OSTP-TECH-2023-0007-0001

Addressing question numbers 4, 7

The Institute for Progress (IFP) is a non-partisan research and advocacy organization dedicated to scientific, technological, and industrial progress while safeguarding America's future. We appreciate the opportunity to provide input on the national priorities for artificial intelligence (AI) and the development of a National AI Strategy. AI has tremendous potential to advance scientific discovery, innovation, and societal well-being, but also creates risks that require responsible governance. In this comment, we specifically focus on the intersection of AI and biosecurity, and discuss the benefits and risks associated with AI within this context.

Intersection of AI and Biosecurity

Some of the chief concerns within the realm of biosecurity are Global Catastrophic Biological Risks (GCBRs), events of catastrophic scale involving biological agents that could lead to a "sudden, extraordinary, widespread disaster beyond the collective capability of national and international governments and the private sector to control." These events would significantly surpass the scale of crises like the COVID-19 pandemic.

AI technologies hold transformative potential for strengthening biosecurity and mitigating the risk of GCBRs. AI can assist in real-time epidemic detection, provide more accurate outbreak risk analyses, aid in gene synthesis screening to prevent harmful biological manipulations, improve disease trend forecasting, and accelerate the development of medical countermeasures. These capabilities could enhance the speed and efficacy of our prevention and response to biological threats and contribute significantly to global health security.

However, AI also presents substantial biosecurity risks. By democratizing access to sophisticated tools and techniques, AI could inadvertently lower the barriers to the misuse of biological agents by a wider array of actors, both state and non-state. Consider, for example, that the historical challenges Iraq's biological weapons program faced were largely due to significant limitations in technical expertise. Today, AI could fill that knowledge gap and facilitate the design, production, and deployment of these agents by actors who previously lacked the necessary expertise or resources. It could also help actors circumvent existing detection mechanisms, such as the Department of Homeland Security's BioWatch program or the Federal Select Agent Program's listed organisms.

Moreover, AI could increase the potency of biological weapons by enabling the design of biological agents with properties designed to inflict maximum damage — including increased transmissibility, greater lethality, and evasion of existing medical countermeasures. By allowing for the creation of threats beyond our current understanding and response capability, AI tools increase the risk of a GCBR materializing.

Given AI's dual-use potential when applied to biosecurity, it is crucial for policymakers to develop a nuanced understanding and a robust strategy. Crafting such a strategy will require a

concerted, interdisciplinary effort involving AI specialists, biosecurity experts, policymakers, and other stakeholders. Through collaborative, informed decision-making, the federal government can leverage AI's full potential for biosecurity, while ensuring the necessary precautions are in place to prevent misuse.

The Positive Role of AI in Biosecurity

One of the most promising roles of AI in biosecurity is its ability to support epidemic detection and disease trend forecasting. With AI, real-time analyses of vast volumes of sequence data, clinical data, social media activity, and other pertinent digital interactions can promptly identify the emergence of potential disease outbreaks. This capability not only allows for more rapid detection of a dangerous novel pathogen but also enables a targeted response. Furthermore, AI algorithms can be trained to predict future disease trends by helping parameterize infectious disease models and even predict the emergence of new biological risks by supporting human forecasters. Such forecasting capabilities offer invaluable lead times, enabling public health authorities to proactively implement preventive measures.

AI's preventive potential extends into the realm of synthetic biology. Sophisticated AI systems can be integrated into the gene synthesis process, making it possible to preemptively identify and flag gene sequences that could yield dangerous biological agents. This includes cutting-edge methodologies, such as random adversarial threshold search for automated and secure DNA synthesis screening, as well as cryptographic protocols that enable screening orders against a database of hazardous sequences (while safeguarding the secrecy of those sequences). As accessibility to gene-editing tools increases, these AI safeguards against potential misuse will become essential.

Leveraging AI-powered tools in protein design and drug discovery can dramatically expedite the development of essential medical countermeasures, such as vaccines and treatments. For instance, AI has already recently enabled the discovery and design of a new drug candidate that is entering Phase 2 clinical trials.

Finally, AI can assist in critically evaluating and fortifying our existing biosecurity protocols, through its ability to identify patterns, evaluate trends, and highlight potential vulnerabilities. In particular, cybersecurity is a crucial aspect of AI-enhanced biosecurity. Ensuring the security of AI systems, the data they access, and the insights they generate is paramount. AI can detect anomalies or breaches in digital systems that could indicate cyber threats, securing both sensitive bio-information and the AI systems themselves.

Risk Landscape & Strategies to Address it

Biosecurity risks associated with AI include potential misuse of Large Language Models (LLMs), along challenges posed by Biological Design Tools (BDTs) — a useful differentiation recently developed by Jonas Sandbrink. These risks, along with cybersecurity vulnerabilities, are formidable but not insurmountable, given a proactive and comprehensive risk mitigation strategy.

Large Language Models (LLMs)

LLMs, AI systems that can generate human-like text, carry substantial potential for misuse in the realm of biosecurity. Given their ability to efficiently synthesize information and provide insight into complex topics, they can inadvertently lower barriers to dangerous knowledge — thus enabling actors with limited expertise to gain insights into processes relevant to biological weapons development. For example, when prompted about influenza virus production, an open-source LLM could outline important laboratory steps, potentially providing a roadmap for malicious actors.

Furthermore, the ability of LLMs to function as lab assistants by providing tailored instructions and troubleshooting can significantly reduce the tacit knowledge barriers that have traditionally impeded bioweapons programs. By making success less dependent on expertise and skills, LLMs lower barriers for smaller, more concealable teams. In a recent classroom exercise, LLMs enabled undergraduate students to identify four potential pandemic pathogens within just one hour, along with their synthesis and the names of synthetic DNA providers that do not screen orders. The LLM also shared the fact that anyone without relevant scientific skills could contact a contract research organization to assist them.

As one potential intervention, pre-release evaluations of LLMs — ideally, by a government-led or government-supported team — should be conducted to assess potentially dangerous biological capabilities.

Biological Design Tools (BDTs)

BDTs, especially protein design tools, pose another unique set of biosecurity risks. While these tools could enhance biological design capabilities and contribute significantly to biosecurity, they can also advance offensive biological capabilities for sophisticated state and non-state actors.

For instance, BDTs could enable the creation of pandemic pathogens optimized for transmissibility, virulence, and immune evasion. These could be potentially worse than any currently existing threats.

In the near term, BDTs could circumvent sequence-based biosecurity measures. To mitigate these risks, use of BDTs should be monitored and interventions to deter biological misuse may be necessary. OSTP should consider the need for structured access for the narrow set of BDTs with high biosecurity risks, for instance by requiring user authentication and documentation of biosafety and dual-use review for potentially hazardous applications.

AI applications for biosecurity offer significant potential to enhance national security through improved detection and response capabilities. But the new vulnerabilities they introduce could be exploited to compromise national security. It is imperative that the federal government navigate this landscape with thoughtful governance, informed research, and an interdisciplinary approach to maximize the benefits and minimize the risks.