

Deliberative Muscles & AI

*How Can DelibTech
Strengthen Citizens'
Capacities to
Deliberate?*

9 July 2026

Claudia Chwalisz, Founder and CEO,
DemocracyNext

Sammy McKinney, AI and Deliberation
Fellow, DemocracyNext, and PhD
Candidate, University of Cambridge

Jorim Theuns, Co-Founder, Dembrane

Eugene Yi, Impact Research Fellow,
DemocracyNext, and PhD Candidate,
University of Oxford

DemocracyNext

About DemocracyNext

We believe in a more just, joyful, and collaborative future, where everyone has meaningful power to shape their societies. DemocracyNext is an international research & action institute focused on broadening who has power and improving collective decision making with sortition, deliberation, and rotation.

www.demnext.org

About the co-authors

Claudia Chwalisz is Founder and CEO of DemocracyNext and Co-Lead of the Deliberation & Technology (DelibTech) Network that is co-convened by DemocracyNext and the AI & Democracy Foundation. Previously, Claudia led the OECD's work on innovative citizen participation from 2018–2022, where she developed the Deliberative Democracy Toolbox and co-authored key reports and standards. Claudia played a central role in designing the world's first permanent citizens' assemblies and has advised governments and institutions globally on deliberative processes. She is an Obama Leader, and serves on advisory boards including the UN Democracy Fund, The Data Tank, and MIT's Center for Constructive Communication. She is also the author of *The Populist Signal* and *The People's Verdict*.

Sammy McKinney is a PhD candidate at the University of Cambridge and an AI and Deliberation Fellow at DemocracyNext. He co-leads the Deliberation & Technology (DelibTech) Network. His PhD research critically explores the integration of artificial intelligence into processes of democratic innovation, especially citizens' assemblies. This research expands on his master's thesis carried out at the University of Edinburgh, which he published in an adapted form in the *Journal of Deliberative Democracy*.

Jorim Theuns is the CEO and co-founder of Dembrane, a civic tech initiative dedicated to transforming how democratic participation works in everyday life. At the heart of Dembrane's work is the belief that democracy should be more than a political event every few years; it should be part of how we govern our cities, design public services, and make collective decisions.

Eugene Yi is an Impact Research Fellow at DemocracyNext, a PhD Candidate at the University of Oxford, a Visiting Researcher at MIT, and Director of Product Management, AI & Creation Tools at Epic Games. He is an experienced tech entrepreneur, product and media executive, and policy expert. Previously, he co-founded Cortico to help build tech-enhanced deliberative systems for democracy. His experiences range from leading Asia policy at Twitter from 2012 to 2016, and he was previously the Head of Product at Hawkfish, a digital technology and data firm founded by Mike Bloomberg. He served in the U.S. Department of Defense and was a U.S. diplomat early in his career.

How to cite this paper

Chwalisz, Claudia, Sammy McKinney, Jorim Theuns, and Eugene Yi (2026). "Deliberative Muscles & AI: How Can DelibTech Strengthen Citizens' Capacities to Deliberate?", DemocracyNext.

Illustrations are by [Adèle Vivet](#).

Acknowledgements

We are grateful for the constructive feedback and comments that we received in two rounds on earlier versions of this draft, in writing and during two virtual feedback sessions, in May and June 2026 from a wide diversity of democracy practitioners and facilitators, technologists, academics, public servants, and civil society organisers across the globe. Many thanks to:

- Matt Abrams (DemocracyNext)
- Shutaro Aoyama (Plural Reality)
- Pete Bryant (Shared Future)
- Josh Burgess (DemocracyNext, Central Oregon Civic Action Project)
- Wes Chow (MIT)
- Oliver Escobar (University of Edinburgh)
- Daniel Fusca (City of Toronto)
- Lisa Gutermuth (Mozilla Foundation)
- Denise Hearn (The Long Now Foundation)
- Nicole Hunter (Mosaic Lab)
- Luke Kemp (University of Cambridge)
- Scott Lappan-Newton (Gauge Consulting)
- Lawrence Lessig (Harvard University)
- Gideon Lichfield (Harvard University, BlueDot Media)
- Mauricio Mejia (Democracy advisor)
- Kelly McBride (Involve)
- Aviv Ovadya (AI & Democracy Foundation)
- Lex Paulson (UM6P)
- Kyle Redman (AI & Democracy Foundation)
- Lucy Reid (DemocracyNext)
- Nathan Schneider (University of Colorado Boulder)
- Samantha Shireman (Harvard University)
- Andrew Sorota (Office of Eric Schmidt)
- Robbie Stamp (DemocracyNext)
- Jane Suiter (Dublin City University)
- Matt Stempeck (Evens Foundation, Civic Tech Field Guide)
- Alistair Stoddart (Scottish Parliament)
- Audrey Tang (Taiwan's Cyber Ambassador, DemocracyNext)
- Matthew Victor (Partners in Democracy, Harvard University)

