

Global Citizen Deliberation on Artificial Intelligence

Options and design
considerations

Global Citizen Deliberation on Artificial Intelligence

Options and design considerations

Design Lab Report: September 2024

Authors: Tim Davies, Claire Mellier, Kristophina Shilongo, Rich Wilson

Project website: <https://connectedbydata.org/projects/2024-gca-ai>

Developed as part of the **CONNECTED BY DATA** '[Growing Data Governance Communities](#)' programme generously supported by a grant from The Omidyar Network.



About this document

This report explores how global citizen deliberation, particularly drawing on the concept of a global citizens' assembly, could and should shape the future of artificial intelligence. Drawing on an extended design lab of in-depth interviews and workshops that took place in mid-2024, it presents a series of options for bringing the voices of those affected by AI development and deployment into decision-making spaces, through processes that can deliver informed and inclusive dialogue.

The landscape of AI governance is rapidly evolving. There are open questions at many levels, from setting shared values and visions to guide AI development, to designing specific governance mechanisms or safety standards, and shaping the models and rules for individual and localized applications. There is growing consensus that these questions cannot be answered by the technology industry or individual governments alone. Global publics must be meaningfully involved.

Over the last year, figures from academia, industry and civil society have put forward calls for a global citizens' assembly on AI. The concept is a powerful one: inviting individuals from across the globe to join in processes where they have access to expert insights, opportunities to learn, and facilitated space to deliberate together, bringing diverse perspectives and experiences to bear on questions of global importance.

In this report we address how established and emerging sites of global AI development and governance can integrate citizen deliberation, setting out five template options for citizens' assemblies on AI: [deliberative review](#) of AI summits and scientific reports; an independent [global assembly](#) on AI; a series of [distributed dialogues](#) organized across the globe; a technology-enabled [collective intelligence](#) process; and [commissioning](#) the inclusion of AI topics in other deliberative processes.

We present the strengths and weaknesses of these options, and outline additional design considerations they give rise to around recruitment, governance, agenda-setting, transnational dialogue and aggregation of findings, and the use of AI as a delivery tool. In doing so, we aim to support a critical assessment of emerging and future proposals for both citizens' assemblies, and wider forms of citizen deliberation on AI, at both the global and local level.

Table of contents

Introduction	4
Focus	7
Values and vision	8
Governance and regulation	9
Applications and implementation	10
Who sets the agenda?	11
Forum	12
Form	17
Option 1. Deliberative review: regional panels and global delegates	19
Thinking about recruitment	21
Option 2. Global assembly: independent transnational & local deliberation	22
Transnational deliberation & aggregating across borders	24
Option 3. Distributed dialogue: Local, national or regional conversations	25
Assembly Governance and Expert Inputs	27
Option 4. Collective intelligence: digital dialogue at scale	28
AI as an assembly delivery tool	31
Option 5. Commissioning body for public dialogue on AI	32
The costs of deliberation	33
Delivering change	34
From concepts to actions: next steps	37
Method & acknowledgements	39
References	40
About the authors	44

Introduction

Artificial intelligence – the use of advanced algorithms and vast quantities of data to take on tasks formerly thought to require human activity – has rapidly become one of the central global governance issues of our time. AI presents significant opportunities, but also great challenges. Cutting-edge AI development is concentrated in a small number of countries, but its impacts are global. When AI is deployed with effects as diverse as reworking supply chains and labor markets, changing agricultural practices, and reshaping media landscapes, even those who do not directly use AI systems are affected by them. It follows then, that AI should be governed “by and for all”¹.

But how can voices, concerns and ideas from people across the world be brought meaningfully into a complex, contested and often technical landscape of AI governance?

Global citizen deliberation offers one important answer.

Over recent decades, a wave of democratic innovation has seen deliberative practice and citizens’ assemblies deployed in local, national and transnational contexts to address complex topics, from climate change and genomics to social policy^{2,3}.

In a citizens’ assembly, a representative group of delegates are selected, offered balanced expert testimony, and given facilitated space to deliberate together on substantive questions. Their recommendations can inform or guide decision making (depending on the assembly design); offer new solutions; deliver legitimacy to policy choices; uncover key tensions; and/or mobilize wider public dialogue and action on an issue of concern⁴. The evidence is striking: time and time again these processes have shown that ordinary citizens can meaningfully address challenging policy topics – often more effectively than conventional decision-making – and that citizens’ assemblies can act as a valuable component of the wider governance ecosystem⁵.

In this paper, we explore the opportunities and challenges for global citizen deliberation on AI by examining a range of possible citizens’ assembly design options, looking at possible focus topics, potential places to ‘dock’ deliberation with existing decision making structures, and different forms that deliberation could take.

Our goals are to:

- demonstrate the viability of directly including citizens as a stakeholder group in the governance of AI
- advance conversation on the possible designs and delivery of global citizen deliberation on AI

This paper has been prepared for the [Coalition for a Global Citizens' Assembly](#), which is working to establish a standing citizens' assembly infrastructure alongside the United Nations. However, we hope that the options set out here are also of wider relevance.

Although we have focussed explicitly on the 'big ideal' of a global citizens' assembly, many of the ideas and components explored in this report could also be applied to other forms of deliberative participation, and at a firm, regional, national or even local level. Our hope is that this work also inspires anyone grappling with AI governance questions to think about how to bring diverse citizen voices more directly into shaping and making informed and deliberative decisions about our shared technological futures.

The design of global deliberation on AI should consider

Addressing the right questions

Focus



Values and vision



Governance and regulation



Applications and implementation

Docking with decision making

Forum

Multilateral: UN, OECD etc.

Industry and AI labs

Regional bodies

Academia and civil society

Designing for inclusion and impact

Form



Deliberative review



Commissioning body



Distributed dialogue



Collective intelligence



Global assembly

Focus

Global citizen deliberation on AI needs to address the right topics.

Every deliberation has a topic of focus: often presented through a guiding question. Artificial intelligence is a vast topic. A choice needs to be made about which aspects and impacts of AI global deliberation should focus on; how to frame guiding questions; and the kinds of outputs a citizens assembly might produce.

Questions around AI can be framed at one of three levels:

- **values and vision**
- **governance and regulation**
- **applications and implementation**

Drawing on interviews and workshops carried out for this project, we've identified a number of example questions and outputs that an assembly could focus on at each of these levels. These are far from exhaustive. Many other topics were suggested in our research and here we have sought to provide indicative examples from a range of different viewpoints.

We can see consensus that unrelated AI can cause harm, but there is uncertainty about AI's trajectory: even experts can't predict with certainty what AI regulation or alignment will look like in five years. We need quick action on imminent harms, and long-term approaches that lay the groundwork to shape AI according to society values as AI capability grows.

— Interviewee

In particular, contributors have noted the important relationship between AI and data governance, and the need to address the whole lifecycle of AI from research and development, through use and eventual decommissioning⁶, including how this lifecycle is driven by, or affects, different geographies at different points. Although we focus on AI-centered questions below, in some cases deliberation on AI may be better framed within the broader context of digital or data governance.

From my perspective, if we can deal with some of these data governance issues, we're dealing with some of the AI 'risk' issues.

— Interviewee

Values and vision

Although numerous sets of principles and ethical codes for AI exist, few have been developed through truly global and inclusive dialogue, taking into account perspectives from a full range of geographic, cultural, religious, political and socio-economic backgrounds. Global citizen deliberation on values and visions for AI could identify points of variation and consensus, build the mandate for global AI governance, and foster local conversations about how to adapt to a world with AI.

Example questions	Example outputs from a deliberative process
How do we live good lives alongside AI?	Principles and policy recommendations to promote human flourishing alongside AI, including in relation to the world of work.
What values should be built into all generative AI systems?	Perspectives on whether publics prefer monolithic AI that represents all cultures and worldviews, or many localized systems that work differently depending on context.
How can AI best deliver sustainable development and public good?	Public priorities for AI funding, effort and alignment; suggested moratoria on anti-social uses of AI.

We need to think not just about the future of work, but the future of life. If AI has impacts in the future the way some suggest, then we need to consider both the positive opportunities for human growth this offers, and the roles we can decide to reserve for humans.

— **Summary, Participants at Brussels workshop**

The technology has been developed, and it's out there, and governments are thinking of how to use it to improve efficiency ... and productivity. But there's not really been much discussion about how we use it as a society. How is it going to affect me? So involving the public globally to understand what the priority areas of focus should be could be part of the deliberative process.

— **Interviewee**

Governance and regulation

Beyond shaping overarching principles for governance, citizen deliberation could engage directly with the specific forms and features of AI governance and regulation⁷. This may include debating governance frameworks, inputting into specific governance tools (e.g. safety assessments), focussing on the governance of specific international AI uses (e.g. in autonomous weapons systems), or exploring specific AI-related issues such as intellectual property or data governance rules. Global CSO Atlas has proposed that a “*Citizens’ Assembly, informed by experts, [could] draft a Treaty to regulate AI globally*” positioning global deliberation as a powerful counterweight to government and industry initiation of regulation⁸.

Example questions	Example outputs from a deliberative process
What global rules should govern AI?	Proposed common rules for policy makers or firms to follow. Draft treaty on AI governance.
What control should local communities have over AI?	Reporting on regional and national variation with respect to key AI governance and alignment questions.
What should the design of an AI Safety Assessment take into account?	Public priorities to guide AI Safety Institutes; intercultural validation of assessment frameworks; review of international AI standards.
Should we allow Lethal Autonomous Weapons Systems (LAWS)?	Citizens declaration on use of AI in lethal autonomous weapons systems. Input to global standards for use of LAWS.
How should data be governed in order to deliver equitable AI?	Recommendations on data governance and management to mitigate risks and secure benefits of AI

...there are a couple of areas where it will be useful and desirable to have more voices, especially from the global south. One of them is regulation. [Right now] the spaces are ... very narrow, very technical and the conversation is often about replicating what is happening in the north.

— Interviewee

Who controls AI's expansion and decides its use is a critical question. While AI could drive progress, forcing it into spaces that don't want it, and embedding the values of its creators, raises ethical concerns. The challenge is balancing the benefits of AI with the right to opt out, ensuring its adoption isn't imposed on those who don't choose it.

— Interviewee

Applications and implementation

A citizens assembly could also focus on specific questions of AI use: from debating how chat-bots should respond to contentious issues, to discussing the situations and contexts in which AI can be used for front-line public service decision making, or agreeing the protocols that should be followed to test AI products before deployment. Outputs from deliberation on these topics could feed into higher-level regulatory action, into guidance for implementers of AI systems, or even directly into the future training of AI models to address gaps in the prior representativeness of their training material.

