

The United Nations' role in international AI governance

Summary paper of a workshop held on 28 February 2024

Background

This note provides a summary of a workshop discussion exploring proposals for new institutional functions related to international governance of artificial intelligence (AI). The workshop was jointly hosted by the Royal Society, the United Nations (UN) High-Level Advisory Body on AI, and Responsible AI UK on 28 February 2024 at Carlton House Terrace. It was chaired by Dame Wendy Hall FEng FRS, a member of the UN Advisory Body, and was opened with remarks from Dr Amandeep Singh Gill, the UN Secretary-General's Envoy on Technology.

The workshop convened a diverse range of experts in order to scrutinise and suggest improvements to the UN Advisory Body's interim report *Governing AI for Humanity*, published in December 2023. The discussions focused on the following four (out of seven) institutional functions outlined in the report:

- To assess regularly the future directions and implications of AI.
- To develop and harmonise standards, safety, and risk management frameworks.
- To promote international collaboration on talent development, access to compute, infrastructure, building of diverse high-quality datasets, responsible sharing of open-source models, and AI-enabled public goods for the sustainable development goals (SDGs).
- To monitor risks, report incidents, and coordinate emergency response.

These functions were selected to focus conversations and best reflect the expertise of participants invited to attend the workshop. This paper represents a summary of discussions and highlights key themes for consideration. It is not intended as a verbatim record of discussions and does not necessarily represent the views or positions of any participants or organisations who took part. It was drafted by staff at the Royal Society and submitted as feedback to the UN Advisory Body on AI.

The Royal Society

The Royal Society is a self-governing Fellowship of many of the world's most distinguished scientists drawn from all areas of science, engineering, and medicine. The Society's fundamental purpose, as it has been since its foundation in 1660, is to recognise, promote, and support excellence in science and to encourage the development and use of science for the benefit of humanity. The Society's strategic priorities emphasise its commitment to the highest quality science, to curiosity-driven research, and to the development and use of science for the benefit of society. These priorities are:

- Promoting excellence in science
- Supporting international collaboration
- Demonstrating the importance of science to everyone.

Responsible AI UK

Responsible AI UK (RAI UK) is a £33m research and funding programme looking to provide science-based advice to policy makers and industry. RAI UK brings together leaders from across the four nations of the UK to research how the development of AI should be shaped to benefit people, communities, and society.

Summary of key takeaways

- Consensus building on the future implications of AI should not be limited to finding agreement but entail the dedication of a discussion forum to platform a diverse range of viewpoints.
- Revisions to the report would benefit from providing more detail on proposed activities relating to horizon scanning and risk assessment.
- There could be a role for the UN in the organisation of future Global AI Safety Summits, creating a common vocabulary on AI, and coordinating agreements on basic minimum standards for Member States to enforce in their jurisdictions. These do not need to be limited to technical standards but could apply to the monitoring of environmental footprints and human rights impact assessments.
- A detailed report setting out evidence and recommendations on global gaps related to data, compute, and skills capacity could be a useful contribution from the UN.
- The ‘talents’ necessary to foster AI development will be varied and not solely limited to technical expertise. Attention will need to be given to a range of areas including data governance, bias detection, organisational dynamics, and cultural differences.
- As AI is applied across domains and can have an amplification effect on ‘non-AI’ harms, it is not clear what constitutes an ‘AI-specific’ incident for reporting purposes.

INSTITUTIONAL FUNCTION 1: ASSESS REGULARLY THE FUTURE DIRECTIONS AND IMPLICATIONS OF AI

This proposed institutional function entails the regular assessment of future directions and implications of AI, for which there is currently no established global mechanism. Similar to the International Panel on Climate Change (IPCC), this function would involve independent expert-led processes to provide policymakers with scientific insights (potentially at six-month intervals). This function also seeks to enhance public understanding, foster international consensus, and coordinate research efforts on the social impacts of AI.

Given the ever-changing nature of AI developments, in contrast to the relatively stable evidence and indicators on climate change, modelling this function on the IPCC may not be entirely suitable. Instead, the UN may be better positioned to inform member state governance by convening existing (and emerging) AI safety institutes, promoting research, and sharing best practices. These best practices could include those within existing frameworks such as UNESCO’s global agreement on the ethics of AI¹ and the G7’s Hiroshima AI Process.² These frameworks are mentioned in the interim report but it was felt that revisions to the report would benefit from making clear how any new institutional function would interface with them.

