

White Paper: AIxBio Global Forum Structure and Goals

The Challenge

Significant advances in AI in recent years offer tremendous potential benefits for modern bioscience and biotechnology by supporting the rapid development of vaccines and therapeutics, enabling the development of new materials, fostering economic development, and helping to fight climate change. If applied equitably, these benefits have the potential to transform global health and overall well-being.

However, AI-bio capabilities—meaning AI tools and technologies that enable the engineering of living systems—also could be accidentally or deliberately misused to cause significant harm, with the potential to cause a global biological catastrophe. Given the rapid development and proliferation of these capabilities, leaders in government, bioscience research, industry, and the biosecurity community must collaborate to anticipate and address emerging risks by developing strategies to protect against misuse.

Current efforts to govern AI have been led by a handful of mostly western countries, including the US, the UK, and the European Union. The launch of the Global AIxBio Forum marks a significant step. Previously, there was no international platform to convene diverse stakeholders to address pressing biosecurity risks associated with rapid advances in AI-bio capabilities.

Moreover, ongoing efforts to develop technical solutions and guardrails for safeguarding AI-bio capabilities and guarding against misuse have been fragmented, limited to a handful of companies and organizations based mostly in the West, and lacking a comprehensive, strategic approach.

The AIxBio Global Forum will address these challenges by providing a platform for international experts, policymakers, and other key stakeholders to identify and address biosecurity risks posed by the convergence of AI and the life sciences (AIxBio).

Regardless of the current level of AI development and use within their borders, all countries have a stake in the development of safeguards to prevent misuse of AI-bio capabilities because the consequences will almost certainly extend beyond state boundaries. As one participant of the AIxBio Global Forum noted, it will take years of investment to build the systems a country would need to benefit from AI, but the risks will come for free. International collaboration will be important for establishing clear guidelines and effective risk reduction measures that do not stifle innovation.

Proposed Solution

The **AIxBio Global Forum** is designed to address the challenges outlined above by fostering international collaboration to reduce global risks associated with the convergence of AI and the life sciences and to realize the benefits of these technologies. The Forum is designed to be a platform to facilitate collaboration among international stakeholders and experts, including AI model developers in industry and academia, biosecurity experts, life scientists and government policy makers. The Forum will seek to incorporate and account for diverse points of view from across the global north and south with special consideration given to variances in perception of risk and stages of technological development. This Forum will serve as an international hub for activities focused on characterizing biosecurity risks associated with AI advances while also considering the potential benefits of these technologies; developing and sharing tools and best practices for evaluating biosecurity risks associated with AI-bio capabilities and developing guardrails to reduce these risks; and establishing global norms to support these practices.

Scope of the Forum

The long-term goal of the Forum is to develop and strengthen global norms and practical solutions for safeguarding AI-bio capabilities. The Forum will focus on three operational goals:

- 1) Developing a shared vision of emerging risks and benefits associated with the convergence of AI and the life sciences
- 2) Establishing, evolving, and guiding the implementation of a research agenda aimed at reducing biosecurity risks associated with the application of AI in biology
- 3) Disseminating resources, tools, and best practices to help the community develop effective oversight and risk reduction approaches to minimize the risks of the use of AI in the life sciences while maximizing potential benefits

The Forum will consider a wide range of AI tools and capabilities that intersect with biology, including large language models, AI-enabled biodesign tools, automated science, and others. The specific risks that the Forum will address and the risk reduction solutions that it will advance will be driven by a strategic research agenda. These will evolve as technology advances.

Proposed Structure

The AIxBio Forum should prioritize international reach and impact, incorporate technical expertise and the ability to innovate, be agile and dynamic to keep up with fast-moving advances, and be able to interface with other international organizations and national governments. The Forum should be structured in a way that enables it to embody these characteristics. To inform our recommendations, we analyzed the structures of several organizations with related mandates (see Table 1).

