


A Citizens' Track on AI Governance

Agency, Alignment and
Accountability



Tim Davies, Octavia Field Reid, Susan Oman and Rich Wilson



Acknowledgements

Written by Tim Davies, Octavia Field Reid, Susan Oman and Rich Wilson

With contributions from Emrys Schoemaker, Canning Malkin, Friederike Schueuer, Johnny Stormonth-Darling, Pierre Noro, Jeni Tennison, Ming Zhuang and Astha Kapoor

Editorial support: Debbie Kite

Design by: Jesi Carson

The authors would like to thank the many individuals and organisations who have contributed to the development of the citizens' track initiative, this paper, and the accompanying PAVE case book through working group calls and ad-hoc conversations.

This work is informed and inspired by the work of practitioners and community members across the world, from the pioneers of the global citizens' assembly on food systems and climate, to the hundreds of participants who have shared work at the Participatory AI Research and Practice Symposium (PAIRS) exploring how the voices of affected communities can be centred in AI governance.

We also extend our thanks to the organisations that have provided the seed funding to support this work (The Swiss Ministry of Foreign Affairs, the University of Sheffield, AI Collaborative and PAIRS) and those who have provided considerable in-kind time and effort (Connected by Data, Ada Lovelace Institute, University of Sheffield Centre for Machine Intelligence, and Iswe Foundation).

Suggested citation:

Davies, T., Field Reid, O., Oman, S., & Wilson, R. (2026). A Citizens Track on AI Governance: Agency, Alignment & Accountability. Iswe Foundation.
<https://doi.org/10.5281/zenodo.21157507>

Online version:

<https://citizens-track.org/2026/report>

Contents

4	Executive Summary
5	01 AI governance has a democratic deficit
9	02 The UN Global Dialogue and Swiss AI Summit provide a critical moment of opportunity
12	03 Participatory practices can meet the challenge
17	04 A citizens' track provides the infrastructure for insight and action
22	05 Participation focused on outcomes and impact
24	06 An invitation to build
26	Bibliography

Executive Summary

People across the world affected by rapid AI developments must be engaged in meaningfully shaping decisions about its future, otherwise efforts towards both AI deployment and AI governance will lack legitimacy and risk public backlash.

A commitment is needed to embed public participation into the architecture of AI governance. This will link bottom-up capacity building and deliberation with top-down global fora: driving more responsive and inclusive governance decisions while mobilising local communities as allies for governance reform and action. The citizens' track is a mechanism to do this: **listening** to, **connecting** with, and **embedding** spaces for public voice from around the world.

On current trajectories, AI risks being misaligned with public interest, human rights, sustainable development and safety. Despite substantial benefits and harms of AI already being experienced, evidence frequently lags impacts and fails to represent much of the world's population. To date, AI governance fora have excluded much of the globe, and AI-related hopes and fears of affected communities have been given little space.

However, the evidence is clear that participatory processes can successfully bring people into high-stakes, highly technical, global governance discussions. Inclusive models demonstrate that participation builds deep public legitimacy, cuts through geopolitical and corporate gridlock, increases resilience to capture, and ensures that the communities most impacted by systemic transformation have a direct hand in shaping guardrails and goals.

The new UN Global Dialogue on AI Governance, and the 2027 edition of the AI Summit Series in Geneva, Switzerland provide key moments where the current trajectory can be reversed. By leveraging recent developments in participatory democracy, and tapping into a growing network of national, local and grassroots groups organising participatory processes on AI, power-holders in AI governance can create the conditions to build legitimacy into AI.

Drawing on learning from the field of climate governance, and a new case book describing participatory AI governance processes from every

continent, this report sets out a power-aware theory of change for public participation, which is able to deliver:

- Greater **alignment** of both AI models and governance policies with a deeper and more inclusive understanding of diverse global public interests, and of human dignity and flourishing;
- Increased **agency** of individuals and communities, through critical understanding, local decision making, and involvement in global AI governance debate;
- Stronger **accountability** through social license, community audit and hard civic power.

To get there requires complementary actions from three domains:

- **Leaders** of the emerging AI governance architecture must build political will for participation and commit to creating 'docking points' that provide space to hear, consider and engage with evidence from public deliberations across the world;
- **Funders** should support democratic approaches to AI that prioritise inclusion, plurality and agency, by investing in both grassroots community assemblies on AI, and in the institutional infrastructure needed to link top-down and bottom-up processes, including transnational public deliberation;
- **Practitioners** should collaborate in the creation and use of high-quality evidence and resources for informed public deliberation on AI, supporting community assemblies that build critical AI literacy and meaningful participation linked to a powerful theory of change.

We invite you to work with us on taking the next steps towards building a citizens' track: for a future alongside AI that puts people in the lead.

www.citizens-track.org

AI governance has a democratic deficit

There is a democratic crisis at the heart of AI governance. AI governance can no longer be seen as primarily a technical challenge, but needs to be recognised as a political one. Powerful AI companies and nation states are pushing both for and against stronger regulation, while public demand for accountability and democratic oversight continues to grow. If even a fraction of the changes AI is forecast to bring are realised, then people across the world are being asked to rapidly adopt or adapt to technologies over which, at present, they have little meaningful say. Inclusive AI governance requires more than increasing the number of nations at the table in global fora. To ensure public legitimacy, people affected by AI, from all regions, age groups, cultures and contexts, must have a role in shaping our AI futures.

Amidst profound uncertainty about how AI might unfold, impacts from current AI systems are being felt right now across the globe. AI stands to reshape almost every domain of society rapidly, simultaneously, and unevenly. Unequal distribution may – in different contexts – mean exclusion from the benefits of AI, or exposure to harms. The challenge is not simply that the impacts may be large, it is that they are deeply interconnected and highly uncertain.

As a result, societies face a paradox. The potential consequences of AI deployments appear so significant that governance is urgently required. Yet the apparent speed, scale and complexity of the questions AI raises can make institutions, leaders and people alike feel less able to take action. The result is a trajectory for AI that risks being misaligned with public interest, human rights, sustainable development and safety.