Consensus building should not be limited to finding agreement but entail the dedication of a discussion forum to platform a diverse range of viewpoints. Relying solely on seeking agreement as a default approach could constrain the UN’s capacity for forward-thinking initiatives and limit its effectiveness in addressing complex global challenges. The aspect of enhancing public understanding of AI, along with the UN’s role in achieving this objective, remains less developed compared to other components of the report and requires further elucidation.

Tension exists between the need for legitimate international consensus processes and the imperative to work at the pace of AI development. The mandate and scope of a new institutional function should aim to reconcile this tension. As it is, this function risks taking a reactive approach by collating existing insights rather than leadership and agenda setting.

1. UNESCO. 2021 UNESCO member states adopt the first ever global agreement on the Ethics of Artificial Intelligence. See <https://www.unesco.org/en/articles/unesco-member-states-adopt-first-ever-global-agreement-ethics-artificial-intelligence> (accessed 6 March 2024).

2. Government of Japan. Hiroshima AI Process. See <https://www.soumu.go.jp/hiroshimaaiprocess/en/index.html> (accessed 6 March 2024).

The activities described, namely horizon scanning and risk assessment are broad as written and would benefit from more detail. These could include the purpose of activities (for example, risk anticipation, regulation or impact on particular SDGs) and whether this applies to technology development, access, deployment, or downstream impact. Further challenges to horizon scanning will include methodological considerations, including frequency of activity and time frames (near versus long-term); indicators and metrics; scope; and audience (whether scientific community, policy makers or the wider public). In addition, it may be useful to set out how any new horizon scanning activity would differ from existing initiatives such as the OECD's Expert Group on AI Futures.³

Horizon scanning and risk assessment may also encounter cultural challenges and the perceived or measured utility, risk and impact associated with an AI system may be inconsistent in various settings. If undertaken, these foresight activities would benefit from a broad range of expertise including, for example, engineers, social scientists and public engagement professionals. These, however, may require significant resourcing to attract and retain talent if these assessments are to be conducted on a very regular basis.

In addition, there is tension around stakeholder engagement in scanning and assessment. Industry engagement from AI developers will be imperative, but will require clear incentives to participate. Currently, there is no information sharing required to participate in the AI market (unlike, for example, international aviation).

Alternative approaches to conventional horizon scanning could be explored. The UN might instead focus on synthesising others' horizon scanning insights^{4,5} or develop empirical methodologies to assess the implications of AI. The UN could also play a role as an aggregator, compiling an inventory of AI safety and assessment tools to support stakeholders in navigating the complex landscape of AI governance effectively. Moreover, fostering collaboration with trusted organisations and building networks of expertise at local levels could enhance the efficiency and efficacy of AI governance efforts.

INSTITUTIONAL FUNCTION 3: DEVELOP AND HARMONISE STANDARDS, SAFETY, AND RISK MANAGEMENT FRAMEWORKS

Identifying a need to harmonise and align approaches to AI standards, safety, and risk management, this proposed institutional function sets out a role for the UN to bring states together to agree common frameworks. Suggestions in the interim report include networking national AI safety institutions and creating new standards on the environmental impact of AI.

The role of the UN as a global convener presents a significant opportunity. In addition to networking emerging AI safety institutes, there could be a role for the UN in the organisation of future Global AI Safety Summits, creating a common vocabulary on AI, and coordinating agreements on basic minimum standards for Member States to enforce in their jurisdictions. These standards would not need to be solely related to technical issues but could relate to the monitoring of environmental footprints or human rights impact assessments. As well as setting these standards, there may be a role for the UN in terms of monitoring and auditing their enforcement.

3. AI Futures. OECD AI Policy Observatory. See: <https://oecd.ai/en/site/ai-futures> (accessed 6 March 2024).

4. UK Government. 2023 State of the science report to understand capabilities and risks of frontier AI: Statement by the Chair. See: <https://www.gov.uk/government/publications/ai-safety-summit-2023-chairs-statement-state-of-the-science-2-november/state-of-the-science-report-to-understand-capabilities-and-risks-of-frontier-ai-statement-by-the-chair-2-november-2023> (accessed 6 March 2024).

5. World Economic Forum. 2023 Global Risks Report. See: <https://www.weforum.org/publications/global-risks-report-2023/> (accessed 6 March 2024).