A number of interviewees also pointed to the use of AI as a tool that can facilitate participatory and democratic decision making, particularly at scale. While establishing the legitimacy of applying AI within democratic state processes is likely to require national rather than global deliberation, there may be opportunities for a global assembly to establish when and how AI tools can be applied for future large-scale global citizen input and engagement.

Example questions	Example outputs from a deliberative process
How should a chatbot respond to questions about contentious issues?	Principles to guide AI system training, statements for LLM fine-tuning, or raw transcripts representing diversity of perspectives to be used in LLM training.
How should AI be used within participatory or democratic decision making?	Guidelines and mandate for use of AI systems within the wider global participatory processes.
When is it appropriate for AI to be used to make decisions?	Application- (and context-) specific guidance on issues like health, law, education etc.

If you ask a question like ‘should medical assistance be given by AI without humans in the loop’, well, you might get very different answers in Europe or Africa. If in your region there are not enough doctors ... you might want AI to be giving responses....
— **Interviewee**

Deliberation could generate new content to train models on, or it could contribute to greater transparency and explainability of systems: it doesn’t only need to feed into policy-level decision making.

— **Online workshop discussions**

Who sets the agenda?

There is an interplay between the focus question, the forum to dock with, the form that an assembly might take, and ultimately the funding sources that enable it. For many citizens' assembly advocates, there is also a strong argument to be made that question-setting should be, at least in part, bottom-up rather than top-down⁹. This can be achieved through a standing citizens' assembly model (such as that operating in Belgium), in which, over time, alumni of earlier 'edition' of the assembly help shape the questions addressed in subsequent iterations¹⁰ or through a pre-assembly process such as the constituent network of the Democratic Odyssey project¹¹.

The commissioning, sponsoring or funding of an assembly also shapes the kinds of questions and outcomes it may be oriented towards:

- **Political actors** may look to commission processes that provide a public mandate for making a contested decision. Political commissioners may or may not have a preferred outcome from the process, but generally are not looking for processes that generate outputs challenging their existing agendas.
- **Companies** may look to commission processes that demonstrate corporate responsibility, that provide insights that can be fed into product design, that support industry self-regulation, or that feed into global regulatory discussions.
- **Independent funders / coalitions** may look to commission processes that build solidarity and collective power of communities, explore alternative visions of the future that may counter dominant narratives, or help set the agenda and framing¹² of global debates.

International actors often take a combination of these positions: taking account of political concerns, but also looking to demonstrate and advance a normative agenda around inclusive governance, and seeking to support elements of civil society and citizen participation in decision making.

Commissioners and organizers of citizens' assemblies differ as to whether they are complementary tools that are used alongside existing systems of representative democracy and that primarily get their power from this connection, or whether they are an alternative form of democratic engagement that should build power in their own right.

Forum

Global citizen deliberation on AI needs to interface with existing or emerging spaces of AI governance, or to have a mechanism to deliver change.

Docking is the process by which citizen deliberation interfaces with one or more existing institutional structures of power, influence or decision-making. Docking could involve a formal relationship, or an unofficial connection in which a deliberative process develops strategies to influence discussions or decisions in a particular forum^{4,13}.

The current AI governance landscape is complex and evolving. Against a backdrop of competition between big AI powers (both state, and corporate), and an explosion of AI ethics principles and voluntary codes^{14,15}, a range of global institutions have staked a claim to coordinate AI governance.

Veale et. al¹⁶ describe the wide range of different approaches to AI governance in play: from ethical codes and councils (including the widely agreed 2021 UNESCO Recommendations on ethics of AI⁶, and an explosion of voluntary codes since) to industry self-governance, technical standards (via bodies such as IEEE and ISO), international agreements (facilitated by groups such as the OECD, UNESCO, G7 and G20) and extra-territorial domestic regulation (such as the EU's AI Act). Dedicated groupings, such as the Global Partnership on AI (now integrated into OECD AI work) have sought to broaden the inclusion of global majority countries who have historically faced exclusion in the AI governance debate, and emphasis on AI has been growing in existing intergovernmental (e.g. International Telecommunications Union - ITU) and multi-stakeholder (e.g. Internet Governance Forum - IGF) fora. AI has also become a critical topic for thematic governance fora, from the International Labour Organisation to the World Health Organization, and in business-led spaces such as the World Economic Forum.

Although a number of these AI governance processes have been developing for years, breakthrough awareness of generative AI in 2022 accelerated the search for coordinated governance. The 2023 UK Bletchley Park Summit¹⁷ advanced the idea of a distinct regular (six-monthly) global policy forum specifically on AI governance (initially framed around 'frontier AI safety', though quickly broadened), and initiated the creation of an international scientific report on AI safety¹⁸. Drafts of the Global Digital Compact appear to build on this, with calls for a regular global policy dialogue on AI, backed by a scientific report¹⁹.

However, few of these spaces, if any, have established mechanisms for ongoing deliberative citizen input, and there are substantial fears that emerging AI governance spaces are dominated by technology firms and a narrow group of government and civil society actors. Although proposals for future AI governance arrangements increasingly highlight the importance of public participation, and inclusive global dialogue^{1,20,21}, few designs for how this can be achieved have been put forward.

In our interviews and workshops, we explored which fora a global citizens assembly could, or should, dock with. Insights, analysis and options from this are summarized in the table below.

If we want to have change and a more representative democratic governance model for AI then it's important to also look at where small spaces fit in: it's not just about the big conversations that are happening in the multilateral spaces.

— **Interviewee**

I think it's not possible for a private corporation to be conferred the kind of public legitimacy that decisions around governance should have, and that's why an institution like the UN is better suited. I don't think that the UN is the perfect model of representation, but I think it is a model of representation for all peoples that would allow it to serve as a better conduit here.

— **Interviewee**

Multi-stakeholder groups, civil society, and academia play a crucial role in shaping policy by putting people at the center of discussions. Their work is invaluable in addressing unanticipated issues, making these arenas the best spaces for meaningful dialogue.

— **Workshop participant**

In an ideal scenario, we need diverse voices—academics, civil society, and trade unions—at the table, as they often contribute positively to governance. However, when it comes to achieving quick, practical outcomes, multilateral forums and tech companies may have more immediate impact. The real challenge is ensuring these stakeholders actually listen and engage without external regulation, and navigating the complex power dynamics at play. The question remains whether civil society can truly influence the direction of technology in a timely and meaningful way.

— **Workshop participant**

Forum	Description	Approach / opportunity / assessment
United Nations UNGA UNDP UNICEF UNCTAD	<p>The UN General Assembly influences global norms.</p> <p>UNDP provides assistance to countries on policy development and implementation around AI.</p>	<p>Interface with UNGA and other parts of the UN System depends on Secretary General or Member State champions, and so is vulnerable to changes of SG, and power-politics between states ²². Shaping UNDP/WBG actions on AI could have impacts on LMICs, but less global impact.</p>
New UN AI Institutions AI Safety Institutes & Conference	<p>The Global Digital Compact draft calls for a global system of AI Safety Reports, institutes and annual multi-stakeholder dialogue, as well as a fund for AI, Emerging Technology & Sustainable Development ¹⁹.</p>	<p>Embed citizen participation into an emerging institutional structure under UN frameworks.</p> <p>Respond to, and influence, the agenda of regular global dialogues on AI.</p>
UN Specialized Agencies UNESCO ITU ILO FAO WHO WIPO etc.	<p>UNESCO produced the Recommendation on the Ethics of Artificial Intelligence in 2021, adopted by 193 countries ⁶.</p> <p>Specialized agencies address cross-cutting issues like health, agriculture, climate, intellectual property etc. The ITU in particular has been convening work on AI for Good.</p>	<p>While UNESCO has established recommendations, it has limited resources or power to support their implementation.</p> <p>Specialized agencies may have specific questions well suited to citizen assembly input, and/or the agenda-setting of an assembly might require engagement with executive agencies to see recommendations followed up.</p>
Internet Governance Forum	<p>The IGF is a long-standing multi-stakeholder forum addressing Internet issues. The Global Digital Compact calls for a stronger government and private sector IGF engagement in the IGF ¹⁹.</p>	<p>As an open multi-stakeholder process it could be easier for a GCA to put items on the agenda, but as a decision-shaping rather than making fora, influence on policy is unclear ²³.</p>
OECD & Global Partnership on AI	<p>The OECD has AI governance expertise, and is secretariat of the multi-stakeholder GPAI (Co-founded by France and Canada, and with 28 Member Countries + EU). Serbia is the incoming GPAI lead chair.</p>	<p>The role and influence of GPAI is currently unclear. There are relatively few LMIC members. Hosting a GCA on AI, with focus on speaking to nation states, companies and academia could offer a route to raise GPAI profile, and bring authority to the inputs of expert working groups.</p>

Forum	Description	Approach / opportunity / assessment
G20 / G7 / G77	<p>Various G- groups have put forward principles or agendas on AI coordination. summits hosting rotates between countries annually.</p> <p>Groups like the T20, B20 etc. seek to provide stakeholder input. Ongoing working groups sustain advocacy across years.</p>	<p>An individual G- host may be willing to commission a GCA on AI as part of their annual process.</p> <p>Resourcing a standing group on public engagement in AI may be difficult without substantial philanthropic backing. Decision making remains with states, and responds substantially to political pressures.</p>
Regional Bodies EU African Union ASEAN Organization of American States etc.	<p>The EU AI Act has established a leading role for the block. Belgium hosted a deliberative AI Dialogue as part of its 2024 EU Presidency ²⁴.</p> <p>The African Union Continental Artificial Intelligence (AI) Strategy has been endorsed by ICT and Communications Ministers of member states and awaits adoption.</p> <p>Other regional bodies are seeking to develop AI policy frameworks.</p>	<p>Influencing EU policy may be particularly important given the 'Brussels effect' ²⁵.</p> <p>Deliberation could inform and strengthen regional consensus on positions to take within real-politik of international negotiations. Interviewees noted the opportunity for engaging with regional bodies to support south-south cooperation, dialogue and positions in global negotiation. With the urgency assigned to developing AI governance frameworks, availing funds to the AU's Pan-African Parliament could create a standing citizens assembly for member states.</p>
Open Government Partnership	<p>The OGP is a voluntary association of states and sub-national members. A current strategic focus on digital governance addresses AI, and a OGP has a long-standing focus on participatory governance and deliberation.</p>	<p>OGP commitments/challenges could support member collaboration on global or local AI dialogues. OGP membership is not universal (specifically excluding China) and the OGP is best seen as champion of democratic governance rather than forum to influence.</p>
Standards bodies IEEE, ISO, etc.	<p>Standard setting plays an important role in the implementation of AI ethics, principles and values ¹⁶.</p>	<p>The need for more inclusive standardization processes is recognised, but influencing standards requires sustained engagement.</p>