Trust of AI systems and worry about AI in 2022 and 2024




Figure 1: Change in trust in AI systems and worry about AI between 2022 and 2024 based on a multi-country public survey covering 17 (2022) and 47 (2024) countries (Gillespie et al. 2025).

Public confidence in AI is fragile, and experience with AI does not necessarily alleviate concerns. Across the world public trust in AI is falling, and polarisation over AI's role in society is growing.

If people don't have a stake in decisions about the opportunities and implications of AI, its benefits will not be realised. Instead we face a future in which, however well intended, technological change, and the regulation of that change, is experienced as imposition from above: less likely to work for communities, and more likely to be met with rejection and backlash. If people do not have agency to debate and decide on AI, then democracy itself is eroded.

Public excitement about AI is falling

More excited: 

Less excited: 

Figure 2: Based on net change between the proportion of respondents in each country surveyed reporting 'Products and services using artificial intelligence make me excited' in the Ipsos AI Monitor between March/April 2025 and 2026. Adapted from Carmichael, 2026.

Better evidence is needed to guide AI governance, yet scientific evidence frequently lags impacts¹.

Public opinion surveys provide a snapshot of public attitudes at surface-level, but offer little insight into why people feel the way they do, or the policies and actions that might address the hopes and fears they express. Current monitoring mechanisms (from mapping of AI implementation and uptake, to projects such as the OECD Incidents and Hazards Monitor) often lack granularity, or are biased towards issues that attract media coverage, are legible to journalists, or occur in contexts with the greatest private sector interest, or civil society and press freedoms. Majority world perspectives are frequently missing and global surveys often reach only digitally connected elites. AI impacts that are diffuse, cumulative, and that particularly affect marginalised populations are underrepresented in the current evidence-base informing AI governance².

At the same time that AI governance demands more of the global system, multistakeholder governance is under strain. The succession of Safety, Action and Impact summits³ reflects a rapidly broadening agenda, though one still light on coherence, binding commitments and accountability mechanisms. The rapid expansion of AI governance risks fragmentation as frameworks emerge faster than they connect. The Council of Europe inventory alone identifies some 450 distinct AI governance initiatives.

As has already been seen in climate governance, if governance initiatives cannot connect with the people they are meant to benefit, critical loops of feedback accountability will remain broken, and trust in multistakeholder governance may further erode. Current institutions must evolve, and work with and through new partners to deliver change.

¹ The Independent International Scientific Panel on AI has been established to provide more rapid synthesis of scientific literature on AI. However, the co-chairs of the panel noted in their informal briefing to the General Assembly that their work still faces an 'evidence gap' created both by the time taken for rigorous scientific investigations, and the rapid deployment of AI ahead of published research on its impacts.

² (Oman and Cannizzarao 2025)

³ The AI Summit series began with the AI Safety Summit at Bletchley Park initiated by the UK Government (November 2023), followed by the AI Seoul Summit (May 2024), AI Action Summit (Paris, France, Feb 2025), and India AI Impact Summit (New Delhi, Feb 2026). Host countries have designed and managed the invite lists for each event, with an emphasis on government and industry participation.

Innovations in public participation, and in multi-country and transnational deliberation, have demonstrated that participatory processes can successfully bring people into high-stakes, highly technical global governance discussions⁴.

Rather than increasing fragmentation or amplifying the loudest voices, inclusive models of putting “people in the lead” have demonstrated the exact opposite: participation builds deep public legitimacy, cuts through geopolitical and corporate gridlock, increases resilience to capture, and ensures that the communities most impacted by systemic transformation have a direct hand in shaping guardrails and goals. Participatory processes can build an agile evidence base on opportunities and harms and can deliver robust insight on what the public is expecting from AI governance: supporting both efforts to deliver a governance floor for AI⁵, and to articulate a collective vision of sustainable development in an era of AI⁶.

This paper explores a route to engage the global field of public participation to help align AI with human dignity and flourishing, strengthen human agency in an age of AI, and deliver accountability in AI governance. The participatory democracy wave of the last decade⁷ shows that diverse communities can be part of finding policy solutions and breaking governance deadlocks. The latest wave of grassroots participatory activity needs to be matched by powerholders embedding participatory mechanisms into the new infrastructures and decision-making spaces of AI governance.

7243 cases of participatory process (any topic) identified so far




Figure 3: Iswe Foundation's Global Participation Map showing examples of over 7,000 participatory processes across the world (2015-2025). Sources: Community Engagement Hub (1.5%); KNOCA (2.2%); LATINNO (43.2%); OECD 2023 (8.5%); ODP (4.1%); OPSI (4.4%); Participatory Budgeting (8.0%); Participedia (27.0%); People Powered (1.2%)

4 See (Reuchamps et al. 2023; Bächtiger and Dryzek 2024; Dryzek et al. 2011; Curato et al. 2025, 2023; Niemeyer et al. 2024; Chwalisz 2019; Stilgoe 2024)

5 (Schüür et al. 2026)

6 (Melamed and Ladd 2026)

7 (Smith 2009; Chwalisz 2019; OECD 2020)

The UN Global Dialogue and Swiss AI Summit provide a critical moment of opportunity

The international architecture of AI governance is evolving rapidly, and the year ahead presents a unique moment of opportunity to embed a participatory and democratic approach. The first edition of the [UN Global Dialogue on AI Governance](#), created through the [Global Digital Compact](#) and [General Assembly resolution \(A/RES/79/325\)](#) takes place in Geneva in July 2026, establishing the foundation for future negotiations on AI governance in 2027. For the first time, all UN member states have a seat at the table, confronting the structural equity deficit in AI governance to date, where a small number of countries have shaped the agenda, and upwards of 118 countries have been excluded⁸.

Paired with the Independent International Scientific Panel on AI, the UN Global Dialogue seeks to act as a 'dialogue of dialogues', bringing together prior work from ITU, UNESCO, UNDP, G7, G20, the OECD and the AI Summit Series, as well as acting as a convening point for a wide range of multistakeholder initiatives on AI. In preparatory consultations for the UN Global Dialogue hosted by the co-chairs, numerous contributors have underlined the importance of an inclusive process that goes beyond national diversity, to also consider how diverse communities affected by AI can be brought to the table.