Table 1. Analysis of Organizational Structures

Name	Structure	Pros	Cons
Global Partnership on AI (GPAI)	Technical secretariat within OECD	Globally representative group with technical and social expertise	Does not develop standards, tools, or operationalize recommendations
Partnership on AI (PAI)	Independent, non-profit organization with for-profit and non-profit members	Provides research, guidance, and recommendations for businesses engaging with AI	Has not yet developed standards or tools and does not have governmental representation
Coalition for Epidemic Preparedness Innovations (CEPI)	Foundation that accepts donations from private and public sectors to finance independent research projects	Can address risks on the horizon and supports technology development	Does not develop standards or tools but instead shapes markets to incentivize new technology development
World Institute for Nuclear Security (WINS)	Independent, non-governmental membership organization	Supports individuals working with nuclear materials	Does not set standards or support research
The International Civil Aviation Organization (ICAO)	UN agency	Representative organization, produces technical standards, and considers social impact	As a UN agency, standards take more than 2 years to develop
International Biosecurity and Biosafety Initiative for Science (IBBIS)	Independent foundation that engages industry as well as public/private stakeholders	Develops tools and seeks to strengthen standards	May have limited reach without official ties to UN or other multilateral fora
World Economic Forum (WEF)	Global not-for-profit	Promotes public-private cooperation through communities of purpose	Does not set standards or develop tools

Based on our analysis of existing models, we recommend that in the short-term, the AIxBio Global Forum operate in a similar fashion to independent, international organizations like WEF and IBBIS. Key features of these organizational models that would be helpful in advancing the Global AIxBio Global Forum include:

- **Pursues Practical Approaches:** The AIxBio Global Forum has the long-term goal of developing and strengthening norms for safeguarding AI-bio capabilities while working in the short-term to develop practical tools to reduce urgent and emergent risks. This is analogous to IBBIS’s biosecurity mission. To accomplish this, IBBIS convenes a technical consortium composed of substantive experts to advance specific project objectives, which may be a model the Forum can adopt.
- **Independent and Incorporates Multiple Sectors:** Both WEF and IBBIS market themselves as effective conveners of stakeholders across the public and private sectors. WEF is well known for its ability to establish trust among business and political leaders. It will be important for the AIxBio Global Forum to maintain a similar stance, so industry and government alike trust the Forum to provide a platform and tools based on the best scientific advice, which is insulated from undue political influence. To achieve this, the Forum should seek to launch as an independent, non-profit entity and maintain transparency regarding its financial status, similar to IBBIS. It should also seek to include industry and public sector participants and ensure that these sectors are represented as part of the Forum’s governance structure, similar to the model established by WEF.
- **International and Agile:** Although AI model development is currently concentrated in a handful of states, this will not always be the case. Moreover, all countries have a stake in safeguarding the application of this technology in the life sciences, since the implications of misuse or abuse can be far reaching. The Forum should consider how to balance the need for representation across regions with the need to remain action-oriented and agile. The structure of IBBIS as an independent international organization—including the structure of its board and advisory group—is designed to balance these two needs.

Over the longer term, to build international influence and ensure sustainability of this initiative, the structure of the Global AIxBio Forum will need to evolve.

One option for the long-term structure of the Global AIxBio Forum would be to fold it into a UN agency or other international standard-setting body responsible for advancing AI safety and security issues. This option draws parallels to the International Civil Aviation Organization (ICAO) structure and function. The key features, benefits, and challenges of such a structure would be:

- **Credible and Deliberate:** ICAO is a UN agency responsible for developing policies and standards, performing analyses, undertaking audits, and building capacity. Its status as a UN agency lends it the credibility and convening power to help 193 countries to “cooperate together and share their skies to their mutual benefit.” AI safety and security, like aviation safety and security, would be an appropriate focus for a UN agency given the widespread use of the technology and the potential for severe damage if an accident or deliberate event were to occur. The ICAO maintains credibility in its norm and standard development process by

following “a structured, transparent and multi-staged process... involving a number of technical and non-technical bodies, which are either within the Organization or closely associated with ICAO.” This lengthy process takes an average of two years which likely is too slow for the current, rapid state of AI development. However, it may be possible to develop a more streamlined, rapid version of this process that is more appropriate to address current and future AI-bio risks.

- **Inclusive and Impartial:** ICAO was established via a convention that establishes general principles or rules, and all countries that have signed onto that convention (193 countries) are invited to participate in ICAO assembly sessions. However, the Council, which is charged with decision making when the assembly is not in session, is composed of elected states that are of “chief importance” or are large contributors to air transport. This may be a way to balance the need for broad representation with the need to highlight the experience of those with the highest concentration of AI development and innovation. ICAO convenes independent experts to carry out the technical work program of the organization with observers from industry and member states. Areas of focus for this work program are defined by the Council. This can be a model for developing impartial recommendations and for creating and maintaining a focus on reducing biosecurity risks, building on the early progress made by the AIxBio Global Forum.
- **Focus on Responsible Growth:** ICAO is an organization explicitly focused on sustainably growing the global civil aviation system, but it recognizes that safety and security (among other factors) are critical elements of responsible growth. Adopting a similar approach—preserving the beneficial application of AI by safeguarding against misuse—will be important to maintain support from a broad range of stakeholders, including governments, industry, and the life-science research community. Among these, industry support will be vital to the development of effective AI governance approaches. A challenge for any AI-focused forum will be identifying who should be included as industry representatives since the AI field is not as clearly defined as the aviation industry.