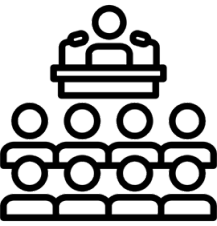
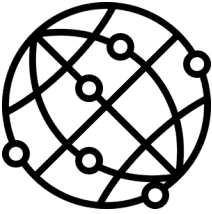
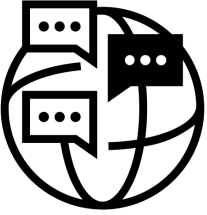
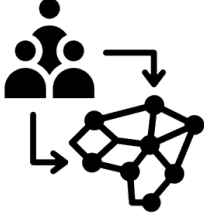

Forum	Description	Approach / opportunity / assessment
Scientific community International scientific report on AI International research programmes	<p>There are calls for an annual 'International Scientific Report on AI' either with UN mandate, or organized by the scientific community collating evidence on AI safety and impacts ^{18,26}.</p> <p>A number of research programmes seek to survey global public opinion on AI on a one-off or ongoing basis.</p>	<p>A public deliberative review of draft Scientific Reports on AI could both provide a clear route to inform publics about the latest evidence (inviting contributors to be expert witnesses to an assembly), and to direct future evidence review to address issues of global public concern (through assembly feedback).</p> <p>Public deliberation could also inform the design of global opinion surveys, ensuring questions reflect diverse public concerns.</p>
Industry Partnership on AI Individual Firms World Economic Forum	<p>Partnership on AI brings together industry, academia and non-profits, and has developed work around stakeholder engagement in AI development.</p> <p>Individual technology firms are also developing mechanisms for public input, in some cases, such as the Meta Oversight Board , endowing independent fora to support governance activities.</p> <p>The World Economic Forum (WEF) is an international advocacy NGO with influence on the business community. WEF has driven dialogue about '4th Industrial Revolution' including AI.</p>	<p>Industry stakeholders (either via partnerships, or individually) have been primarily interested in public input to 'fine tune' AI development, rather than setting overall goals for AI investment, or providing governance and oversight of existing AI systems. However, with the right incentive structures to ensure uptake of assembly recommendations (potentially through regulatory requirements) certain forms of deliberation could feed directly into industry.</p> <p>WEF could be a conduit to influence business debate on AI. The evidence gathered from an assembly could be communicated to WEF stakeholders.</p>
Civil Society Trade unions Consumer groups Social movements	<p>Civil society organizations play a key role in advocating for reforms, and supporting their constituencies to adapt to a world with AI.</p>	<p>Global deliberation can help make connections between national civil society networks, and build a more coherent and inclusive global agenda: supporting solidarity and shared action on AI governance.</p>

Form

Global citizen deliberation on AI needs to be designed for inclusion and impact.

The democratic innovation field is booming, although democratic culture itself faces significant threats. Recent years have seen substantial experimentation with both global deliberation, and public deliberation on AI. Many of the components for a global citizens' assembly on AI are now tried-and-tested. The question is how these might be combined so that form follows function, and an assembly has the best chance to both influence existing decision making fora, and catalyze wide impacts.

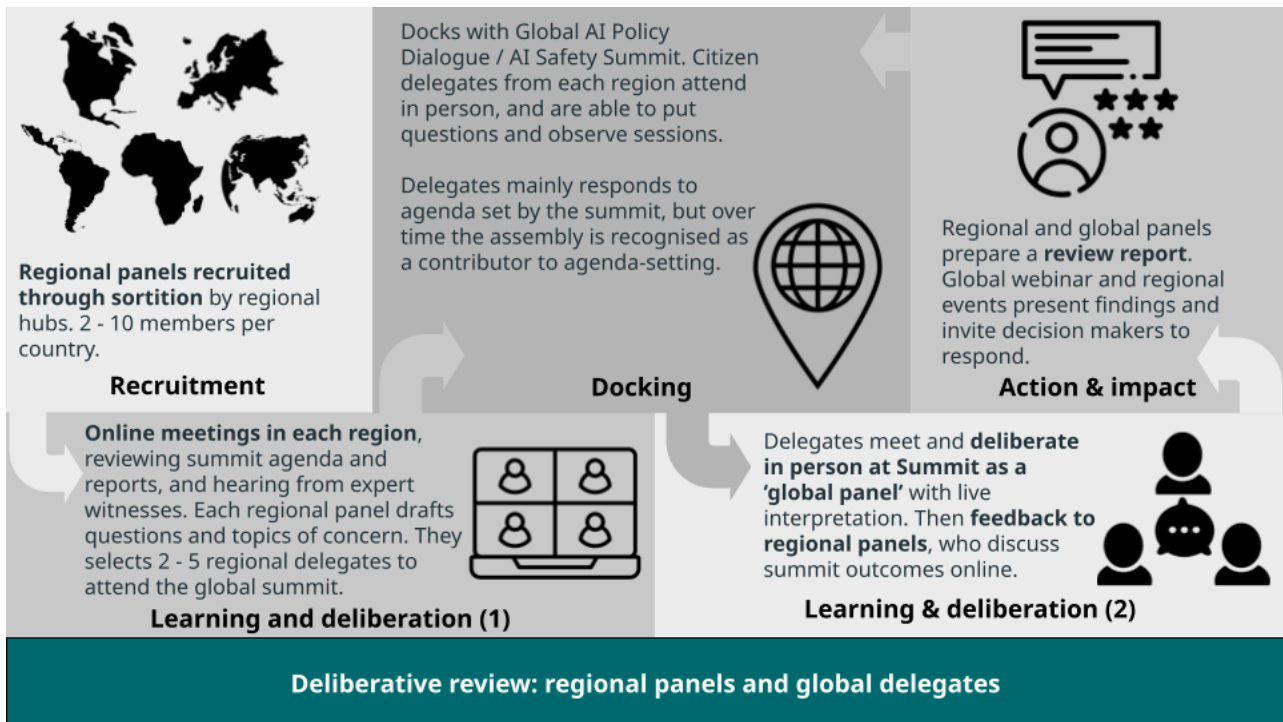
Based on desk research, workshops and interviews, we developed five stylised options for a global citizens' assembly on AI, each one placing emphasis on different areas of focus, fora, and assembly design. Through subsequent workshops and interviews, we refined these into the four assembly options, and one commissioning structure presented below.

 <p>1: Deliberative Review Providing public inputs to scientific reports and/or global summits, through pre-event regional panels and public delegates/ rapporteurs</p>	 <p>2: Global assembly Independent core sortition-selected transnational assembly, and self-organized community assemblies. Consensus reports and 'cultural wave' to shape AI narratives and policy.</p>
 <p>3: Distributed dialogue A network of partners organize local dialogues following a common framework. Global report aggregates findings, and dialogue partners engage to influence change locally.</p>	 <p>4: Collective intelligence Using digital platforms with open and intentional participant recruitment to support dialogue at scale, and provide inputs to industry and policy.</p>
 <p>5: Commissioning body for public dialogue An expert group that is resourced to commission public deliberation or dialogue on AI topics from other existing deliberation processes. For example, funding and providing expert input for environmental impacts of AI to be addressed by the standing citizens assembly on climate change.</p>	

While set out as discrete forms, these options are not necessarily mutually exclusive. Components of each may be combined. At the same time, there are real choices here. For example, some interviewees suggested it would be impossible to build powerful citizen narratives on the future of AI without total independence from technology industry stakeholders and funding, whereas for others, working with industry was seen as the only viable means to influence change over the short to medium term.

In the following pages we also consider a number of cross-cutting design issues for any assembly to address, such as recruitment, governance, transnational deliberation, and the use of technology as an enabler of dialogue.

Option 1. Deliberative review: regional panels and global delegates



Global summits act as a key focal point for international negotiation: as sites of both decision-shaping and decision-making. They also have a rhythm of preparation, meeting and follow-up that can be leveraged to structure the learning and deliberation components of a citizens' assembly²⁷. Docking into an existing summit can be official when supported by the summit organizers, or unofficial, by using public access components of a summit (webcasts, open meetings, reports), or through engaging with an open multi-stakeholder fora like the Internet Governance Forum.

Picture:

- **A series of regional panels** created through sortition with between 2 and 10 representatives from each participating country, meeting online over 4 to 6 weeks to learn about an upcoming AI Summit; review background materials; hear from expert speakers; and identify questions and topics of concern to the panel. Each regional panel would be organized by a local coalition of partners, following a common approach.
- **Nominated citizen rapporteurs from each region attend the AI Summit/Policy Dialogue** and (a) have the opportunity to put forward questions decided by their regional panel in discussions/debates; (b) act as observers and report back to their regional panels; and (c) hold facilitated deliberation (with live interpretation) between regional delegates to feed into the post-summit review.

- **After the summit**, the panels reconvene and produce regional and global **review reports** that (a) summarizes the debate for a wider group of citizens; and (b) provides a judgment on whether the summit process and outcomes are delivering against citizen priorities.
- **Summit hosts** are invited to consider the review reports in planning future editions of the summit.

A deliberative review model could also be applied to the preparation of a global scientific report on AI, with citizen panels hearing evidence from scientific experts, posing questions, and offering commentary on the final report.

Advantages	Challenges
<ul style="list-style-type: none"> ■ Enables regional blocks to develop distinctive voice in the AI debate, addressing current global north biases. ■ Some logistical, interpretation and translation costs are borne by the anchor Summit. ■ Ease. In essence this fits with the established approaches to tech governance. 	<ul style="list-style-type: none"> ■ Summit agendas are often finalized last-minute, creating practical challenges. ■ Expense of bringing rapporteurs physically together at a summit. ■ Strongest version requires buy-in from Summit hosts (may be tricky to sustain if rotating chair). ■ Often limited scope for significant impact due to capture by 'status quo' mindsets and unclear connections / commitments by powerholders to enact proposals.

Inspirations: UK Peoples' Panel on AI; Oregon Citizens' Initiative Review.

The Peoples' Panel on AI²⁷ convened a sortition-selected group of 11 people to attend sessions of the AI Fringe conference hosted alongside the 2023 Bletchley Park AI Safety Summit, and to review livestream recordings and output documents from the summit. The panel produced a review during the summit week, reporting back to industry, academic, government and civil society stakeholders with recommendations for future action on AI governance. The design of the People's Panel on AI built on the Oregon Citizens' Initiative Review²⁸, in which citizen groups produce an assessment of proposed legislation, used to educate other citizens about their choices in a local referendum. A review does not have to reach consensus, and can report on differences of opinion amongst the deliberating group.

Thinking about recruitment

Sortition (or civic lottery) is a central concept for citizen assembly recruitment. Sortition combines a principle of random selection, with the goal of securing a demographically representative or diverse group of participants. Well-designed sortition addresses self-selection biases, and maximizes the chance that any individual member of the population *could* be invited to participate in an assembly. This is often central to the legitimacy claims of a citizen assembly to represent informed views of the general public.