As the institutional structure of the UN Global Dialogue develops, member states and other stakeholders have a unique opportunity to build in citizen participation from the start: listening to existing public voices, connecting with the global participatory democracy field, and embedding formal mechanisms to engage with public perspectives on AI as a complementary source of insight and direction-setting, alongside inputs of the Independent International Scientific Panel on AI.

In the field of climate governance, it has taken almost 30 years for a recognition that the scientific and diplomatic pairing of IPPC and COP needs to be accompanied by a participatory citizens' track: tapping into participatory processes to both inform, and activate, agendas for change⁹. AI governance can take on this learning and seize this moment to centre public voices from the start.

⁸ ('Governing AI for humanity' 2024)

⁹ (Oliveira 2025)

On who should be at the table



“People from rural communities are missing. In Uganda, most people do not live in Kampala. They live in villages. They farm. They trade in local markets. AI is already affecting their access to loans, health information and government services. But nobody is asking them what they think about it.”

— Flavia Ouma, Uganda Bureau of Statistics, Africa

“The 2,000-plus languages spoken across Africa, the Pacific, and Indigenous communities globally are almost entirely absent not just as subjects of discussion, but as languages of participation.”

— Roland F. Ganafa, AI Studio Uganda, Africa

“Indigenous voices are often treated as data sources rather than sovereign partners. Their perspectives on communal privacy and the preservation of low-resource languages are vital to prevent AI from becoming a tool of cultural erasure.”

— Dr. Shaista Hussain, SAIF CHECK, Asia and the Pacific

“Women and girls, particularly from underrepresented and developing regions, remain significantly excluded from AI governance discussions, alongside grassroots educators and community-led innovators.”

— Paulette Watson MBE, Academy Achievers, Western Europe and Other States

Figure 4: Summary of insights from written submissions to the *UN Global Dialogue on AI Governance open consultations*. Source: *UN Global Dialogue on AI Governance Joint Secretariat. Prepared based on 1,532 written submissions (Joint Secretariat of the Global Dialogue on AI Governance 2026)*

A note on language

The term ‘citizen’ has a range of meanings. State citizenship is experienced by some as an exclusionary or colonial category: where citizens have greater rights and entitlements than non-citizens. For others, the idea of citizenship is not primarily a status ascribed by the state, but draws on a recognition of universal rights and responsibilities, and the belief that all individuals can participate in civic solidarity whether at the collective level of a local area, nation or planet¹⁰. When we talk about citizen participation and engagement, we draw on this second sense of universal citizenship.

The term ‘citizen’ and ‘person’ can be read as equivalent. Our use of the term ‘citizen’ is to maintain continuity with existing discourse on citizens’ assemblies and participation, and to emphasise that participation is by right not by the benevolence of power-holders.

¹⁰ (Kabeer 2005)



The 2027 edition of the AI Summit Series, hosted in Switzerland, provides a complementary opportunity to drive innovation in participatory AI governance.

Building on the Swiss tradition of direct democracy, and on successive civil-society led pilots that have surfaced public inputs for previous editions of the AI Summit series, the next AI Summit can lead by example and centre public voices.

This could take the form of commissioning distributed public dialogues feeding into the Geneva Summit, and providing support for a flagship transnational convening of citizens from across the world at the Summit to deliberate together and share their review of proceedings¹¹. This would act as a powerful demonstration of inclusive capacity building, and of developing AI futures with and for people.

Meeting this moment requires action from governments, funders, industry and practitioners. As the next section of this report outlines, the foundations of participatory practice are here, and offer a cost-effective approach to build capacity and coordinate legitimate action. However, funding for participatory practice remains insufficient, and rhetorical claims about the importance of listening to the people have not yet been met with the institutional and infrastructural support required to make participation meaningful. Grassroots practitioners need resources to enable informed, actionable and joined-up dialogues on AI that build agency and connection, rather than risking disempowered forms of engagement that may reinforce disaffection and frustration.

¹¹ See (Davies 2024) for details of deliberative review piloted at the 2023 AI Safety Summit, and public participants' call for future iterations to include input from citizens from across the world.

Participatory practices can meet the challenge

Placing public voices at the heart of governance can bring legitimacy and impact to a field dominated by powerful political, economic and technological interests. Successive waves of participatory practice have delivered robust methods to bring citizen voices to the table around complex technology and policy topics. The next generation of practice recognises public deliberation as both an input to global governance and an engine for action. This link between participation and power directly addresses concerns about the legitimacy of governance responses to interconnected global challenges.

Since the 1970s, tens of thousands of participatory processes have engaged citizens in decisions on budgets, planning, and infrastructure. More than 11,000 participatory budgeting initiatives have been implemented worldwide, enabling communities to

decide how public resources are spent¹². The OECD has tracked more than 700 citizens' assemblies and mini-publics, with participation accelerating sharply in the past decade. However, this figure likely underrepresents activities, as democratic innovation theorists point to the flourishing of new channels of citizen participation, often outside of visible spaces of political engagement¹³.

This infrastructure exists across global north and south alike. Data from the [V-Dem Project](#)¹⁴ shows that strong citizen participation is not limited to wealthier countries: many global south states score highly, reflecting long traditions of community organising and co-governance. Together, these practices demonstrate that people have the capacity, skills, and appetite to participate directly in shaping policy.

Civil Society Participation Index, 2025

Our World
in Data

Data by V-Dem. Expert Estimates of the extent to which citizens are active in diverse organizations which choose and influence policy-makers. The index ranges from 0–1 (most active).




Figure 5: Civil Society Participation Index 2025.
Source: [Our World in Data](#)

¹² (Hagelskamp et al. 2021)

¹³ (Warren 2025)

¹⁴ (Michael Coppedge et al. 2026)

Defining Participation, Deliberation, Assembly

Participation invites non-experts and affected communities to inform, shape or make decisions about AI-related policy, system design or technology deployment. Participation can take many forms, such as live consultations, surveys, online polling with opinion mapping, interactive workshops and deliberative dialogues.