Table of contents

Executive summary	5
1. Introduction	9
2. Why deliberative muscles matter	14
3. Seven deliberative muscles	19
3.1. The self-reflection muscle	23
3.2. The reasoning muscle	26
3.3. The dialogue muscle	28
3.4. The vulnerability muscle	30
3.5. The collaboration muscle	33
3.6. The imagination muscle	35
3.7. The facilitation muscle	37
4. Investing in civic infrastructure and civic AI	40
5. Three future trajectories to advance the deliberative muscles approach	44
5.1. Exploring and measuring how different process designs affect deliberative muscles	46
5.2. Communication guidance about using AI in deliberative processes	47
5.3. Exploring the political economy dimensions of AI in deliberation	48
6. Conclusion	49
7. References	52

Executive summary

Democracy is under strain, and one of the most promising responses to that strain is the growing global movement around deliberative assemblies: citizens' assemblies, citizens' juries, and related forums that bring randomly selected, broadly representative groups of people together to weigh evidence, listen to one another, and make shared decisions on complex public issues. Over 1,000 such processes have now been run worldwide, and a growing body of evidence suggests they depolarise opinion, generate well-reasoned recommendations, build trust, and reconnect people to political life.

However, these processes are also resource-intensive, slow, and hard to scale, and have thus become a site of intense interest for AI integration. The pitch from many technologists, practitioners, and funders is consistent: AI can make deliberation cheaper, faster, more accessible, and more scalable.

In this paper, we argue that AI, when designed with care, can indeed play a powerful role in strengthening deliberation. But the very efficiencies that make AI attractive also risk undermining what deliberation is for in the first place. Whether AI strengthens or weakens deliberation is not predetermined, however; it is a matter of design.

Our starting point is that deliberative assemblies are not decision-making machines whose sole value lies in the recommendation they produce. They are also spaces in which participants exercise and develop the civic capacities that democratic life depends upon. If we automate too much, we may end up with smoother processes that hollow out the productive friction that makes them valuable, while simultaneously reducing people's ability to participate in democratic life.

These considerations are relevant to all places where deliberation takes place – workplaces, schools and universities, museums, financial institutions, corporations and cooperatives, membership-based associations, and other organisations.

We make three contributions.

First, **we argue that one of the most important and most overlooked virtues of deliberative assemblies is that they build deliberative muscles:** the cognitive, dispositional, and relational capacities that citizens need to do the work of democracy together. We use the language of muscle deliberately. A muscle is not an idea one holds; it is a capacity one maintains through practice, that weakens when unused, and improves when trained.

Second, **we offer a typology of seven deliberative muscles:** **self-reflection** (examining one's own values and beliefs), **reasoning** (engaging critically with evidence and expertise), **dialogue** (listening attentively, responding, and giving reasons), **vulnerability** (sharing feelings and reflections, tolerating conflict, feeling the weight of others' experiences), **collaboration** (moving from individual reasoning to shared judgement), **imagination** (envisioning futures and alternatives concretely enough to deliberate about them), and **facilitation** (guiding small-group deliberation productively and inclusively).

For each muscle, we identify the practical challenges that make practitioners reach towards AI, and we map AI use cases along a competitive–complementary spectrum drawn from David Krakauer's distinction between cognitive artefacts that strengthen underlying human capacities and those that substitute for them. Context matters throughout: the same tool can strengthen a muscle in one setting and substitute for it in another.