- The People’s Panel on AI used an opt-in e-mail list built by Sortition Foundation through previous recruitment processes to identify people available to take part in a short-notice deliberation, and then ran a stratified lottery to select 12 diverse people from the 490 expressions of interest.
- The Belgian Citizens Dialogue on AI sent mail to 16,200 randomly selected postal addresses in the country, with additional sampling of German speaking areas. A stratified lottery was then used to select 60 participants from respondents²⁴.
- The Global Assembly on Climate and Ecological Emergency used geographic and demographic sortition, selecting 100 points on earth, and then identifying local partner organizations within 200km of each point, before working with these organizations to recruit 4 to 6 candidate participants, of which one was then selected to create demographic diversity across the whole 100-person assembly^{13,29}.

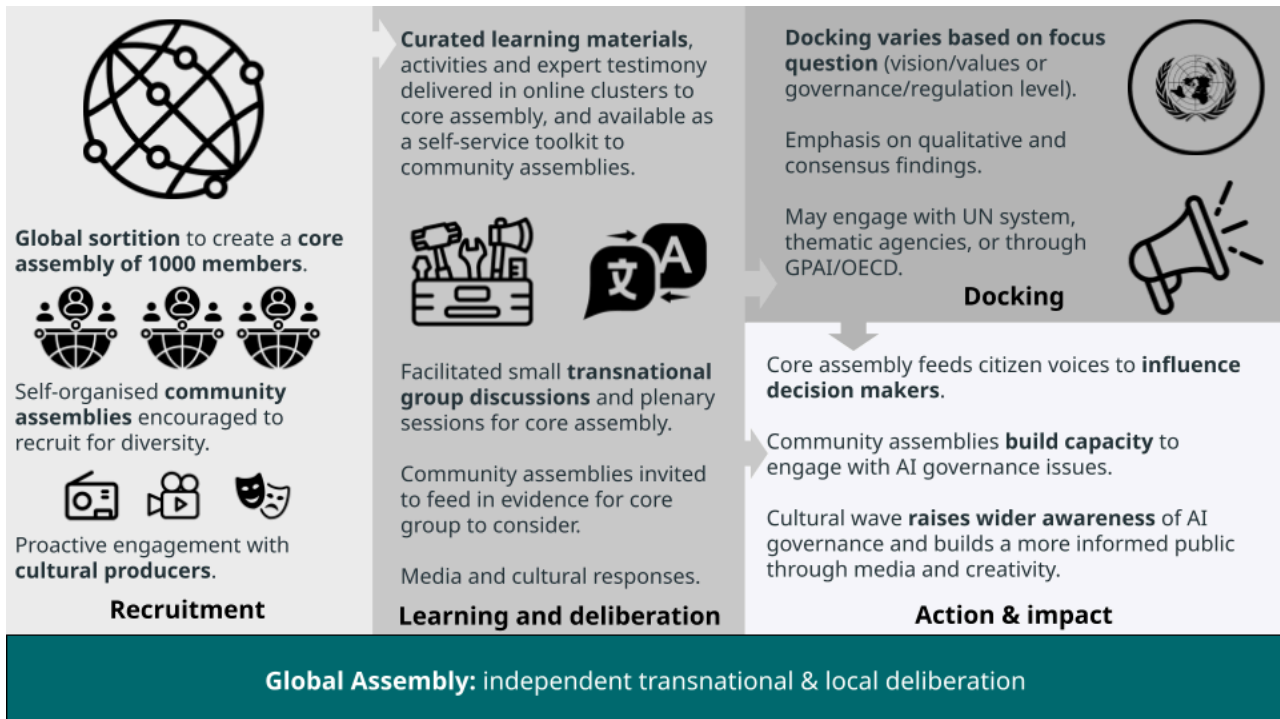
Sortition algorithms can be weighted to over-sample particular demographics, or to increase diversity of perspectives: for example, increasing representation of historically excluded groups, or to ensure certain demographics are represented by at least two members. In some cases, particularly to support the involvement of children and young people who may find different kinds of conversations and exercises more engaging, it might be appropriate to create parallel assemblies or demographic sub-groups.

In a global deliberation, recruitment practices may need to be adapted in different geographies to recognise practical and cultural constraints, ranging from the absence of an established list of addresses to select from, or safety concerns with recruiting, to cultural differences in approaching potential participants via heads of households, or only with consent of community leaders.

It's not clear to me that you necessarily just want random draws. You may actually want to over-represent some groups, however defined, whether it's locally, nationally, however you want to define those groups to make sure that, because they've been excluded in the past, that they actually get more weight in the conversations right now.

— Interviewee

Option 2. Global assembly: independent transnational & local deliberation



An independent assembly would have its own funding endowment and resources to organize a large-scale global sortition-selected assembly with members coming together from different countries, as well as to provide support for community-based assemblies, and to catalyze a cultural wave of activities to popularize the learning components and recommendations of the assembly.

Picture:

- **A core assembly of 300-1000 members** selected through global sortition and meeting regularly online in linguistic or regional clusters, and in transnational sessions, to consider expert learning inputs, testimony and insights from members. By maintaining relationships with key UN Agencies, the assembly would consider questions covering general AI governance, and specific topics, such as AI impacts on health, work or climate.
- **A community assembly toolkit** would use the expert inputs and questions presented to the core assembly as materials to support locally-organized community assemblies. Community assemblies would be invited to provide evidence and reports that would be summarized for the core assembly to consider.
- **Catalyzing a cultural wave** that supports media partnerships/coverage of the core and community assemblies, and that engages with artists and creative practitioners to create and disseminate work in response to themes and questions considered by the assembly.

The core assembly would rely on trained facilitators and interpreters. Each iteration (e.g. theme or question) of the assembly would produce a report, and nominate a number of members to present this at a relevant forum (e.g. ITU AI For Good summit; UN meeting). Community assemblies and the cultural wave should contribute to developing broader understanding of AI, more citizen-centered narratives about our AI future, and cross-connected community-level initiatives on AI governance³⁰.

Advantages	Limitations
<ul style="list-style-type: none"> ■ Places emphasis on transnational deliberation. ■ Links mini-publics with macro-publics, fostering both policy impact and public education on AI. ■ Working with local partners to support assembly member participation and community assemblies can promote inclusion and build civil society capacity. 	<ul style="list-style-type: none"> ■ Weighting by population / geography in ways that may be appropriate for discussing climate issues, may be less relevant to capturing diversity of communities affected by digital environments. ■ Significant budget requirement (particularly if seeking endowment to build independent/ongoing infrastructure).

Inspirations: Global Citizens Assembly; European transnational assemblies.

The Global Citizen Assembly on Climate and Ecological Crisis²⁹ took place in 2021, bringing together a core assembly of 100 participants for transnational dialogue, and supporting community assemblies involving 1000s of participants. The project evaluation notes that the model that was adopted sought to “address structural constraints of convening global citizen deliberation in a highly unequal world” but that “Not all constraints can be overcome”, highlighting the importance of clear governance¹³.

A number of transnational citizens’ assemblies have also taken place within the European Union, generating considerable learning on facilitating deliberation across languages³¹.

The theory of change for an assembly cannot rely on action from the commissioners alone. Experience from the last decade of research into citizen deliberation tells us that even when politicians say they will listen and implement recommendations from the public, in reality they often don’t. We also need mechanisms for movement building, and for people to act independently based on assembly outcomes: power literate theories of change.

— Interviewee

Transnational deliberation & aggregating across borders

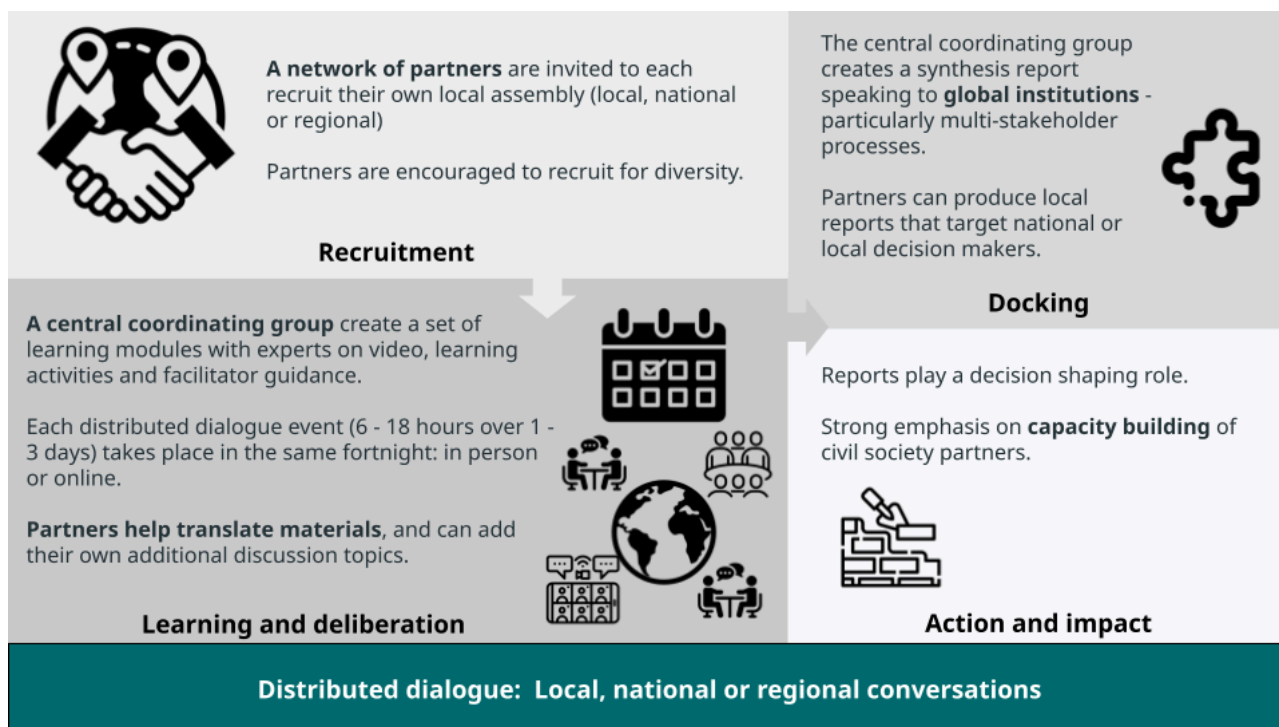
A key design consideration for any global deliberation on AI is how to foster conversations across borders.