To ensure that standard-setting on AI is not dominated by certain countries (eg in the Global North⁶), the UN has an important part to play in giving a voice, and role, for countries in the Global South⁷ who may otherwise have their needs and concerns excluded from approaches to AI governance. The inclusion of a broad range of stakeholders in standard setting, more generally, will also be necessary. Across all nations, in the Global North and Global South, there will be significant differences in views on AI as well as specific cultural and sectoral requirements. These requirements could affect, for example, ethical standards on fairness and bias in AI systems.

There was considered to be a lack of clarity in the interim report as to who the audience is for a global, harmonised approach to standards, safety, and risk management frameworks. Standards could be set at an international level for governments to adopt, at a sectoral level for industry actors, or at a hyperlocal level for individuals and professional bodies to apply. Understanding who this audience is will be necessary to understand where action needs to be taken, the design approach (e.g. domain-specific or domain-agnostic), and capacity-building across nations and sectors. By sector, approaches to AI governance will vary significantly due to different levels of sensitivity in data, existing regulations, and inherent complexities of domain-specific systems.

A domain-agnostic approach to governance (e.g. through mandating the use of model cards which explain what AI systems do, how they were constructed, and what data they were trained on) is a potential alternative that the UN could explore. It is also an approach which may lead to the provision of information with greater utility to key stakeholders in AI development and AI assurance.

Lessons may be learned from existing global harmonisation efforts. These include the Internet Governance Forum; the IEEE Standards Association; the OECD's expert group on AI risk and accountability; the International Labour Organization; the International Organization for Standardization (ISO), and the ISO's national counterparts. In particular, there may be lessons regarding enforcement best practice and learnings from failures to prevent existing technological harms.

Enforcement may present the most significant challenges in harmonisation efforts. Key issues relate to state sovereignty; the formation of public/private collaborations; global information exchange between regulators; the lack of recognised professional ethics norms in AI and data science; and resource disparities between nations for enforcement and capacity building. The resource disparity challenge is one where the UN could have a positive role to play and is one which will be necessary to solve if standards are to be globally enforceable and interoperable.

While there are very clear and specific challenges with this proposed institutional function, the convening role of the UN may be useful for the coordination of accountability across borders; the exchange of best practice between regulators and AI safety institutes; and for mandating minimum standards and reporting on the environmental and human rights implications of current and future AI developments.

6. The Global North refers to countries or regions typically classified as 'developed'. These regions include Europe, North America, Australia, New Zealand, and Japan. See: <https://www.un.org/en/development/desa/population/migration/publications/migrationreport/docs/MigrationReport2015.pdf> (accessed 6 March 2024).

7. The Global South refers to countries or regions typically classified as 'developing'. These regions include Africa, Asia (excluding Japan), Latin America, the Caribbean, Melanesia, Micronesia, and Polynesia. See: <https://www.un.org/en/development/desa/population/migration/publications/migrationreport/docs/MigrationReport2015.pdf> (accessed 6 March 2024).

INSTITUTIONAL FUNCTION 5: PROMOTE INTERNATIONAL COLLABORATION ON TALENT DEVELOPMENT, ACCESS TO COMPUTE INFRASTRUCTURE, BUILDING OF DIVERSE HIGH-QUALITY DATASETS, RESPONSIBLE SHARING OF OPEN-SOURCE MODELS, AND AI-ENABLED PUBLIC GOODS FOR THE SDGS.

Focused on addressing the UN's SDGs⁸ (which cover a range of issues including poverty; hunger; gender equality; innovation; and climate action), this proposed institutional function aims to improve global data access, compute resources, and training. Drawing parallels with the European Organization for Nuclear Research (CERN), the European Molecular Biology Laboratory (EMBL), and the International Atomic Energy Agency (IAEA), the proposed function seeks to incentivise the pooling of resources around shared goals.

The lack of relationship between global AI development and the SDGs suggests there may be a role for the UN to play in ensuring a greater link between the two. Whether achieved through a new institution or through existing institutions, key priorities would include the need for high-quality national and international datasets; investment in capacity building; and greater collaboration between research groups across the world. Greater collaboration could expand access to resources (e.g. funding, cloud storage) for researchers in the Global South, particularly in the absence of major investments in new, local, compute infrastructure.