Third, **we argue that the exercise of these seven muscles, sustained across many people and many processes over time, is what produces the societal capacity that resilient democratic systems require.** This has implications for civic infrastructure: for the training, standards, and ecosystem-building bodies that sustain deliberative practice; for how we think about AI ownership, provenance, and dependency; and for the case for a civic AI future in which communities are not merely end-users of DelibTech but central to the conditions that make it work at all.



Overview table: Seven deliberative muscles & AI impacts

Muscle	Description	Competitive AI use	Complementary AI use
1. Self-reflection	Examining one's own values and beliefs before, during, and after a deliberative exchange, with openness to changing one's mind.	Chatbots that infer a participant's position from demographic data, or "AI agent" systems trained on stated preferences to represent them in simulated deliberations.	Structured pre-deliberation prompts (e.g. Plural Reality's Cartographer, Revel and Pénigaud's AI Reflector) that ask participants to articulate their own values and surface tensions in their own reasoning.
2. Reasoning	Engaging critically with evidence and expertise; evaluating competing claims; updating beliefs when warranted.	AI that presents pre-digested conclusions, unilaterally identifies knowledge gaps, or weighs competing perspectives on citizens' behalf.	AI that transforms dense information packs into multi-modal formats (podcasts, videos, images), or that filters and prioritises citizen-generated questions for experts.
3. Dialogue	Listening attentively, responding, taking turns, and providing reasons for one's claims.	AI that drafts a participant's contribution for them, producing polished statements presented as their own; "AI agents" that exchange arguments in place of participants.	Real-time multilingual translation and live captioning that let participants engage directly with the substance and texture of others' contributions; private prompts that ask participants what value or experience they are drawing on.
4. Vulnerability	Courageously sharing one's feelings and reflections; tolerating conflict; feeling the weight of others' experiences.	"F-bomb translators" that smooth harsh language; AI-simulated stakeholder personas replacing lived-experience testimony; consensus generators that synthesise across views without human encounter.	AI that surfaces perspectives the group has passed over, flags where an emerging consensus has not engaged with a strong objection, or helps participants articulate the values beneath their disagreement.
5. Collaboration	Moving from individual reasoning to shared judgement; experiencing collective ownership of an outcome through sensemaking together.	Handing transcripts to a language model and presenting the resulting thematic summary back to the group as "what you said" – creating what practitioners describe as a "layer of plastic wrap" between people and their own outputs.	Facilitator-reviewed first-pass summaries and translations that give the group a shared starting point to revise together; argument-embedding visualisations that let a group explore the shape of its own thinking.
6. Imagination	Envisioning futures, counterfactuals, and alternative arrangements concretely enough to deliberate about them.	Asking a model what the future could look like and deliberating about the answer; using AI as an oracle that narrows the space of imagined futures while presenting the narrowing as breadth.	AI as a queryable index of prior attempts, precedents, and analogous projects; complex-systems dynamics modelling so participants can see potential impacts of a proposed change; Socratic tools that return questions rather than answers.
7. Facilitation	Guiding small-group deliberations productively and inclusively; navigating conflict; building the capacity of others to do the same.	AI moderators that decide who speaks when and enforce rigid speaking-time allocations regardless of context, emotion, or communication style.	Conversation-guide drafting tools; AI transcription that frees a facilitator to attend to the group; private nudges to novice or citizen facilitators that build their capacity over time.

