. Biological terrain mapping will not provide a definitive answer to whether a particular sequence is a threat anytime soon. But it represents a research direction worth pursuing.

Building a handshake

Even if researchers become much better at understanding what a biological design could do, that design's sequence or predicted function still cannot indicate what someone wants to do with it. That requires looking beyond the biology—to the customer, the transaction, the surrounding behavior, and the steps required to turn information into operational capability.

Much of the infrastructure needed already exists. Enterprise systems determine who is authorized to purchase what, from whom, for which project, and where it should be delivered. E-commerce platforms authenticate customers and establish relationships between buyers and sellers.

Payment systems look for indications of fraud. Logistics companies track how goods move through complex networks.

Universities and biotechnology companies already participate in much of this infrastructure through purchase orders and tracking systems, lists of approved vendors, and purchasing authorities. The biosecurity community must consider how to adapt this existing infrastructure to sequence ordering and other biological procurement.

One strategy is a *digital handshake* between biological suppliers and their customers. This would not require unrestricted disclosure of information. Instead, it could provide limited confirmations: Is this individual affiliated with the institution? Are they authorized to make this type of purchase? Is the transaction associated with a legitimate institutional account or project? Does the delivery location make sense? Most transactions would presumably produce unremarkable answers and proceed unencumbered.

As with suspicious activity, responses should be graduated. A new shipping address, an atypical equipment purchase, or another unusual signal might prompt authentication. Multiple signals might justify correlation with additional information to improve the signal-to-noise ratio. A stronger pattern should warrant human assessment, consultation with the institution or vendor, and, when appropriate, referral to law enforcement.

The biosecurity equivalent of suspicious-activity analysis should not be an automated system that declares someone a bad actor or mechanically stops a purchase. In fact, trying to train AI to recognize bad actors may prove inherently difficult to scale. Known threat indicators have value, but they inevitably reflect what is already known. As technologies and markets evolve, new applications—and new opportunities for misuse—emerge. Biotechnology evolves, and a biosecurity system that cannot adapt and iterate alongside it is likely to reproduce the same known-knowns problem that confronts sequence screening.

A more durable approach would be to develop a network or ecosystem capable of understanding normal activity well enough to recognize meaningful aberrations. Connecting existing transactional systems could gradually establish a baseline of legitimate activity across the biological marketplace. That baseline would demonstrate how technologies are actually being used, how markets are evolving, and when combinations of activities begin to depart from expected patterns.

This kind of network could also provide an increasingly valuable landscape for AI-assisted analysis. Rather than teaching AI what a criminal or bioterrorist looks like, it would give AI enough contextual information to help answer a more useful question: Does this activity make sense in the context of legitimate research use? The detection architecture could learn along with the market it is observing.

There may be an opportunity to extend this approach as laboratory automation advances. Commercial and

government research efforts are exploring how AI can be used in the lab not simply to perform experiments, but to learn from the results and inform subsequent research. That same learning environment could provide an opportunity for biosecurity: As AI systems learn how to design and conduct biological research, detection systems could learn alongside them, gaining context about legitimate activity and improving their ability to recognize meaningful departures.

Properly designed, this would protect legitimate science rather than burden it. Better fidelity should reduce indiscriminate scrutiny and focus limited human attention on the relatively small number of patterns that actually warrant review.

And someone must be designated to receive and follow leads. During my FBI career, including 15 years at the Weapons of Mass Destruction Directorate, I championed the notification of FBI WMD coordinators—personnel who are trained to assess and respond to WMD threats and who partner with local law enforcement, first responders, and the public health community—whenever suspicious biological orders surfaced. Specialized information is useful only if it reaches a designated person who is also capable of understanding why it matters, asking informed questions, and determining whether further scrutiny is warranted. I was dismayed earlier this year to learn the FBI would dissolve the directorate and fold its functions into general counterterrorism efforts. Without that clear responsibility for assessment and enforcement, digital handshakes cannot provide protection.

Beyond the sequence

The biosecurity community needs to bring together the capabilities of AI to gain insight into biological function and human behavior and to harness the power of AI for detecting and mitigating biotechnological threats. These two approaches—digital handshakes and biological terrain mapping—deserve as much attention in the biosecurity community as the question of how generative AI might create new biological threats.

But more than relying on AI to recognize a dangerous sequence or identify a bad actor, the community should be asking how AI can help create an ecosystem of trust, validation, and detection—one capable of revealing both the biological terrain and the operational landscape surrounding it.

Nearly two decades after I began my career in biosecurity, the fundamental questions remain: What are people ordering? With what intent? But the approach must change to expand how risk is evaluated and bring broader expertise to the table. Otherwise, efforts will simply refine policy without solving the problem.

Edward H. You is the founder of EHY Consulting LLC. From 2005 to 2025, he worked as an FBI special agent specializing in bioweapons and emerging technologies.