Deliberation is a particular form of participation, ideally suited to addressing topics like AI governance which involve technical complexity, values-based trade-offs, rapid change and uncertain timelines, trajectories and impacts. Deliberation involves intentional participant recruitment, a learning phase, and facilitated dialogue that draws on both expert inputs and lived experience.

Assemblies bring people together synchronously online or in-person to engage in collective dialogue. A fundamental component of social and political life, assembly has gained renewed focus in recent years as a vital corrective to democratic distance created by globalisation¹⁵.

- **A community assembly** is a locally-organised gathering, run in ways that reflect local cultures of convening and dialogue. They provide space to address both shared questions and local concerns. They can identify messages for higher-levels of decision making and organise around actions that can be taken locally.
- **A citizens' assembly** is a formal convening that invites a 'mini-public', often matching the demographics of a wider population, to deliberate over multiple days on a focussed topic. Citizens' assemblies have been run at city, national and transnational levels, creating a synthesis of recommendations for policy¹⁶.

A citizens' track provides an institutional interface between ad-hoc public participation and ongoing spaces of governance (such as the UN Global Dialogue on AI Governance). It connects with community assemblies, and may commission or support a citizens' assembly to support sense-making across diverse public inputs.

The design of a citizens' track incorporates values of subsidiarity (making decisions at the lowest feasible level, which may be local or global); solidarity (recognising global inter-relationships); plurality (building from diversity of perspective, culture and context); and power (adopting power-literate theories of change and avoiding tokenism).

¹⁵ (Davidian and Jeanpierre 2023)

¹⁶ A 2026 review by People Powered of nearly 70 studies into Citizens' Assemblies identified evidence for 19 key impacts, including: increased political knowledge and efficacy; reduced polarization and increased openness to opposing views; shifts of attitude on specific issues; influence on the broader public attitudes and discourse; provision of legitimacy and a public mandate for difficult decisions; influence on policymaking culture and institutional openness; education of policymakers and even policy shifts through introducing new perspectives; reshaping of traditional lobbying and advocacy dynamics; and catalysis of local action through bypassing of institutional inaction (Gem and Milosh 2026).

The field of grassroots participatory AI is rapidly developing. In an initial research sprint for this paper, we crowdsourced and compiled case studies on over 80 participatory and deliberative processes focussed on AI in the last five years. These have included over 200,000 participants from more than 95 countries, and involved over 280 different organisations from governments, industry, academia and civil society.

This mapping reveals the spread of public engagement on AI, as well as geographic and thematic gaps that need attention and investment to ensure inclusive inputs to future global processes.

PAVE Case Book of Public Participation




Figure 6: Mapping from the PAVE Case Book of public participation on AI, drawing on crowdsourced contributions (May/June 2026) describing selected examples of participatory AI governance that are exemplary for their context-specific rigour in inclusive practice, meaningful engagement and/or impact. The full case book is available to browse online at <https://pave.pairs.site> and further details of the rationale are published in [Oman et al. 2026](#).

PAVE Case Book Examples: who is hosting AI-focussed deliberation?

- **Governments.** The Government of Belgium organised a multi-lingual Citizens Panel on AI as part of the Belgian presidency of the European Union in 2024. 60 people deliberated over 40 hours to deliver 9 key messages. [\[Read more\]](#)
- **Companies.** A cross-industry community forum commissioned by Meta, Oracle, DoorDash, PayPal, Cohere and Microsoft used deliberative polling to involve 503 people, connecting online through a digital facilitation platform, to shaping recommendations on the regulation of advanced AI agents. [\[Read more\]](#)
- **Workers.** Mozilla Foundation and Berggruen Institute brought together 91 creative workers in Los Angeles through five assemblies, exploring how the ‘dawn of AI’ was impacting upon creative practice, and putting forward “Hollywood’s 8 Rules For AI”. [\[Read more\]](#)
- **Researchers.** The EPFL AI Center organised the Citizen’s Assembly on AI in French Speaking Switzerland in 2025, delivering 19 recommendations on AI strategy, literacy and governance after 40 people took part in four days of learning and deliberation sessions. [\[Read more\]](#)
- **Communities.** In Kalinga, in the Cordillera highlands of the northern Philippines, a ‘Digital Arang’ process built on traditional cultures of deliberation to produce a local AI Policy Ordinance on the use of Indigenous data in AI systems. [\[Read more\]](#)

Digital platforms and AI can help scale and deepen participation, but the real strength of participatory practice comes from fostering solidarity and connection.

Projects are experimenting with the use of AI to support participant recruitment and deliberation planning, to create accessible learning materials, to scaffold process facilitation, to transcribe and translate discussions, to aggregate across deliberations and generate consensus statements, and to communicate deliberation findings¹⁷. Industry-led community fora have drawn on custom video conferencing platforms to facilitate hundreds of live, parallel, small group discussions¹⁸, and many initiatives are inspired by pioneering uses of collective intelligence platforms such as Pol.is or Decidim that draw on bridge-based algorithms to help process-leaders find common ground around contested issues¹⁹.

These approaches can bring inputs from thousands of individual voices, but effective deliberation must retain the ability for marginalised voices to be heard, and for participants to explore collective perspectives, fostering and drawing on relational interactions. In exploring models for local and global deliberation on AI governance²⁰, large-scale online participation can offer a ‘feeder lane’ to a citizens’ track, potentially supporting the formation of inclusive local community assemblies.

17 (McKinney 2024)

18 (Fishkin et al. 2019; Gelauff et al. 2023)

19 (Mikhaylovskaya 2024)

20 (Davies et al. 2024)

Online tools can also present possibilities for validating and sharing outcomes from community and large scale assemblies, delivering wider reach for the outputs of deliberative mini-publics. This view is supported by insights from DemNext, who describe how *“scaling democratic deliberation is not a technological challenge alone, but one that requires a diverse repertoire of technological applications to be developed and fruitfully combined with strengthened civic infrastructure”*²¹.