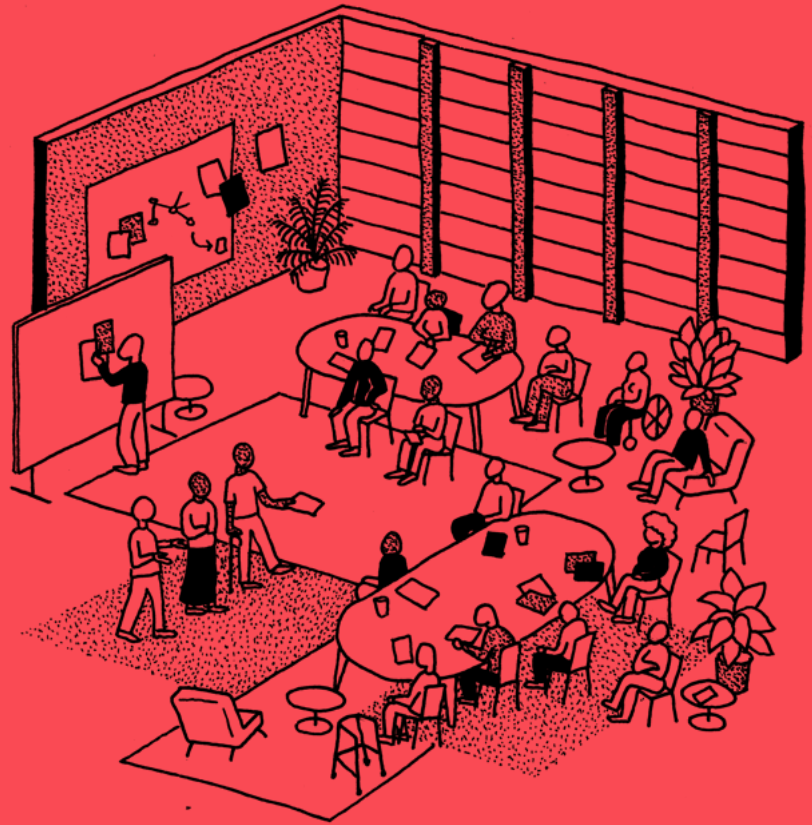
The central design question for anyone building, commissioning, or funding deliberative technology is therefore not whether AI produces a recommendation faster, but whether it leaves participants, facilitators, and the wider institutional ecosystem more capable at the end of the process than at the start.

The stakes are high: will we develop and deploy technologies that strengthen our individual and collective capacity to enable democratic flourishing, or will we hollow out those capacities through technological shortcuts? The answer will depend on choices made by many actors:

- **Commissioners and practitioners** should resist easy efficiency gains that bypass hard work and capacity building.
- **Public actors** should invest in digital public infrastructure and civic AI that distributes power rather than concentrating it.
- **Technologists** should reflect on how their tools shape the capacities of the people who use them.
- **Scholars** should support the mapping of deliberative muscles and how technology affects their cultivation.
- **Communities and civil society** should demand and design deliberative processes that sustain and strengthen the muscles democracy needs to be resilient.

These structural pressures, however, require systemic change and support, which is why we also focused on the investment needed in civic infrastructure and civic AI if we are to enable conditions that make deliberative technology worth using at all. Efforts across these multifaceted domains will support us in navigating one of the most pressing challenges facing humanity: building a world where technological innovation and democratic innovation work in harmony.





CHAPTER 01

Introduction

Democracy is under strain. Polarisation is hardening, trust in institutions is eroding, and global challenges from ecological collapse to mass migration are outpacing the capacity of our political systems to respond. Part of why many people are frustrated is that problems they care about are not being addressed, and they feel a lack of agency to change this. In our view, the lack of responsiveness is largely because decision-making processes are broken – people in power are hugely unrepresentative of, and therefore disconnected from, the problems that many people face, and centralised, top-down, winner-take-all debates dominate.

In the face of this, a growing movement of social entrepreneurs, civil servants, and democratic innovators has been advancing one promising response: deliberative democratic innovations such as citizens' assemblies and citizens' juries. These deliberative forums bring randomly-selected, broadly representative groups of people together to weigh evidence, listen to one another, and work towards shared decisions on complex public issues. In other words, they broaden who has power, and improve how people decide together. These considerations are relevant to all places where deliberation takes place – workplaces, schools and universities, museums, financial institutions, corporations and cooperatives, membership-based associations, and other organisations.

These processes have shown real promise as they can promote trustworthy, legitimate outcomes, while also strengthening people's capacities as citizens for a more resilient, anti-fragile society. There are over 1,000 of them worldwide to date – used to inform decision making by governments and other organisations like schools and universities, museums, workplaces, banks, pension funds, and others (OECD, 2020; Buergererrat, 2026, DemocracyNext, forthcoming). Research from across decades of practice suggests they can depolarise opinion, generate well-reasoned policy recommendations, build trust, and reconnect people to political life (Reuchamps, Vrydagh, and Welp, 2023; Fishkin et al., 2024).