Practical issues, such as differences of language, culture and either working digitally across time-zones, or managing the costs of bringing large groups of people together in person make transnational deliberation challenging. Power dynamics also have to be considered, in particular to ensure majority-world participants experience global deliberation as a space for active voice, rather than solely engaging with deliberation in the mode of students or learners¹³.

However, to explore the extent of global consensus on responses to AI, and to make sure the proposals arising from deliberations take into account lived experience from across the world, it is vital to have mechanisms that bridge dialogue in different locations.

- **Common background materials** can be translated into different languages, tailored to local cultures, and developed with balanced expert inputs from all regions.
- **Transnational clusters** might bring together a small subset of different language groups to discuss common issues, supported by live interpreters.
- **Regional exchanges** might invite representatives from one regional cluster to present their discussions and findings to other regional clusters.
- **Digital aggregation** may make use of large language models to translate and summarize discussions between different local discussion groups.
- **Local facilitators, or participant-selected rapporteurs** may meet together to draft a synthesis report that reflects different local discussions.

Option 3. Distributed dialogue: Local, national or regional conversations



In a distributed dialogue, a central coordinating group would identify questions, prepare learning materials and create facilitation guides, and then invite local organizations to host their own local, national or regional deliberative events.

Picture:

- **A dialogue toolkit for a selected AI governance topic**, consisting of background documents, recorded expert testimony (available to translate to local languages), recommended learning activities, facilitators guide, and guidance on recruiting a demographically diverse participant group.
- **A network of local partners**, each organizing their own dialogue days (1 to 3 days) with a mix of grant-funding from the central organizers and self-funded partners using independent resources. National hubs and local partners would be offered the opportunity to include their own tailored questions within the dialogue, responding to local concerns.
- **A focus fortnight** when partners are encouraged to run their distributed dialogue events, and report findings back to the central coordinating team. Reporting back may take the form of structured surveys, event transcripts, and summary statements or recommendations from each site.
- **A global report** presented to key stakeholders, and the option of partners producing local, national and regional reports and advocacy events.

Advantages

- Adaptable to local contexts, and able to influence at local, national and regional as well as global level.
- Contributes to network-building and education amongst civil society hosts of each dialogues

Limitations

- Lack of cross-border / transnational deliberation;
- Difficulty of quality-controlling deliberation at different sites, and risk of capture by partner organizations;
- Risks that global agendas distract from local discussions.

Distributed dialogue offers a valuable opportunity for localized action. There are many things communities or regions can do to shape AI based on their local norms and values, which can be reflected in AI deployments, such as through fine-tuning models and shaping system prompts.

— Interviewee

Inspirations: We The Internet; World Wide Views; Global Data Barometer regional hubs

We the Internet convened a global citizen dialogue in October 2020, engaging over 5000 members of the public through more than 70 dialogues organized by a network of partners in different countries. Each dialogue followed a common structure of nine modules, involving scene-setting video, learning inputs, group discussion and individual tasks. Local dialogues could also add a country-specific module. Quantitative and quality results from local dialogues were aggregated and reported alongside outputs from a stakeholder dialogue³².

World Wide Views is a methodology for global consultation, first used in 2009, and most recently in 2015 with almost 10,000 citizens across 76 countries. The method combines a common structure for a consultation day, run in parallel across many different locations. Following common information videos, participants deliberate in small groups, and then vote via the World Wide Views website, with results aggregated and reported to policy makers³³.

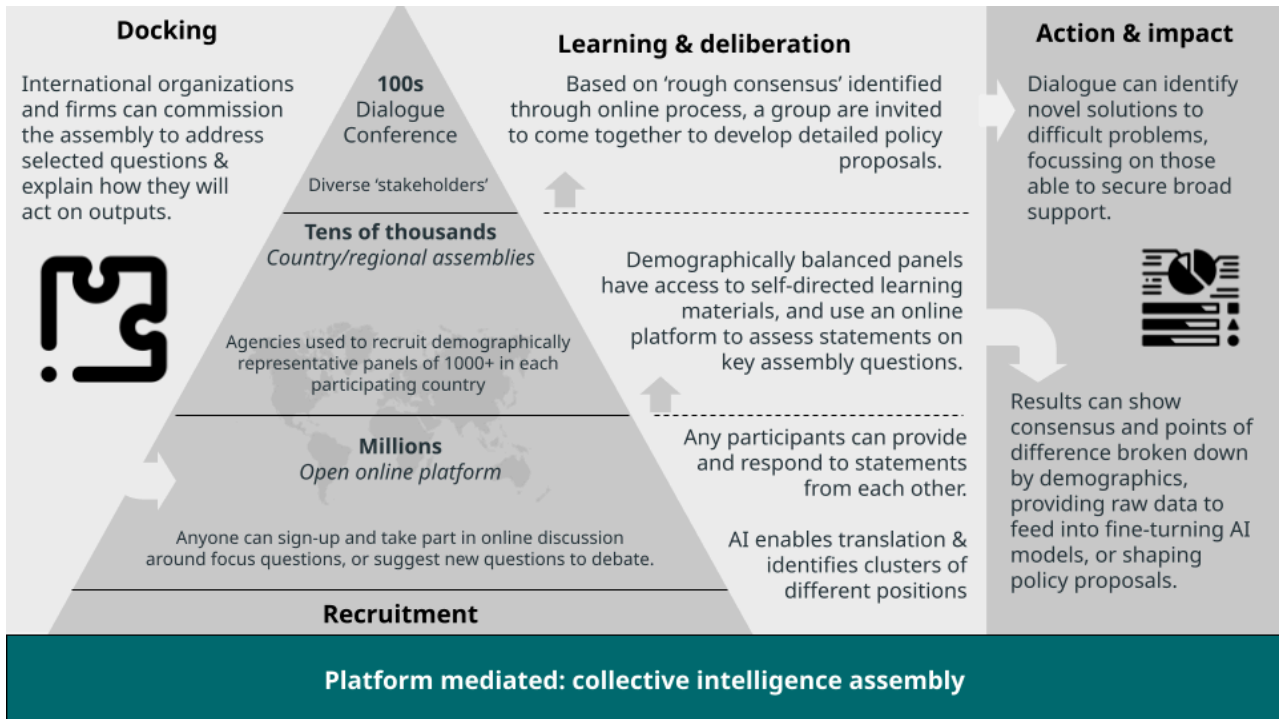
The Global Data Barometer³⁴ is an expert survey-based study on data for the public good. Whilst not a form of public dialogue, it is organized through regional hubs that play a key role in governing the project by helping to recruit country-level participants, shape global research questions and tailor the presentation of findings to local culture and contexts. The host organization for the Global Data Barometer team is based in the majority world.

Assembly Governance and Expert Inputs

Governance is vital to ensuring the quality and legitimacy of a global citizen deliberation on AI. Most of the deliberations used as inspiration for this report have steering, oversight or governance groups that include a range of respected experts. Some may also have scientific councils or groups that were involved in selecting or reviewing learning materials and experts presenting to participants. It will be particularly important for a global citizen deliberation on AI to have *global* governance, and to ensure governance and expert testimony is not dominated by global north participants¹³.

Although domination by the global north may not be intentional, it can be the common outcome when histories of political and socio-economic privileges are not addressed. For instance, when some stakeholders in a deliberation have greater proximity to funders, this can affect team power dynamics. The selection, and presentation, of expert inputs must also take into account historical oppression that impacts on the way in which participants show up within global spaces, and how they engage when their knowledge systems, languages and cultures have historically been subject to what Frantz Fanon calls cognitive imperialism³⁵. Scaling from a local to a global dialogue on AI is likely to involve not just transformation of scale, but also critically thinking about framing, content and processes that respond to current and historical injustice.

Option 4. Collective intelligence: digital dialogue at scale



Collective intelligence approaches make use of technology for “surfacing and combining group beliefs, goals, values, and preferences” at scale³⁶, relying upon the diversity of participants, and the design of the platform, to yield outcomes that represent either unified, or multiple, points of consensus around policy options or practical action^{37,38}. In this option, we emphasize how technology platforms, supported by use of AI, could facilitate large-scale public inputs to AI governance, and the potential for processes to dock directly with decision making within AI firms.

Picture:

- **A series of framing questions**, both those commissioned by different decision-making institutions (including AI companies, and international organizations), and a selection proposed through a bottom-up process.
- **An open online platform accessible to anyone**, making use of AI to translate contributions and find clusters of opinions in response to prompt questions. Participants interact by indicating agreement or disagreement with propositional statements, or offering their own suggestions (This might draw on the pol.is platform³⁹, designed to prioritize ‘bridging statements’ that build consensus across different clusters of diverse views).
- **A series of country or regional panels** recruited to be demographically representative by recruitment agencies. Panelists are paid an honorarium to take part, and asked to consider background materials, before responding to, or providing, statements. Statements from this demographically representative group can also be fed into the open process.

- **A dialogue conference** that brings together stakeholders (including organizations that commissioned the process, and individuals representing different positions in the discussion) to discuss potential ways forward.

Data about public views on AI could also be ‘tapped’ from the process at a number of stages: from looking at the clusters of opinion gathered in an open dialogue, to exploring disaggregated data on how different demographics feel about certain issues, through to exploring transcripts from a dialogue conference. Data could feed directly into fine-tuning of AI models, or to making AI models more accurately reflect cultural differences around the world. These outputs could be made available as a public good for use by open and public AI model developers⁴⁰, as well as by commercial model providers.

Advantages	Limitations
<ul style="list-style-type: none"> ■ Highly scalable. ■ Enables deliberation participants to have some agenda-setting power. ■ Supports a degree of direct cross-language engagement (via machine translation). 	<ul style="list-style-type: none"> ■ Requires high levels of connectivity and literacy. Needs to address cultural variations in how people communicate online. ■ Industry is unlikely to implement uncomfortable recommendations without being pushed to by external forces (consumer pressure, regulatory action etc.).

One of the benefits of collective intelligence and deliberative forums is the cross pollination and diverse ideas and people experiences, which also builds legitimacy.
 — **Workshop participant**

Inspirations: vTaiwan; Meta Community Forum; Collective Intelligence Project Constitutional AI project; Open AI Democratic Inputs to AI

vTaiwan is described as “a decentralized open consultation process that combines online and offline interactions, bringing together Taiwan’s citizens and government to deliberate on national issues... [A] model for People-Public-Private Partnerships (PPPP), involving government ministries, elected representatives, scholars, experts, business leaders, civil society organizations, and citizens in crafting digital legislation.”⁴¹. vTaiwan has made extensive use of the pol.is platform, combined with recruitment outreach and face-to-face engagement. It was also used in one of the 2023 OpenAI funded **‘Democratic Inputs to AI’** experiments that sought to gather public input addressing

questions such as whether and when AI systems should provide medical or legal advice^{42,43}.