Insights from PAIRS and PAVE: What can participation address?

In 2025 and 2026 the Participatory AI Research & Practice Symposium (PAIRS) brought together over 200 researchers and practitioners to share insight from participatory AI projects across the world, exploring how public engagement can shape AI systems, governance rules, and contestation over AI adoption and deployment. The papers presented at PAIRS, and case studies in the PAVE Case Book, demonstrate the use of participatory processes to address:

- **Capacity building:** supporting communities to develop critical AI literacy, including both practical skills for AI use and awareness of the consequences of AI adoption that need to be considered.
- **Design decisions:** involving affected communities in decisions about the purpose, design and governance of AI specific systems.
- **Deployment decisions:** shaping local and national decisions about the deployment of AI within public services or workplaces.
- **Policy design:** co-designing policies on AI, and policies responding to the impacts of AI.
- **Accountability:** involving affected communities in red-teaming, evaluation and audit of AI systems, and mobilising resistance to imposition of discriminatory forms of AI.

Over the year ahead, millions of citizens around the world could be taking part in participatory processes on AI. From community-level deliberations, to trade-union led dialogues, digitally-mediated conversations and national engagements, people will be learning and talking about AI. This provides a key opportunity to include these conversations as part of the UN Global Dialogue ‘dialogue of dialogues’, and to make them as informed, inclusive and connected as possible.

Harnessing participatory practice for AI governance does not require the heavy lift of creating a new participation field, or building participation from scratch: rather, it requires organising and connecting with an existing and growing field to collectively address one of the defining governance challenges of our time.

21 (McKinney and Chwalisz 2025)

A citizens' track provides the infrastructure for insight and action

A citizens' track involves going beyond rhetoric about public interest or human centrism, to action that includes the public in debate and decision-making on AI. This calls for the creation of clear **docking points** within the developing AI governance architecture that provide space to hear and consider evidence from public engagement. It requires an **institutional interface** that enables two-way connectivity between global governance, and regional, national and local processes: helping to frame accessible questions that communities can consider, to facilitate balanced expert inputs for local deliberation, and to work with a **coordination layer** to aggregate and translate evidence from community deliberations into clear messages for decision makers. And it requires increased support for grassroots public engagement, addressing gaps in the current landscape to ensure an inclusive conversation, and building the capacity of **community assemblies** to be active change agents, working in step and solidarity with a global dialogue of dialogues.

The exact design of the citizens' track infrastructure interface will depend on available resources, partnerships and institutional support. Here we set out an ambitious, but achievable vision that addresses how bottom-up and top-down processes can meet and become more than the sum of their parts²².

Figure 7 provides a stylised presentation of the layers of the citizens' track. In the following paragraphs we describe this from bottom to top, finishing with a description of how power-holders can accelerate the development of the citizens' track in three stages: **listen, connect** and **embed**.

²² The model put forward here, of distributed community assemblies, linked through both global aggregation and a global citizens' assembly, draws on the five design options set out in [Global Citizen Deliberation on Artificial Intelligence: Options and design considerations](#) (Davies et al. 2024)

Components of a citizens' track




Figure 7: Components of a citizens' track

Community assemblies

Community assemblies provide the foundation of a citizens' track, reflecting a plurality of lived experiences, hope for, and concerns about AI. Prepared toolkits of learning resources, facilitation guides and question banks²³, as well as catalytic funding (particularly for assemblies in under-represented communities), can increase their breadth, depth and focus and increase legitimacy. Investing in the capacity of community assemblies to connect participation and power can build local impact loops: strengthening subsidiarity in AI decision making²⁴, while keeping interoperable and globally-connected approaches in view.

Coordination hubs

Within the coordination layer, national, regional and thematic hubs would organise and facilitate 'distributed dialogues'. Drawing on a diversity of funding sources, hubs can support focussed deliberation on priority topics, aggregate insight and recommendations, and support coordinated action in response to public input. By providing impetus and support for participation on specific themes or in specific areas, hubs energise dialogue, and work to make sure that messages reach relevant decision makers. By surfacing diverse lived experience and common concerns across communities, hubs help build solidarity and strengthen geopolitical coherence.

As a complementary element of coordination, a global citizens' assembly would synthesise public priorities across countries and regions. At critical moments of global decision-making, transnational citizen deliberation can help make sense of different perspectives and bring forward a focussed set of recommendations. A global citizens' assembly could periodically convene to consider AI, bringing together 100 to 500 sortition-selected citizens from across the world. They would hear expert inputs (including from the Independent International Scientific Panel on AI) and aggregated insights from community assemblies, before deliberating on recommendations and reporting on the degree of consensus the assembly was able to reach.

Institutional interface

An institutional layer provides the continuity of the citizens' track, and acts as an interface between AI-focussed governance fora and public participation. As a facilitation layer, rather than a gatekeeper, this would provide a standing secretariat, able to maintain the resources and relationships needed to ensure public voices are included within, and are legible to, sites of AI power: from multilateral institutions, to industry players and open source public AI collaborations. The institutional layer would host a federated database of community assembly outcomes, drawing on the foundations of the [PAVE Case Book](#) and work on [participation platform interoperability](#), and would support resource mobilisation to address gaps in thematic, geographic or community representation.

²³ For example, accessible fact sheets, short videos of diverse expert perspectives, and suggested hands-on learning activities, translated into local languages, might be provided to support a distributed dialogue process.

²⁴ (Marcucci et al. 2025)

Docking points

Existing AI governance fora, and AI power-holders need to create docking points for a citizens' track.

The UN Global Dialogue on AI Governance and AI Summit Series should support coordination with a citizens' track: ensuring diverse citizen perspectives are invited and profiled at key decision-shaping and decision-making moments. This involves providing agenda space, recognition of public participants, and other practical support. This will require two-way communication between powerholders and citizens: with a feedback loop to hubs and community assemblies that informs future rounds of distributed and local dialogue, and that engages communities as allies in action to take forward governance reform.