Investment in infrastructure, skills, and data in the Global South may be necessary to ensure that benefits from AI development can be experienced by all. There are clear challenges, however, related to funding, infrastructure time lags, and clarity on what resources would be needed. A detailed report setting out evidence and recommendations on global gaps related to data, compute, and skills capacity could be a useful element of this institutional function. Insights on the quality and quantity of non-English language data required to train local AI models could be a useful chapter for such a report. In the meantime, it is likely that AI development will continue to be driven by those who can afford access to compute power and high-quality training data.

Robust national data protection standards and intellectual property protections will also be critical to enabling global collaboration. This will, in part, be a question of regulation, but will also include challenges related to training, accountability, enforcement, and access mechanisms.

CERN is considered to be an exemplar for international collaboration that works well. For AI development, collaboration could include building shared models, training frameworks, and benchmarks. On the responsible sharing of open-source models, there is potential for the UN to play a role in convening expert bodies to share best practices as well as guidance on notification procedures on high-risk systems. This could include an exploration of anonymous disclosure mechanisms for documenting incidents across regions.

The 'talents' necessary to foster AI development will be varied and not solely limited to technical expertise. Solving the impact of issues such as deepfakes, for example, will require an understanding of social dynamics, information literacy, and educational approaches. Similarly, understanding the implications of AI on jobs will require the experience and understanding of workers, who are likely to have better insights on organisational dynamics than technical AI experts. The interim report was critiqued for the lack of depth it went into on issues related to data-labelling practices and the potential 'race to the bottom line' which could lead to companies seeking workers in countries with low pay and weak employee protections. A better understanding of these impacts may help determine whether or not AI developments equate to economic growth.

8. Adopted in 2015 by all UN Member States, these 17 goals aim to provide 'a shared blueprint for peace and prosperity for people and the planet, now and into the future'. See: <https://sdgs.un.org/goals> (accessed 6 March 2024).

INSTITUTIONAL FUNCTION 6: MONITOR RISKS, REPORT INCIDENTS, COORDINATE EMERGENCY RESPONSE

This proposed function addresses the need to monitor and respond to risks posed by the borderless nature of AI tools. These tools raise concerns about access to weapons of mass destruction, cyberattacks on critical infrastructure, and the potential misuse of AI for lethal autonomous weapons. To mitigate these risks, a global framework is proposed, drawing inspiration from macro-prudential frameworks in central banking, focusing on human rights principles. Additionally, reporting frameworks could mirror established practices from organisations like the IAEA and WHO, ensuring mutual reassurance and effective surveillance of AI-related risks.

The UN's role in defining, monitoring, and responding to AI incidents is important, as these incidents may require a coordinated global response. However, the function's scope lacks clarity. First, any definition of AI should account for the fact that AI is not a discrete sector, but a set of tools that can be applied to any domain or application. Second, in considering the 'borderless nature of AI tools' the UN should also account for global disparities in AI access and training data representation where AI is deployed in global contexts. Further, the UN might consider its role in identifying and mitigating these wider inequalities as AI risks.

It is unclear what constitutes an AI-specific incident⁹. Further detail around the definition of AI incidents will be key, specifically, whether 'incidents' include 'near-misses' (i.e. unplanned events that do not result in harm, but which could indicate a future risk). Response mechanisms and responsibilities could be further clarified. Autonomous vehicle crashes, for example, may fall under the purview of both AI and transport incident management. In these and other cases AI can have an amplification effect on non-AI risks, for example, the spread of online misinformation or cyberattacks. This emphasises the need for clear delineation of what constitutes an AI incident.

Risk conceptualisations vary regionally. As presented, the notion of risks focuses on AI application. However, there are also risks associated with global disparities in access to compute, data and talent. Considering risks and benefits requires examining trade-offs, especially between global north and global majority settings. Emerging AI technologies will be seen to have different advantages and disadvantages according to varying social and political contexts. The deployment of general-purpose AI further complicates risk assessment, obscuring potential outcomes and cascading effects that may not be foreseeable in advance. Information sharing concerning AI incidents and risks presents a challenge, as it may require defining incident levels and establishing detailed emergency response mechanisms to encourage reporting with clear outcomes. Currently, the purpose or outcome of reporting to the UN, as well as the potential for regulation and enforcement, remains unclear.

The UN could play a useful role in categorising AI incidents by assembling diverse expertise across technical, social, and environmental domains to understand risks. As with existing incident monitoring systems, distinguishing between incidents caused by deliberate actions and unintended consequences is essential. These unintended consequences may be multifaceted, stemming from human error, system failure, or unforeseen outcomes during system deployment. This underscores the need for interdisciplinary expertise in establishing evolving risk models, for example, drawing from human factors research.