A number of other technology firms have also commissioned or carried out projects to gather greater public input to AI development or governance. **The 2024 Meta Community Forum**, run in partnership with the Stanford Deliberative Democracy Lab and the Behavioural Insights Team used deliberative polling to engage 1,500 people across the USA, Brazil, Germany and Spain in digitally facilitated discussions on how AI chatbots should behave⁴⁴. In late 2023, **Anthropic worked with the Collective Intelligence Project** to gather inputs from 1,000 US citizens through the pol.is platform, creating and expressing agreement or disagreement with statements that were then used to fine-tune a large language model (LLM) chatbot⁴⁵.

These industry-commissioned processes have tended to draw on market research recruitment agencies, and to prioritize scale of participation over depth or duration of dialogue. To enable light-touch participation, these processes tend not to include a substantial expert-input stage, and focus on outputs (statements, polling questions etc.) that can reported on in aggregate, or analyzed to understand the views of different sub-populations (e.g. to look at differences in support for certain governance measures between countries, age groups or other groups exposed to dialogue on AI, vs. those who have not had space to deliberate).

AI as an assembly delivery tool

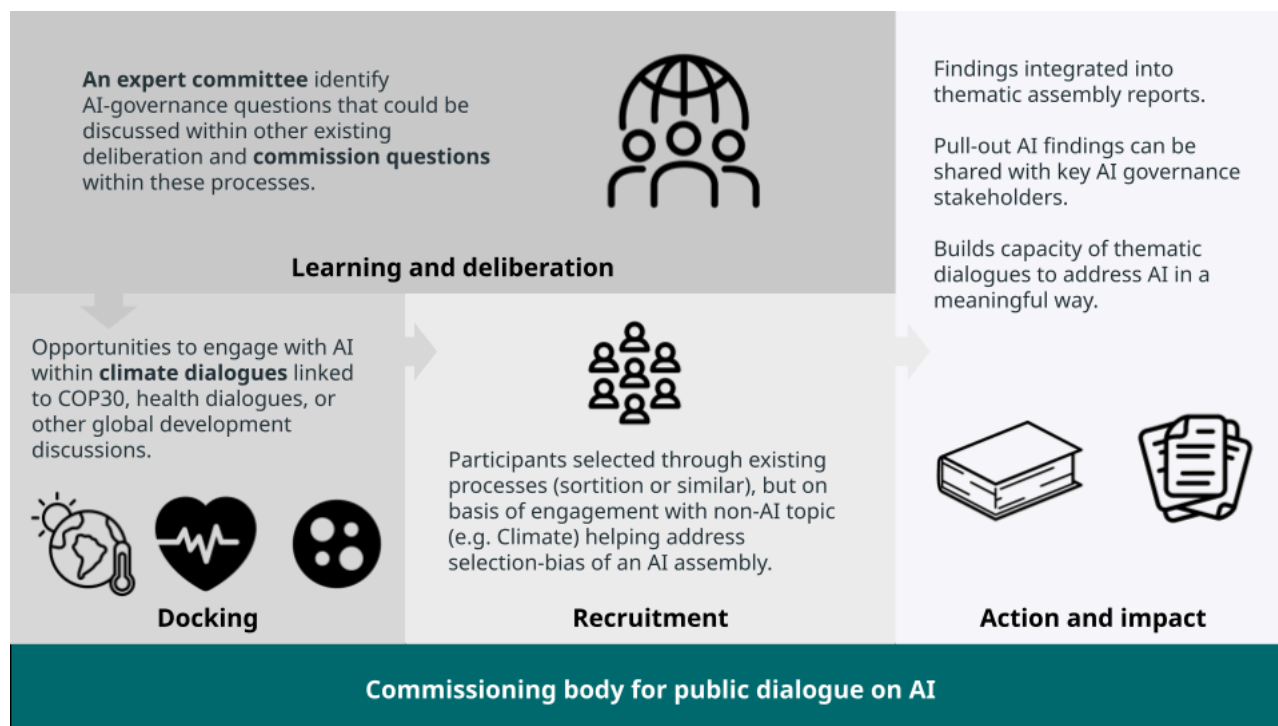
While we have focussed in this paper on how global citizen deliberation could be part of governing AI, there is also significant interest in how AI might be used to support citizen deliberation, both at local and global levels. McKinney⁴⁶ highlights the potential of AI tools to support deliberation:

- **recruitment and planning**, including clustering inputs to support agenda setting
- **learning**, particularly through simplifying and summarizing materials, assisting with question generation, and acting as Q&A systems
- **deliberation and decision-making** including automated facilitation; measuring deliberative quality; playing devil's advocate; aggregating across deliberations; and generating consensus statements
- **follow-up**, including communicating outputs with the public

Deliberation designers have also expressed an interest in how AI tools could support transcription of discussions; provide live translation or interpretation for multilingual processes; and provide live synthesis of qualitative materials across parallel discussions.

While AI technology has significant promise, in the context of global dialogue a number of limitations are important to keep in mind. The organizers of the Belgian Citizen Panel on AI²⁴ found transcription quality was highly dependent on use of good microphones, and AI models' capabilities and performance vary across languages depending on their training data. There is currently inadequate evidence on how model biases may affect machine-driven summarisation of discussions⁴⁷.

Option 5. Commissioning body for public dialogue on AI



Many of the most important global governance questions about artificial intelligence may be best discussed in the context of wider global challenges (e.g. discussions of climate, trade or migration), or in application-focussed discussions (e.g. discussion of rules on self-driving vehicles, or autonomous weapons systems). An expert committee on public dialogue on AI could commission, fund or support AI components within other public dialogue processes.

Picture:

- **An expert committee on public dialogue on AI** which works with global and regional assemblies on topics such as climate, health and international cooperation to frame questions (e.g. When and how should AI be used to address climate change? What should be done to address climate impacts of AI?), curate expert inputs, and translate outputs into AI policy fora.
- **Participants in these existing assemblies are provided expert inputs** on AI, and given space to discuss these questions. Because they were not recruited directly to an AI-focussed discussion, some of the self-selection biases around interest in the AI topic that can remain even with effective sortition are addressed.
- **Outputs feed both into thematic assembly agendas, and to AI-related processes.** For example, the committee may commission discussions that inform a chapter of the International Scientific Report on AI, or feed into the work of UN Institutions developing financing for AI-related development.

Advantages

- Reduces self-selection bias in AI-focussed deliberation.
- Grounds discussions of AI within wider issues.
- Provides access to expertise on designing effective deliberation on AI.

Limitations

- Commissioning body needs adequate resources to bring together expertise and fund processes.
- Relies on other assemblies existing *and* having space on their agenda for AI issues.

Inspirations: Sciencewise; KNOCA

We were not able to locate examples of bodies that commission specific questions *within* other public deliberation. However, there are some examples a model could build on.

The Sciencewise programme, funded by UK Research and Innovation, provides support to government bodies in the United Kingdom to commission deliberative public dialogue on science topics, including data governance and AI. The support they offer includes providing up to 50% co-funding, expert advice and guidance⁴⁸.

The Knowledge Network on Climate Assemblies (KNOCA) works “to improve the commissioning, design, implementation and impact of climate assemblies, using evidence, knowledge exchange and dialogue.” Members learn from each other on how to improve the design and delivery of citizens’ assemblies focussed on climate change⁴⁹.

The costs of deliberation

Existing experiments in global deliberation such as We The Internet or the Global Assembly on Climate and Ecological Crisis have been delivered with budgets of around \$1m. Moving beyond these pilots and ensuring partners and citizen participants are properly paid for their time and effort may require an uplift to these budgets, but global deliberation is a comparatively marginal sum compared to the \$millions and even \$billions being invested in training data and compute for modern AI models.

Delivering change

Any instance of global citizen deliberation on AI needs to have a clearly articulated theory of change.

Battles for the future of AI will be hard fought. There are huge vested interests at stake, with some of the most powerful companies and countries across the globe approaching AI as a strategic interest and key tool in how they build power and shape the world. Many of the governance issues the rapid development of AI has raised currently fall into an institutional void⁵⁰, where questions of jurisdiction and enforcement remain unsettled. Against this backdrop, notions that citizen deliberation can shape AI can seem naive.

Fortunately, recent decades have generated a wealth of learning on what it takes for deliberative processes to lead to action. With a robust issue-by-issue analysis of where power currently lies, awareness of the rapidly evolving institutional landscape, and serious strategies for influencing and rebalancing power, deliberative practice on AI can deliver change.

Building on an in-depth design process, the Global Citizens' Assembly for People and Planet launching in September 2024 alongside the UN Summit of the Future articulates five desired impacts: institutional action; citizen action; solidarity (local and global); learning at scale; and inclusion.



Figure: Global Citizens' Assembly for People and Planet Impacts (ISWE Foundation)

Plans to deliver any specific option for global citizen deliberation on AI should adopt a comparable impact framework. With potential impacts identified, the theory of change for a deliberation will determine which of the impacts are prioritized. To a large extent, solidarity, learning, inclusion and even citizen actions lie within the locus of control of the deliberation design team and the available operating constraints such as budget. Influencing institutional actions, which is often the impact most people care about, is likely to prove most difficult. Whether shaping global regulatory frameworks or robustly influencing how big AI Labs choose to develop and deploy their technologies, these actions often fall outside the direct control of deliberation organizers.

Building on the overall design decisions for a deliberation outlined above, the delivery of a deliberation will involve a range of tactical choices, combining ‘insider’ and ‘outsider’ tactics to achieve influence.