A second option would be to maintain the AIxBio Global Forum as an independent, international organization. To justify this option, the Forum will need to demonstrate its ability to develop effective risk-reduction approaches and that key stakeholders are adopting these solutions. The benefit of this approach is that the Forum could maintain its agility and its focus on the AIxBio issue set. The downside might be limited ability to influence global norms, which could be mitigated through partnerships with the UN or other industry standard-setting organizations.

A third option would be to fold the AIxBio Global Forum into an AI-focused initiative like the Partnership on AI. The benefit of this approach is that the Partnership has already convened diverse stakeholders and technical experts to address AI-enabled challenges and opportunities. Although the Partnership does not currently have an explicit focus on reducing biosecurity risks, this is an area of work that could be built out. The potential downside of this option is that the Partnership is not a UN agency so it may not have the reach or credibility to influence governance structures of a diverse set of countries.

Although AI may pose severe risks if misused, the odds of reaching consensus among countries to form a new UN body or standard-setting organization focused on AI governance is likely to be low in the current geopolitical environment. If such consensus is not attainable, a more feasible option—at least for the foreseeable future—may be maintaining the Forum as an independent organization focused on the intersection of AI and the life sciences with the goal of influencing broader AI safety and security norms through relationships with other standard-setting bodies. Over time, the Forum could pursue its goals of advancing international governance of AIxBio capabilities through a ‘coalition of the willing’ (or mini-lateralist) approach, with support from a growing group of nation states. The evolution of the AIxBio Global Forum structure will depend on the changing needs of relevant stakeholder communities and the state of overall AI governance. Going forward, a key consideration and priority should be identifying ways that the AIxBio Global Forum can improve its ability to influence global norms, by shaping both its institutional structure and its activities.

AI-Bio Forum Working Groups

Borrowing from the WEF and IBBIS models, the Forum will convene working groups focused on developing practical solutions to emerging risks at the intersection of AI and the life sciences. The working groups will be time-bound, and their areas of focus will be shaped in consultation with international experts via meetings of the Forum and/or external consultations. The Forum will develop and maintain a roadmap that outlines how the working-group priorities contribute to a broader research agenda aimed at reducing risks at the convergence of AI and the life sciences. Based on consultations with founding members of the AIxBio Global Forum, initial working group activities could include:

1. Building a shared understanding of the threat model for AIxBio capabilities, horizon scanning for emerging AIxBio capabilities, and sharing best practices for conducting evaluations to assess biosecurity risks associated with AI models.
2. Developing guardrails for biodesign tools (BDTs) to prevent misuse while safeguarding beneficial applications.

Membership of working groups should be based on interest and ability to substantively contribute to development of risk-reduction solutions. Working groups can recommend and facilitate technical partnerships for further tool development and refinement. The outcomes of the working group should be shared among the full Forum, considering the various levels of AI development and awareness among international stakeholders. Working groups should aim to develop solutions that are effective in terms of their risk-reduction potential but realistic in terms of the likelihood for adoption across diverse environments. As working groups are established, they should outline what their initial goals are and how they plan to pursue those goals.

Avoiding Information Hazards

A key challenge for the AIxBio Global Forum will be to achieve the appropriate balance between openness and transparency—which are critical for international engagement and building effective norms—and securing hazardous information that could enable malicious actors. Information hazards

are likely to arise in discussions and resources that identify biological hazards, concerning outputs from AI models, or details of how guardrails are implemented. Initial convenings of the Forum will seek to establish a shared understanding of how information hazards should be managed in this context.

It may be difficult to share resources and best practices for AI guardrails among AI model developers because this information could be seen as proprietary intellectual property. Providing this information could also constitute an information hazard, so it will be important for the Forum to establish practices for sharing.

Financial Support

NTI, a non-profit global security organization, is taking the first steps to establish the AIxBio Global Forum, but the goal is for the Forum to operate as an independent entity with dedicated staff in the next 2 to 3 years. Funding for the Forum may come from philanthropies or governments, but NTI or any future leadership of the Forum should take steps to insulate it from undue influence of any one organization or government.