Other thematic governance institutions, industry actors, and open source efforts (e.g. Public AI initiatives) can also create docking points for engaging with distributed and global public deliberation: recognising that not all the perspectives aggregated through a citizens' track will be addressed by global governance institutions, and that there are other avenues for public voice and insight to contribute to change.

Building the track: listen, connect, embed

A citizens' track for AI governance is already under construction. Grassroots public deliberation on AI is steadily growing, and there have been ad-hoc spaces for public voice in global AI governance²⁵. However, to realise the full potential of participatory engagement for AI governance requires three steps:

- **Listen** - Governance institutions must demonstrate they are meaningfully engaging with existing insights, recommendations and testimony from public participation processes.
- **Connect** - Governance institutions must invite and support public deliberation on key topics. For example: co-chairs articulating open questions for community assemblies to consider between editions of the UN Global Dialogue, and providing access to experts to inform these discussions.
- **Embed** - Governance institutions must build requirements for public participation into business-as-usual, ensuring clear feedback loops, and giving weight to participatory inputs in decision making.

Figure 8 shows how the UN Global Dialogue on AI Governance provides a key opportunity to move from simply listening to public voice, towards embedding ongoing public engagement within a sustained governance cycle.

25 The open stakeholder consultations for the UN Global Dialogue on AI Governance for example offered an opportunity for perspectives from grassroots communities to be heard, although there was no specific invitation or facilitation focussed on participatory public inputs.




Figure 8: Illustration showing the growth of grassroots public participation in AI, and how this can progressively be embedded as a key element of the global AI governance architecture. Showing early public inputs to the AI Summit Series, the opportunity for the Global Dialogue to connect back with grassroots public voice, and the embedding of both grassroots participation and transnational deliberation in a future governance cycle.

The citizens' track is engaged with the 'Partnership Hub for Participatory AI Governance' proposed to the UN Global Dialogue on AI's continuity mechanism.

This brings together complementary initiatives strengthening the participatory AI field, including the [Coalition for Inclusive AI](#), [Stanford Deliberative Democracy Lab](#), [MetaGov](#) and the [Worldwide Alliance for AI and Democracy](#), and will expand to work with related initiatives including those looking at how participatory practice could be embedded in the work of AI labs.

The distinct contribution of the citizens' track is a focus on the connection between (a) top-down engagement rooted in global multistakeholder governance; and (b) bottom-up engagement from community assemblies that are sites of action as well as recommendations.

It is also important to note that a citizens' track adds to, but does not substitute for, formal civil society involvement in AI governance. Direct citizen participation is a fourth strand in multistakeholder governance alongside government, industry and civil society participation.

Participation focused on outcomes and impact

Implementing a citizens' track demands a power-aware theory of change.

Democratic and participatory AI governance can mean different things to different people, and sometimes too much energy has been focussed on proving that participation is possible rather than ensuring that it is powerful. Effective and meaningful participation requires maintaining context-specific, flexible approaches while avoiding tokenism, 'participation washing', or capture of participatory processes.

Participation, designed well, can contribute towards institutional impact (influencing decisions, strengthening democratic norms, legitimacy and trust), citizens' impacts (increased agency and collective efficacy), and substantive change (reducing AI harms, and increasing benefits). Building an appropriate institutional interface requires paying attention to the overlapping routes to impact that a citizens' track can draw on, and on setting out clear impact metrics that can guide iteration, investment and evaluation. In considering how to get there, we look at three outcome areas a citizens' track should focus on: alignment; agency and accountability.

1. Alignment

An effective citizens' track will demonstrate how AI (governance) outcomes are better aligned with the views and needs of diverse publics.

- **Aligned models (research and industry):** Independent and globally diverse inputs aggregated through a citizens' track will help proprietary and public AI labs to increase the legitimacy of technical model alignment efforts.
- **Aligned advocacy (civil society):** Deliberative approaches will build wider civil society capacity to engage with AI, expanding engagement with grounded thematic concerns (health, education, human rights), while maintaining the coherence of the AI governance agenda.
- **Aligned policy (government):** The synthesis functions of the citizens' track will help governments and multi-stakeholder fora to make sense of public priorities and understand the strength of public mandate for governance measures.

2. Agency

An effective citizens' track will demonstrate how it has developed both individual and collective agency.

- **Individual agency:** Deliberation will increase public understanding and AI literacy, supporting individual choices about whether, when, and how to refuse or use specific AI applications.
- **Collective adaptation:** Participatory dialogue will help communities move from being passive users of AI tools, to becoming co-creators of locally established norms and governance practices²⁶.
- **Collective advocacy:** Deliberation will generate sharpened policy messages and demonstrate their political saliency and legitimacy. Many community and citizens' assembly processes involve moments where findings are presented to power-holders, who are asked to respond to citizens²⁷.

3. Accountability

An effective citizens' track will demonstrate how it supports accountability mechanisms that deliver robust AI governance.

- **Social license:** Enabling communities to negotiate over permission for corporate data mining or resource use by data centres, securing transparency, setting terms for local benefits, and ensuring ongoing compliance with these agreements²⁸.
- **Audit and social audit:** Participatory audit processes involve the public in assessing claims made about AI impacts and equity. In social audit, official evaluations (evals) and audit outcomes are presented to affected communities, giving them the right of reply, and creating an accountability loop.
- **Hard civic power:** Drawing on participatory processes to ground legitimate civic actions: directing funding, shaping public narratives, and using complaints, investigations, and litigation to ensure corporate and government power-holders abide by agreed-on AI governance rules.

²⁶ Patra et. al describe this as “no-code” governance, drawing on a case study of how the Ugandan Women in AI community have used workshops, dialogue and co-creation to establish shared-boundaries for AI use. (Patra et al. 2026)

²⁷ Power-holders are often also invited to observe assembly processes, invited to think along with, and deepen their understanding of public perspectives.

²⁸ (S. G. Verhulst 2024; S. Verhulst et al. 2025)

An invitation to build

Creating the citizens' track on AI governance is up to all of us. With concerted action and focussed funding, over the next year many thousands of community assemblies could take place, building the capacity of groups across the globe to engage critically with AI, and increasing our chance of securing a future where AI is aligned with human flourishing, with strong human agency, and clear accountability. With institutional commitment, the deliberative wave can be directed to bring greater legitimacy to AI governance, bringing hundreds of diverse public voices into the rooms where decisions are shaped, and creating a two-way connection between global dialogue and communities around the world.