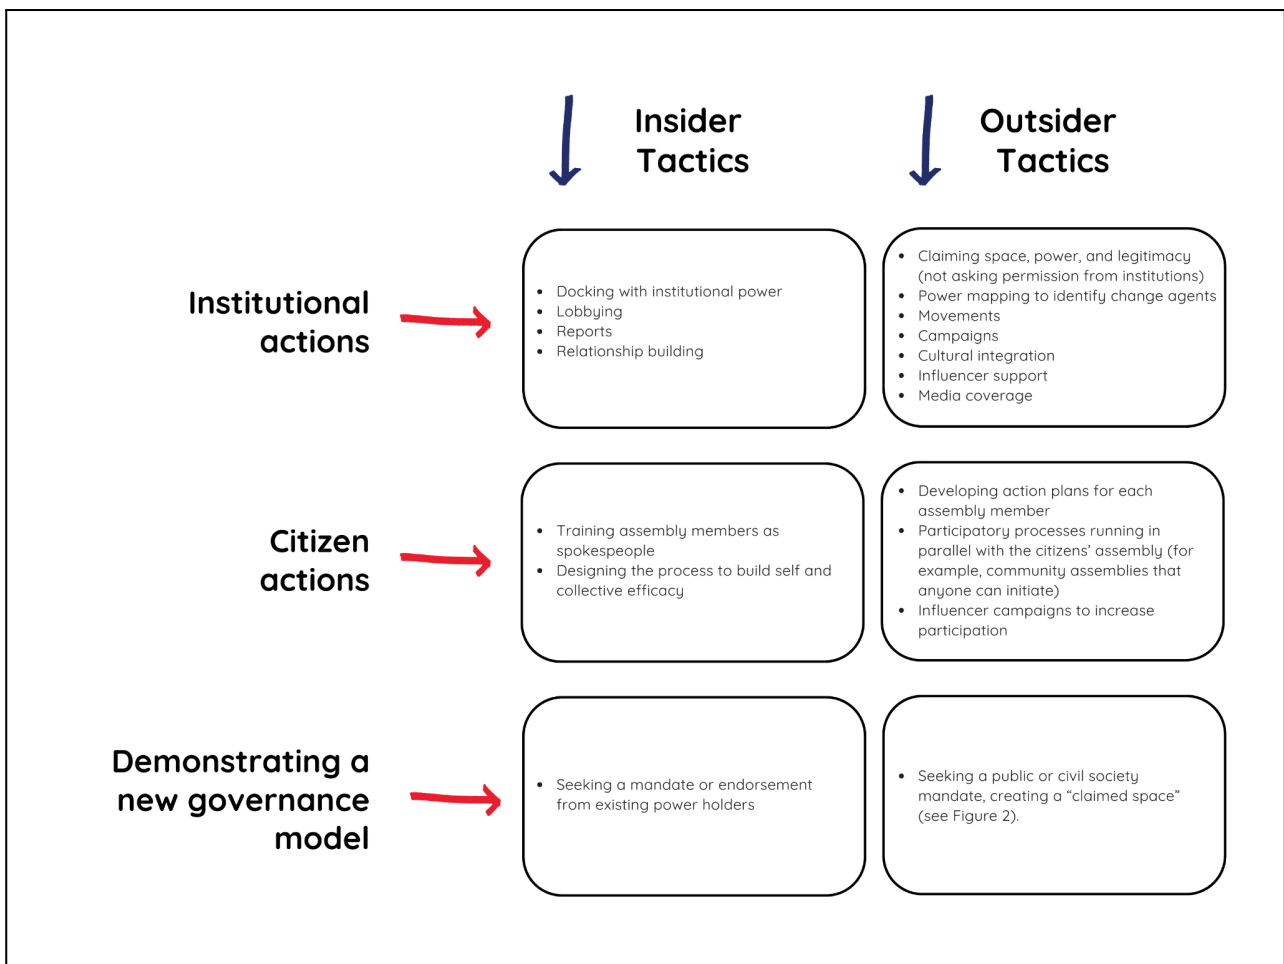


Figure: Insider and outsider theories of change⁵¹

Insider tactics include getting mandates from formal power structures, lobbying decision-makers, writing reports or relationship building. This is how the vast majority of citizen deliberation projects have been configured in recent years. By contrast, outsider tactics may draw on traditions of citizen deliberation such as the Senegalese people's

assemblies and Brazilian National Public Policy Conferences that have a lower emphasis on demographic representation, but involve large numbers of citizens to agree policies and then campaign for them from outside formal institutions. Efforts to establish community-led permanent citizens' assemblies also offer an approach based on building alternative models of governance, and new institutions that 'claim space' rather than wait for it to be granted.

The interaction between public deliberation on AI and social movements is also critical to explore. In the past, when the world has confronted powerful and contested new technologies, such as nuclear weapons, decades of protest and civil disobedience have played an important part in shifting national and global policy. The educational and activational role of social movements are often overlooked, but are likely to be critically important for AI, given often low public awareness and the highly contested nature of how AI should develop. In the climate field, citizens' assemblies have been the demand of social movements: and assemblies have also sought to feed into social movement activity⁵². How this interaction will play out around AI is unclear, but important to consider.

As we will outline below, one key aspect of delivering change through citizen deliberation is ongoing evaluation and iteration of different designs, strategies and tactics. However, at the early design stage, the following questions may help to sharpen any initial approach:

- **How can the assembly effectively navigate and influence different levels of power (local, national, global) using both insider and outsider tactics?**
- **What strategies will be employed to ensure both visible and hidden forms of power are addressed in the deliberation process?**
- **How can marginalized voices claim or create spaces within the assembly to challenge institutional power and shape AI governance?**

From concepts to actions: next steps

Bold steps are needed to embed citizen voices from across the globe in defining the future of AI.

During workshops and interviews, participants called for an ambitious programme of action that can meet the challenges of the current moment of AI development and deployment. As the examples through this report show, efforts to expand deliberative governance of AI are already underway, but they remain ad-hoc and do not yet add up to a sustainable and transformative programme of work to center the voices of affected communities in AI governance.

We believe that, regardless of the exact combination of options for global deliberation on AI pursued, there are four key areas for action:

1. Develop institutional docking – emerging AI governance institutions must commit to engage with processes of global citizen deliberation

As new institutions and structures are established for global governance of AI, aspirational language on democratization, public participation or talk of AI governance ‘by all’ must be accompanied by concrete plans to create docking points that will bring diverse, inclusive, informed public inputs into the heart of discussions and decision-making. Securing this calls for government champions who can ensure space is created for public input as new institutional arrangements are negotiated, as well as for industry and civil society to recognise and support the complementary role of citizen deliberation within multi-stakeholder governance.

2. Invest in deliberation at scale – funders should pool resources for a range of independently governed experiments in global deliberation on AI

As we have set out above, there are many options that can be taken to bring deliberative public inputs into AI governance, and many different decisions that public participation should influence. Compared to vast investments in AI development, AI safety, and wider AI governance efforts, citizen participation programmes have received negligible funding to date. Funding should be ring-fenced to support deliberative public engagement with AI governance, including standing infrastructures that can support a range of deliberative processes such as citizens’ assemblies.

Building on existing evidence of what works, and seeking to address a range of focus topics and fora, we need to see a set of strategic experiments that can iteratively refine effective approaches for meaningful and globally inclusive citizen involvement in AI governance.

To move beyond ad-hoc pilots, and to secure the independent governance of processes while allowing corporate funding, consideration should be given to the creation of a sector-wide

endowment for global deliberative inputs to AI, drawing on the precedent of the \$250m+ firm-specific Meta Oversight Board endowment.

3. Create infrastructure for learning – practitioners and researchers need opportunities to learn, develop and evaluate work together

In developing this report we have surfaced a range of overlapping communities working to promote public inputs to AI policy and practice, as well as public participation communities with deep experience of different approaches to citizen deliberation. There is an ongoing need to build bridges between these communities, connect learning and support robust critical and constructive evaluation of different approaches to deliberation on AI.

Constructive and formative independent evaluation has been key to the development of the climate assembly field, and will be critical for AI dialogues. Evaluation should bring an emphasis on the extent to which each initiative has, and achieves progress against, clear theories of change, and the extent to which they deliver inclusive global participation with full representation of the majority world.

4. Connect with movement building – local knowledge, confidence and movements will be critical to delivering change

In parallel to institutionally linked citizen deliberation, progressive change that centers global citizen needs also requires a focus on how deliberative processes positively contribute to the capacity of individual citizens, social movements and grassroots organizations to confidently engage with the topic of AI governance. Deliberative processes should be approached as a powerful tool in developing global skills and knowledge on AI in ways that help empower citizens to shape the future: whether in their own workplace and neighborhood, or on the global stage.

Method & acknowledgements

This paper draws upon an extended design lab process coordinated by **CONNECTED BY DATA** between May and August 2024. Drawing on a rapid literature review on citizen deliberation and AI governance, and brief case studies of a number of existing public deliberations, we organized a series of 15 semi-structured interviews. These invited experts on public engagement and AI policy from across the world to reflect on the issues, opportunities, challenges and design features that may be important to address in developing a global citizens assembly on AI. Interviews took place online, and were recorded, machine-transcribed and coded for analysis.

We also held a 2-hour workshop with seven technology policy and public engagement practitioners in Brussels on 24th May 2024 to explore design options, and then presented and received feedback on draft options at two workshops, one in Oxford on 18th July (approx 20 participants), and in an online workshop on 12th August (7 participants).

We have included selected illustrative quotes in this report, and have attributed these on the basis of whether they were from in-depth interviews or workshops. In some cases, quotes have been paraphrased and edited for clarity and flow.

The report solely represents the view of the authors, and should not be taken to represent the position of any interviewee or workshop participant.

Contributors

With thanks to everyone who has contributed to the research interviews, workshops and reviews for this project, including: Adam Cantwell-Corn, Andrew Sorota, Anna Beckett, Anna Colom, Anthony Zacharzewski, Antoine Vergne, Audrey Tang, Aviv Ovadya, Bob Fay, Bruce Schneier, Charlie Pownall, Connor Dunlop, Elinor Carmi, Emily Macaulay, Fernando Perini, Flynn Devine, Giovana Fleck, Gloria Guerro, Gryffin Wilkens-Plumley, Helena Hollis, Helene Landemore, Isabelle Ferreras, Jeni Tennison, Jenna Slotin, Laura Sandor, Lena Slachmuis, Matt Ryan, Meg Young, Moiz Shaikh, Nardine Alnemr, Nigel Shadbolt, Omer Bilgin, Polen Türkmen, Raegan MacDonald, Reema Patel, Renata Avila, Richard Wilson, Sophie Tomlinson, Thai Jungpanich & Tiago Piexoto.