However, they are also resource-intensive, slow, and difficult to scale. The exponential growth of the capabilities of deep neural networks like Large Language Models (LLMs) have led to a growing interest in what this new wave of Artificial Intelligence (AI) might contribute to addressing these challenges. We will use "AI" as shorthand to refer to this class of generative deep neural networks.

Growing numbers of people are interested in how (emerging) technologies can enhance deliberative assemblies. Practitioners are experimenting with AI for transcription, translation, summarisation, and to help create multi-modal learning-materials. Technologists are building AI applications that promise to facilitate small-group dialogue, simulate absent stakeholders, or generate consensus statements. International networks connecting actors across these ecosystems are being established. Funders are excited. The pitch is consistent: AI can make deliberation cheaper, faster, more accessible, and more scalable.

At the same time, public authorities and regulators are lagging behind – lacking understanding about how and if technology should support deliberative processes, with procurement procedures making experimentation harder. There is a [backlash](#) growing against AI in some parts of the world, especially in Europe and North America – with AI perceived as extractive, disempowering, and exploitative (Ipsos, 2025). We must take this AI backlash happening in some places seriously, as introducing technology in contexts where there is high distrust could undermine trust in the whole deliberative process itself. We have seen this with electronic voting – the technology exists; the barriers to its adoption are not technical, but infrastructural, cultural and trust-based.

In this paper, we argue that, when designed with care and intentionality, AI can play a powerful role in strengthening the goals of deliberation. The catch is that the very efficiencies that make AI attractive can undermine those goals in the first place if we are not deliberate about how we use them.

We also argue that deliberative processes are not decision-making machines whose sole value lies in the end result. They are also spaces in which participants exercise and develop the civic capacities that democracy depends upon. The trustworthiness, rigour, and legitimacy of a decision matter, and our hypothesis is that these are stronger when a deliberative process engages and strengthens people's deliberative capabilities. If we automate too much, we may end up with smoother processes that hollow out the hard work and productive friction which make them valuable, while simultaneously disempowering people by reducing their democratic capabilities and their ability to participate in democratic life.

We make three contributions to this debate:

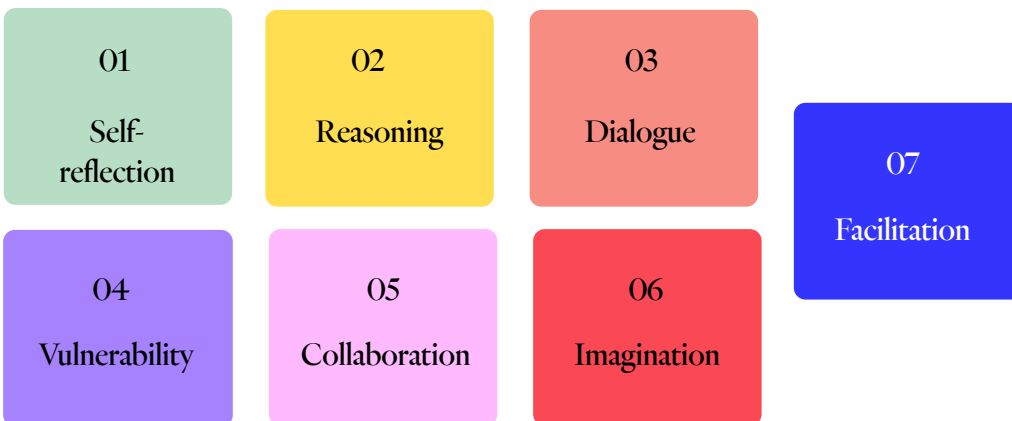
First, **we argue that one of the most important virtues of deliberative assemblies is that they build deliberative muscles:** the cognitive, dispositional, and relational capacities that citizens need to do the work of democracy together.

Second, **we offer a typology of seven deliberative muscles – self-reflection, reasoning, dialogue, vulnerability, collaboration, imagination, and facilitation** – and show, for each, where AI integration risks substituting for the muscle versus where it may be able to strengthen it. Whether AI supports or diminishes deliberation is therefore not predetermined; it depends on how, where, and why it is introduced.