Through citizens-track.org, the [Participatory AI Research and Practice Symposium \(PAIRS\)](#), the [PAVE Case Book](#), the [Assemblis](#) platform, and an open, collaborative approach, we're inviting partners to join us in realising this potential. Since the start of 2026, we've brought together a number of components:

- **Citizens-Track.org** provides a coordinating space to bring together governments, civil society, funders and other stakeholders to focus on establishing the institutional layer needed to amplify existing participation, catalyse gaps-filling engagement, and to embed within global governance fora.
- **The PAVE Case Book and Assemblis** offer a prototype of the knowledge architecture needed to make sure citizen views flow from grassroots to global, and to inspire and support effective local practice.
- **The Participatory AI Research and Practice Symposium** offers an evidence layer, supporting research, evaluation and learning on inclusive and powerful forms of participatory AI.

To build further, there are three key gaps to address:

1. To date, **funding for democratic approaches to AI remains scarce**. If even a fraction of the amounts currently allocated to AI development, AI safety, or AI for Good, were put towards people-power in defining the future of AI, impacts from participatory AI would accelerate substantially²⁹. The full potential of a participatory approach to AI governance ultimately cannot be realised without long-term and diversified funding.
2. **Inclusion and plurality** must underpin the citizens' track. Taking a global approach to AI governance means recognising and valuing different approaches to solving problems in a variety of contexts. Representation requires investment: as the examples in this document and the PAVE Case Book show, participatory approaches are being employed in nations and communities across the world, but there is more work to be done to make sure these are represented in the leadership, stewardship and governance of a citizens' track itself.
3. Lastly, and most importantly, is **political will**. Current power-holders must recognise that the democratic crisis at the heart of AI can not be addressed through technocratic measures. Nor can issues of AI governance be isolated from wider concerns about public interest, human rights, safety, and sustainable and equitable development. Sharing power can be challenging, particularly in the context of such an apparently polarised and rapidly changing topic as AI. But trust in AI and public legitimacy can only be built through listening to and learning from people, by embedding participation into the new decision-making spaces of AI governance.

Our aim is that the evidence and ideas presented in this report leave you excited about co-creating a citizens' track, and we invite you to get in touch to help build a more participatory AI future: contact@citizens-track.org

²⁹ PAIRS has been delivered annually on a budget between \$15k and \$50k. Single-country or thematic distributed dialogue processes can take place with seed funding from \$50k - \$250k; a lightweight secretariat might cost \$250k - £500k/year; a flagship global citizens assembly on AI could be delivered for \$1.5 - \$3m. A fund of \$10m/year dedicated to democratic AI could deliver a step-change in practice across the world. This year, tech firms are expected to collectively invest over \$650 billion in AI infrastructure alone.

Bibliography

- Bächtiger, André, and John S. Dryzek. 2024. *Deliberative Democracy for Diabolical Times*. Cambridge University Press.
- Carmichael, Matt. 2026. Ipsos AI Monitor 2026. Ipsos. <https://www.ipsos.com/en/global-attitudes-ai-2026-wonder-vs-worry-divide-deepens>.
- Chwalisz, Claudia. 2019. *A New Wave of Deliberative Democracy*. Carnegie Europe.
- Curato, Nicole, Pierrick Chalaye, Wendy Conway-Lamb, et al. 2023. *Global Assembly on the Climate and Ecological Crisis Evaluation Report*. University of Canberra.
- Curato, Nicole, John S. Dryzek, David Levai, Claire Mellier, Melisa Ross, and Rich Wilson. 2025. 'Where next for Global Climate Deliberation: From Proof of Concept to a Role in Transformation'. *Earth System Governance* 25 (August): 100280. <https://doi.org/10.1016/j.esq.2025.100280>.
- Davidian, Anne, and Laurent Jeanpierre, eds. 2023. *What Makes an Assembly?: Stories, Experiments, and Inquiries*. Sternberg Press.
- Davies, Tim. 2024. *Involving the Public in AI Policymaking - Experiences from the People's Panel on AI*. Connected by Data. [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).
- Davies, Tim, Claire Mellier, Kristophina Shilongo, and Rich Wilson. 2024. *Global Citizen Deliberation on Artificial Intelligence: Options and Design Considerations*. Connected by Data and ISWE.
- Dryzek, John S., André Bächtiger, and Karolina Milewicz. 2011. 'Toward a Deliberative Global Citizens' Assembly'. *Global Policy* 2 (1): 33–42. <https://doi.org/10.1111/j.1758-5899.2010.00052.x>.
- Fishkin, James, Nikhil Garg, Lodewijk Gelauff, et al. 2019. 'Deliberative Democracy with the Online Deliberation Platform'. Paper presented at 7th AAAI Conference on Human Computation and Crowdsourcing. HCOMP 2019. <https://doi.org/https://doi.org/10.25740/cz055qj0441>.
- Gelauff, Lodewijk, Liubov Nikolenko, Sukolsak Sakshuwong, et al. 2023. 'Achieving Parity with Human Moderators: A Self-Moderating Platform for Online Deliberation'. In *The Routledge Handbook of Collective Intelligence for Democracy and Governance*. Routledge.
- Gem, Ryan, and Maria Milosh. 2026. 'Citizen's Assemblies - Impact Data Research Brief 2025'. <https://www.peoplepowered.org/research-data>.
- Gillespie, Nicole, Steven Lockey, Tabi Ward, Alexandria Macdade, and Gerard Hassed. 2025. *Trust, Attitudes and Use of Artificial Intelligence: A Global Study 2025*. The University of Melbourne. 4974511 Bytes. <https://doi.org/10.26188/28822919>.
- 'Governing AI for humanity: final report'. 2024. With UN. Advisory Body on Artificial Intelligence. UN. <https://digitallibrary.un.org/record/4062495>.