Enhancing the capacity of existing regulators and governance mechanisms in specific sectors, such as health or bioweapons, to address AI-related incidents may prove more beneficial than creating separate AI-specific reporting and response structures. This could involve aligning response mechanisms and standards where reporting mandates already exist.¹⁰ Reporting should be managed sensitively and incentivised, for example, through anonymous whistleblowing channels (akin to cybersecurity reporting mechanisms for software incidents). Where commercial sensitivities are likely to disincentivise information sharing, tools should be put in place to collate information in an anonymous way.

9. For example, the OECD AI Incidents Monitor applies broad definitions of AI incidents and hazards. See: <https://oecd.ai/en/incidents-methodology> (accessed 6 March 2024).

10. World Health Organisation. 2023 Pandemic prevention, preparedness and response accord. See: <https://www.who.int/news-room/questions-and-answers/item/pandemic-prevention--preparedness-and-response-accord> (accessed 6 March 2024).

The distinct characteristics of AI incidents and risk reporting should be delineated. There may be lessons to be learned from established systems in other sectors, including central banks, aviation and health research.¹¹ Cybersecurity and nuclear safety also rely on self-report systems and systemic analyses to diagnose faults. It will be important to establish which factors are exceptional and specific to AI and develop agile evaluation metrics and benchmarks as there may be limits to transferrable lessons, particularly where AI is general purpose. Identifying ‘points of failure’, tolerance limits, and ‘never events’ for AI (serious incidents that are generally preventable) could be useful. Rather than spearheading response efforts, a UN-level body could serve as a hub for knowledge transfer to relevant international or civic response units (akin to medicines¹²).

Some industry players may be motivated to develop self-regulation to forestall stringent government oversight. A survey of existing industry practices, particularly where there is regional or global alignment, would be worthwhile in the development of global AI governance frameworks. There is room to further incentivise industry to share near-misses, particularly when dealing with frontier models. Drawing from examples like aviation, where companies are motivated to mitigate risks to benefit the public and avoid fines, such incentives could be extended.

There is an opportunity to monitor global, gradual risks linked to the SDGs including those with diffuse social impacts that may only become apparent through societal interaction, such as effects on mental health or equitable livelihoods. This encompasses the sustainability aspect of AI deployment, which can be resource-intensive, prompting the UN to enhance awareness and identify ‘energy-proportionate’ AI use cases likely to yield greater benefits than environmental drawbacks.

Other general reflections and considerations

Framing of AI and risk

A robust ontological framework is required to address the multifaceted aspects of AI, AI risks and response at the multinational level. The report could better distinguish between data and AI functionalities, particularly regarding the role of sovereign data, as well as implications of open-source AI tools. Some examples of AI risks in the report are examples of misuse of purpose rather than a technological limitation. There are inconsistencies between the report’s content and its depiction in the pyramid model, particularly concerning the emphasis on peer review.

There is a risk associated with narrowing the scope of these functions to address only current anxieties, which are at the forefront of policy concerns. This focus may lead to overlooking the establishment of functions to address more challenging and complex issues, such as long-term social and environmental harms, that may not be currently in the spotlight but are nonetheless crucial for consensus-building. Additionally, there is a danger of labelling certain topics as ‘out of scope,’ potentially resulting in discussions being limited to comfortable or easier subjects. Furthermore, the report could be more explicit about actionable items, such as automotive and medical devices, rather than emphasising broad aims like peace and security.

Tensions in global AI governance

The UN has an important role in establishing universal norms and values, but when delving into specific details, responsibility shifts to individual nation-states. This presents a challenge in formulating norms that are not solely based on Western ideals, considering the diversity of values across different cultures and regions. Implementing metrics to monitor progress, particularly regarding inclusivity, is crucial, although not all aspects are quantifiable, underscoring the importance of interdisciplinary methodologies and expertise. Building public confidence involves not only addressing risks but also championing the opportunities provided by AI, a task that requires explicit acknowledgment and promotion.

11. UK Government. 2022 Better, broader safer: using health data for research and analysis. See: <https://www.gov.uk/government/publications/better-broader-safer-using-health-data-for-research-and-analysis> (accessed 6 March 2024).