Icons from thenounproject.com used under license. Infographic by [Convincible.media](https://convincible.media)

References

1. *Interim Report: Governing AI for Humanity*. United Nations Advisory Body on Artificial Intelligence; 2023. Accessed July 8, 2024.
https://www.un.org/sites/un2.un.org/files/un_ai_advisory_body_governing_ai_for_humanity_interim_report.pdf
2. OECD. *Innovative Citizen Participation and New Democratic Institutions: Catching the Deliberative Wave*. Organisation for Economic Co-operation and Development; 2020. Accessed July 10, 2024.
https://www.oecd-ilibrary.org/governance/innovative-citizen-participation-and-new-democratic-institutions_339306da-en
3. Smith G. *Democratic Innovations: Designing Institutions for Citizen Participation*. Cambridge University Press; 2009. doi:10.1017/CBO9780511609848
4. Reuchamps M, Vrydagh J, Welp Y, eds. *De Gruyter Handbook of Citizens' Assemblies*. De Gruyter; 2023. doi:10.1515/9783110758269
5. Lerner J. *From Waves to Ecosystems: The Next Stage of Democratic Innovation*. SNF Ithaca; 2024. <https://udspace.udel.edu/items/95bf3dbb-9990-483c-9040-197f42060df1>
6. UNESCO. Recommendation on the Ethics of Artificial Intelligence. Published online 2021. <https://unesdoc.unesco.org/ark:/48223/pf0000381137>
7. Ovadya A. Reimagining Democracy for AI. *Journal of Democracy*. 2023;34(4):162-170.
8. Atlas. The Survival Bill: Control Artificial Intelligence. Published online 2024. Accessed September 4, 2024.
<https://web.archive.org/web/20240515232758/https://www.atlasmovement.org/issues>
9. Malkin C, Maier F, Alnemr N. *Agenda-Setting in Transnational and Global Citizens' Assemblies*. Global Citizens' Assembly Network; 2023. Accessed July 11, 2024.
<https://glocan.org/wp-content/uploads/2024/01/Technical-Paper-2-2023-Malkin-et-al.pdf>
10. Niessen C, Reuchamps M. Designing a Permanent Deliberative Citizens' Assembly: The Ostbelgien Modell in Belgium. Published online 2019. Accessed July 11, 2024.
https://researchportal.unamur.be/files/43760057/Niessen_Reuchamps_2019_Designing_a_permanent_deliberative_citizens_assembly.pdf
11. Democratic Odyssey. European University Institute. Accessed July 11, 2024.
<https://democraticodyssey.eui.eu/home>
12. Tanner J, Bryden J. *Reframing AI in Civil Society: Beyond Risk and Regulation*. Rootcause; 2023. <https://rootcause.global/framing-ai/>
13. Curato N, Chalaye P, Conway-Lamb W, et al. *Global Assembly on the Climate and Ecological Crisis Evaluation Report*. University of Canberra; 2023.
14. Fjeld J, Achten N, Hilligoss H, Nagy A, Srikumar M. Principled Artificial Intelligence: Mapping Consensus in Ethical and Rights-Based Approaches to Principles for AI. Published online January 15, 2020. doi:10.2139/ssrn.3518482
15. Corrêa NK, Galvão C, Santos JW, et al. Worldwide AI ethics: A review of 200 guidelines and recommendations for AI governance. *Patterns*. 2023;4(10):100857. doi:10.1016/j.patter.2023.100857
16. Veale M, Matus K, Gorwa R. AI and Global Governance: Modalities, Rationales, Tensions. *Annual Review of Law and Social Science*. 2023;19(Volume 19, 2023):255-275. doi:10.1146/annurev-lawsocsci-020223-040749
17. AI Safety Summit: introduction. GOV.UK. October 31, 2023. Accessed July 11, 2024.

- <https://www.gov.uk/government/publications/ai-safety-summit-introduction>
18. Yohsua B, Daniel P, Tamay B, et al. *International Scientific Report on the Safety of Advanced AI*. Department for Science, Innovation and Technology; 2024. Accessed July 11, 2024. <https://hal.science/hal-04612963/>
 19. United Nations. *Global Digital Compact: Rev 2 Draft.*; 2024. Accessed July 11, 2024. https://www.un.org/techenvoy/sites/www.un.org.techenvoy/files/GlobalDigitalCompact_rev2.pdf
 20. Cass-Beggs D, Clare S, Dimowo D, Kara Z. Framework Convention on Global AI Challenges. Published online 2024.
 21. 2024 World AI Conference and High-Level Meeting on Global AI Governance. Full text: Shanghai Declaration on Global AI Governance_Ministry of Foreign Affairs of the People's Republic of China. Published online 2024. Accessed August 4, 2024. https://www.fmprc.gov.cn/eng/zy/gb/202407/t20240704_11448351.html
 22. Machani A. *Designing Global Citizens Assemblies for Impact: Power Mapping of the United Nations System*. Global Citizens' Assembly Network; 2024. <https://glocan.org/wp-content/uploads/2024/03/Technical-Paper-3-2024-Machani.pdf>
 23. Estier M. *The Relevance of WSIS & IGF for International AI Governance*. Simons Institute; 2024. Accessed August 12, 2024. <https://www.simoninstitute.ch/blog/post/the-relevance-of-wsis-igf-for-international-ai-governance/>
 24. Belgian Citizens' Panel on AI. *A Citizens' View of Artificial Intelligence within the EU*. Secretary General of FPS Foreign Affairs, Foreign Trade and Development Cooperation, Belgium; 2024. Accessed July 11, 2024. <https://belgian-presidency.consilium.europa.eu/media/lzxauu4i/rapport-ia-en-web-v2.pdf>
 25. Siegmann C, Anderljung M. The Brussels Effect and Artificial Intelligence. Published online October 14, 2022. doi:10.33774/apsa-2022-vxtsl
 26. Pouget H, Dennis C. *The Future of International Scientific Assessments of AI's Risks*. Carnegie Endowment for International Peace; 2024. Accessed September 4, 2024. <https://carnegieendowment.org/research/2024/08/the-future-of-international-scientific-assessments-of-ais-risks?lang=en>
 27. Davies T. *Involving the Public in AI Policymaking - Experiences from the People's Panel on AI*. Connected by Data; 2024. Accessed July 12, 2024. [https://connectedbydata.org/assets/projects/peoplespanel/2024%20-%20Peoples%20Panel%20on%20AI%20-%20Final%20Report%20\(10%20Pages\).pdf](https://connectedbydata.org/assets/projects/peoplespanel/2024%20-%20Peoples%20Panel%20on%20AI%20-%20Final%20Report%20(10%20Pages).pdf)
 28. Participedia. Oregon Citizens' Initiative Review – Participedia. Accessed July 12, 2024. <https://participedia.net/method/592>
 29. Global Assembly. *Report of the 2021 Global Assembly on the Climate and Ecological Crisis.*; 2022. Accessed July 12, 2024. <https://globalassembly.org/resources/downloads/GlobalAssembly2021-FullReport.pdf>
 30. Srivastava L. Building Community Governance for AI (SSIR). *Stanford Social Innovation Review*. Published online 2024. Accessed September 5, 2024. <https://ssir.org/articles/entry/ai-building-community-governance>
 31. Participedia. Transnational Citizens' Assemblies – Participedia. Accessed July 12, 2024. <https://participedia.net/collection/8376>
 32. Vergne A, Potet M, Verhulst B, Tran-Karcher H, Schurhoff F, Tazi M. *We The Internet Results Report 2020 Edition - What Citizens of the World Say on the Future of the Internet*. Missions Publiques; 2020. Accessed July 12, 2024. <https://wetheinternet.org/wp-content/uploads/2020/12/WTI-final-results-report-v1h.pdf>

33. World Wide Views. The World Wide Views method – World Wide Views. n.d. Accessed September 5, 2024. <https://wwviews.org/the-world-wide-views-method/>
34. Davies, Tim, Fumega, Silvana. *Global Data Barometer - First Edition*. Zenodo; 2022. doi:10.5281/ZENODO.6488349
35. Battiste M. *Decolonizing Education: Nourishing the Learning Spirit*. Purich publishing limited; 2013.
36. The Collective Intelligence Project. *The Collective Intelligence Project: Whitepaper.*; 2023. Accessed July 10, 2024. <https://cip.org/whitepaper>
37. Boucher S, Hallin CA, Paulson L. *The Routledge Handbook of Collective Intelligence for Democracy and Governance*. 1st ed. Routledge; 2023. doi:10.4324/9781003215929
38. Landemore H. *Democratic Reason: Politics, Collective Intelligence, and the Rule of the Many*. First paperback printing. Princeton University Press; 2017.
39. Polis. Accessed September 5, 2024. <https://pol.is/home>
40. The Public AI Network. *Public AI. A New Approach to Public-Interest AI Investment.*; 2024. Accessed September 5, 2024. <https://drive.google.com/file/d/1BPonjBZlIV9O5OTLYPQdz3KL1a7kjtWN/view>
41. vTaiwan. vTaiwan. Accessed September 5, 2024. <https://info.vtaiwan.tw/#five>
42. Zaremba W, Dhar A, Ahmad L, et al. OpenAI: Democratic inputs to AI. OpenAI Blog. 2023. Accessed December 18, 2023. <https://openai.com/blog/democratic-inputs-to-ai>
43. Malkin C, Alnemr N. *Big Tech-Driven Deliberative Projects*. Global Citizens' Assembly Network; 2024. <https://glocan.org/wp-content/uploads/2024/05/Technical-Paper-5-2024-Malkin-and-Alnemr.pdf>
44. Chang S, Ciesla E, Finch M, et al. *Meta Community Forum on AI: Results Analysis.*; 2024. Accessed July 11, 2024. <https://cddrl.fsi.stanford.edu/publication/meta-community-forum-results-analysis>
45. Ganguli D, Huang S, Lovitt L, Siddarth D. Collective Constitutional AI: Aligning a Language Model with Public.... Anthropic. 2023. Accessed December 21, 2023. <https://www.anthropic.com/index/collective-constitutional-ai-aligning-a-language-model-with-public-input>
46. McKinney S. Integrating Artificial Intelligence into Citizens' Assemblies: Benefits, Concerns and Future Pathways. *Journal of Deliberative Democracy*. 2024;20(1). doi:10.16997/jdd.1556
47. Wilson C. AI worse than humans in every way at summarising information, government trial finds. *Crikey*. <https://www.crikey.com.au/2024/09/03/ai-worse-summarising-information-humans-government-trial/>. September 3, 2024. Accessed September 17, 2024.
48. About Sciencewise. Sciencewise. Accessed September 9, 2024. <https://sciencewise.org.uk/about-sciencewise/>
49. KNOCA - Knowledge Network On Climate Assemblies. Accessed September 9, 2024. <https://knoca.eu/>
50. Hajer M. Policy without polity? Policy analysis and the institutional void. *Policy Sciences*. 2003;36(2):175-195. doi:10.1023/A:1024834510939
51. Mellier C, Wilson R. Getting Real About Citizens' Assemblies. European Democracy Hub. October 10, 2023. Accessed September 11, 2024. <https://europeandemocracyhub.epd.eu/getting-real-about-citizens-assemblies-a-new-the-ory-of-change-for-citizens-assemblies/>
52. Mellier C, Smith G. *The Changing Shape of Environmental Activism: Disruption, Engagement, Backlash.*; forthcoming.

About the authors

Connected by Data is an independent non-profit working to transform narratives, policy and practice to center communities in the governance of data and AI. We want communities to have a powerful say in decisions about the future of technology, so that it is used to create a just, equitable and sustainable world.

Find our more at connectedbydata.org

Iswe Foundation is an independent democratic innovation center. Our mission is to make governance systems work for people and planet. We are headquartered in London but operate internationally.

We are experts in democratic innovation and social change. We believe that many of our biggest challenges are symptomatic of a crisis in governance. We therefore work closely with academics, governments, movements and the media to accelerate democratic systems change. Find our more at iswe.org