- Hagelskamp, Carolin, Josh Lerner, and Nikhil Kumar (People Powered). 2021. 'Introducing a Global Theory of Change for Participatory Budgeting'. October 4. <https://www.peoplepowered.org/news-content/introducing-a-global-theory-of-change-for-participatory-budgeting>.
- Joint Secretariat of the Global Dialogue on AI Governance. 2026. 'Insights from Written Submissions'. Global Dialogue on AI Governance, Global Dialogue on AI Governance. <https://www.un.org/global-dialogue-ai-governance/en/insights-written-submissions>.
- Kabeer, Naila. 2005. *Inclusive Citizenship: Meanings and Expressions*. Zed Books.
- Marcucci, Sara, Natalia González Alarcón, and Stefaan Verhulst. 2025. 'AI Localism: The Emergence of AI Governance at the Local Level'. In *Handbook of Platform Urbanism*, edited by Annelien Smets and Pieter Ballon. Edward Elgar Publishing. <https://doi.org/10.4337/9781035313761.00020>.
- McKinney, Sammy. 2024. 'Integrating Artificial Intelligence into Citizens' Assemblies: Benefits, Concerns and Future Pathways'. *Journal of Deliberative Democracy* 20 (1). <https://doi.org/10.16997/jdd.1556>.
- McKinney, Sammy, and Claudia Chwalisz. 2025. 'Five Dimensions of Scaling Democratic Deliberation: With and Beyond AI'. <https://www.demnext.org/uploads/DemocracyNext-Five-dimensions-of-scaling-democratic-deliberation-paper-June-2025.pdf>.
- Melamed, Claire, and Paul Ladd. 2026. *People, Power, and AI: Rethinking Development for the New Era*. United Nations Foundation. <https://unfoundation.org/what-we-do/issues/ai-data-and-digital-cooperation/working-together-to-build-a-positive-digital-future/people-power-and-ai-rethinking-development-for-the-new-era/>.
- Michael Coppedge et al. 2026. 'V-Dem Dataset V16'. *Varieties of Democracy (V-Dem) Project*. <https://doi.org/10.23696/VDEMDS26>.
- Mikhaylovskaya, Anna. 2024. 'Enhancing Deliberation with Digital Democratic Innovations'. *Philosophy & Technology* 37 (1): 3. <https://doi.org/10.1007/s13347-023-00692-x>.
- Niemeyer, Simon, Francesco Veri, John S. Dryzek, and André Bächtiger. 2024. 'How Deliberation Happens: Enabling Deliberative Reason'. *American Political Science Review* 118 (1): 345–62. <https://doi.org/10.1017/S0003055423000023>.
- OECD. 2020. *Innovative Citizen Participation and New Democratic Institutions: Catching the Deliberative Wave*. Organisation for Economic Co-operation and Development. https://www.oecd-ilibrary.org/governance/innovative-citizen-participation-and-new-democratic-institutions_339306da-en.
- Oliveira, Marcele. 2025. 'COP30: We Need Citizens' Assemblies for Climate Negotiations'. *openDemocracy*, November 20. <https://www.opendemocracy.net/en/cop30-climate-change-people-seat-negotiations-climate-justice/>.
- Oman, Susan, and Sara Cannizzarao. 2025. *People's Feelings about AI: What the Evidence Tells Us*. The University of Sheffield. <https://digitalgood.net/dg-outputs/peoples-feelings-about-ai/>.

- Oman, Susan, Emily Rempel, Erin McCloskey, Octavia Field Reid, Tim Davies, and Canning Malkin. 2026. PAVE (Public AI Voice and Engagement) - Case Book Rationale. The University of Sheffield. 116347 Bytes. <https://doi.org/10.15131/SHEF.DATA.32840456>.
- Patra, Evelyn Asio, Viola Babirye, Francis Ssekitto, et al. 2026. 'From Users to Governors: Integrating Ugandan Women's Lived Experiences into Participatory Governance'. Paper presented at Participatory AI Research and Practice Symposium, New Delhi. PAIRS 2026. <https://www.pairs.site/From-Users-to-Governors-integrating-Ugandan-Women-s-Lived-Experiences-into-Participatory-Governance-2fd260e24e1a8134b0a7e6a621c30194>.
- Reuchamps, Min, Julien Vrydag, and Yanina Welp, eds. 2023. De Gruyter Handbook of Citizens' Assemblies. De Gruyter. <https://doi.org/10.1515/9783110758269>.
- Schüür, Friederike, Michael Birtwistle, Octavia Field Reid, Lara Groves, and Cam Rincon. 2026. 'The UN Global Dialogue on AI Governance: The Case for a Global AI Governance Floor'. Ada Lovelace Blog, May. <https://www.adalovelaceinstitute.org/blog/the-case-for-a-global-ai-governance-floor/>.
- Smith, Graham. 2009. Democratic Innovations: Designing Institutions for Citizen Participation. Cambridge University Press.
- Stilgoe, Jack. 2024. 'AI Has a Democracy Problem. Citizens' Assemblies Can Help.' Science 385 (6711): eadr6713. <https://doi.org/10.1126/science.adr6713>.
- Verhulst, Stefaan G. 2024. 'Beyond Checking a Box: How a Social Licence Can Help Communities Benefit from Data Reuse and AI'. The Conversation, December 15. <https://doi.org/10.64628/AAK.53kxefqu5>.
- Verhulst, Stefaan, Andrew J. Zahuranec, and Adam Zable. 2025. 'Participatory Approaches to Responsible Data Reuse and Establishing a Social License'. In Global Public Goods Communication: Mapping Actors, Policies, and Narratives, edited by Sónia Pedro Sebastião and Anne-Marie Cotton. Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-90667-1_10.
- Warren, Mark E. 2025. 'Democratic Innovation and Representative Democracy'. Perspectives on Politics 23 (1): 7–14. <https://doi.org/10.1017/S1537592724002111>.

People in the lead

**CONNECTED
BY DATA**

i-we



Public AI
Voice and
Engagement



**University of
Sheffield**

Centre for
Machine
Intelligence