12. Medicines, for example, are monitored through various reporting channels; in the UK this includes: Medicines and Healthcare products Regulatory Agency. See: <https://yellowcard.mhra.gov.uk/> (accessed 6 March 2024).

Annex

The Royal Society would like to thank the following workshop participants who have contributed to the development of this note.

Workshop participants

Name	Organisation
Professor Abigail Sellen FEng FRS	Microsoft Research Cambridge
Aidan Peppin	Cohere
Dr Alex Mankoo	The British Academy
Dr Alexandra Smyth	Royal Academy of Engineering
Professor Alison Noble FEng FRS	University of Oxford
Dr Amandeep Singh Gill	United Nations
Ankur Vora	Google DeepMind
Anna Bacciarelli	Human Rights Watch
Brittany Smith	Open AI
Byonghyon Bak	Embassy of the Republic of Korea
Dr Chris Meserole	Frontier Models Forum
Dr Cosmina Dorobantu	Alan Turing Institute
Dr Declan McDowell-Naylor	Information Commissioner's Office
Dr Farzana Dudhwala	Meta
Professor Gina Neff	Minderoo Centre for Technology and Democracy
Professor Gopal Ramchurn	Responsible AI UK
Dr Hema Sridhar	International Science Council
Imran Shafi OBE	Department for Science, Innovation and Technology
Dr Jack Stilgoe	University College London
Dr James Dracott	UK Research and Innovation
Professor Jon Crowcroft FEng FRS	University of Cambridge
Kate Dearden	Community Trade Union
Dr Kris Shrishak	Irish Council for Civil Liberties
Professor Marta Kwiatkowska FRS	University of Oxford
Michael Birtwistle	Ada Lovelace Institute
Professor Dame Muffy Calder DBE FEng FRSE	University of Glasgow
Professor Nick Jennings FEng FRS	Loughborough University
Oliver Buckley	Partnership on AI
Father Paolo Benanti TOR	UN AI Advisory Board
Professor Petri Myllymaki	UN AI Advisory Board
Philippa Green	Government Office for Science
Rachel Coldicutt OBE	Careful Trouble
Professor Shivaji Sondhi	University of Oxford
Rashik Parmar MBE	British Computer Society
Rebecca Finlay	Partnership on AI
Professor Robert Trager	University of Oxford
Dr Rumman Chowdhury	Humane Intelligence
Sabhanaz Rashid Diya	Tech Global Institute
Sahar Danesh CEng	British Standards Institute

Name	Organisation
Professor Sana Khareghani	King's College London
Sarah Bérubé	OECD
Sasha Rubel	Amazon Web Services
Dr Seán Ó hÉigartaigh	University of Cambridge
Shahzeb Mahmood	Tech Global Institute
Dr Simon Staffell	Microsoft
Professor Stuart Russell OBE	University of California, Berkeley
Sue Daley	TechUK
Professor Tom Crick	Department for Culture, Media and Sport
Dr Vicki Seward	National Physical Laboratory
Professor Virginia Dignum	UN AI Advisory Body
Professor Dame Wendy Hall DBE FREng FRS	UN AI Advisory Body
Yasmin Ibison	Joseph Rowntree Foundation

Royal Society staff

Name	Role
Areeq Chowdhury	Head of Policy, Data and Digital Technologies
Caitlin Neuwelt-Kearns	Policy Adviser, People and Planet
Charise Johnson	Policy Adviser, Data and Digital Technologies
Claire Hurst	Public Affairs Adviser, Science Policy
Denisse Albornoz	Senior Policy Adviser, Data and Digital Technologies
James Hannaford	Programme Coordinator, Data and Digital Technologies
Dame Julie Maxton DBE	Executive Director
Dr June Brawner	Senior Policy Adviser, Data and Digital Technologies
Luke Clarke	Head of International Affairs, Americas/Multilaterals/Africa
Nicole Mwananshiku	Policy Adviser, Data and Digital Technologies
Poppy Joyce	Policy Adviser, Americas and Academic Freedom
Dr Rupert Lewis	Chief Science Policy Officer
Sumaiya Zahoor	Senior Programme Manager, Science Policy

© The Royal Society. Issued: March 2024 DES8911

The text of this work is licensed under the terms of the Creative Commons Attribution License which permits unrestricted use, provided the original author and source are credited. The license is available at: creativecommons.org/licenses/by/4.0