Third, **we argue that the exercise of these seven muscles, sustained over time and across many people, contributes to the civic societal capacity that resilient democratic systems require.**

Our overarching argument is that **a central design criterion for deliberative technologies is that they should strengthen citizens’ deliberative capacities, rather than weaken them.** In other words, deliberative technologies (DelibTech) should be sustaining and developing citizens’ ability to do the work of citizenship together, strengthening their democratic deliberation muscles.


Figure 1. Seven deliberative muscles



In previous DemocracyNext research, McKinney and Chwalisz (2025) [explored five ways that deliberative assemblies can scale](#): scaling out (more participants), scaling up (higher governance levels), scaling across (more processes), scaling deep (increasing impact) and scaling in (improving deliberative quality). They identified different ways that AI could support these dimensions of scaling and argued that technological innovation needs to be combined with robust civic infrastructure building if promising scaling pathways are to materialise. **The scaling lens and the deliberative muscles approach are synergistic: in combination, they provide a vocabulary for reflecting on how different approaches to scaling affect deliberative muscle building.** As discussed in section five, future research and practice should grapple with the way different scaling approaches affect the exercise of deliberative muscles. In this paper, our focus is more modest: to introduce the deliberative muscles framework and establish its importance for thinking about AI's role in deliberative assemblies.

This paper is for people with the authority and responsibility of deciding how technology is going to be deployed in a deliberative process – commissioning authorities in government and other organisations, as well as democracy practitioners, designers, and facilitators. We are also writing for technologists deciding what to build, for funders deciding what to back, and for scholars mapping a fast-moving field.

We want to be explicit that we are not addressing infrastructural questions, such as compute sovereignty, environmental cost, or supply chains (although in section four, we briefly argue that these considerations only strengthen the case for the deliberative muscles approach). We are also not addressing the broader landscape of AI in democracy beyond deliberative practice. In the conclusion, we outline three trajectories that could extend the deliberative muscles approach in future research and practice.



CHAPTER 02

Why deliberative muscles matter

Asking what deliberative assemblies are for is a question with numerous answers. Some advocates emphasise their epistemic potential, meaning that diverse, informed groups can reach better-reasoned conclusions than narrower expert panels. Others emphasise their legitimacy-generating function – they give voice to citizens who would otherwise be absent from decisions affecting them. Others still focus on policy impact, or the transformative effects that participating in a deliberative process has on the people involved, awakening a sense of agency.

In this paper, we focus on an **important virtue of deliberative democratic innovations that often gets overlooked: they build citizens' deliberative muscles. In other words, they build the individual and collective skills and capacities that enable high quality deliberation to thrive, both inside and outside formal decision-making spaces.** Strong deliberative muscles are helpful to citizens at an individual level, and there are collective societal benefits of having a citizenry with strong deliberation muscles. We argue that this is essential (though not the only thing needed) for bolstering democratic resilience.

Deliberation has three main requirements: it requires citizens to provide reasons and justifications for their claims regarding collective action, it requires citizens to listen attentively and respectfully to others, and it requires the outcome reached to be a shared decision. Deliberation therefore requires that people exercise a variety of skills, including critical thinking, open-mindedness, listening, reason-giving, and the willingness to work through disagreement with others. Hosting [high quality deliberative processes](#) also requires institutional capacities, such as engaging diverse stakeholders, assimilating expertise on complex issues, facilitating inclusion, communicating effectively with the public, and much more (OECD, 2020).

Deliberation has three main requirements: it requires citizens to provide reasons and justifications for their claims regarding collective action, it requires citizens to listen attentively and respectfully to others, and it requires the outcome reached to be a shared decision. Deliberation therefore requires that people exercise a variety of skills, including critical thinking, open-mindedness, listening, reason-giving, and the willingness to work through disagreement with others. Hosting high quality deliberative processes also requires institutional capacities, such as engaging diverse stakeholders, assimilating expertise on complex issues, facilitating inclusion, communicating effectively with the public, and much more (OECD, 2020).

Why is maintaining and strengthening these capacities important? **We argue that for a resilient democratic society, strong deliberative muscles are not a nice-to-have; they are a necessity.** In a democratic climate increasingly defined by polarisation, manipulation, and hostility, deliberative assemblies create unique spaces that allow people to cultivate skills at the heart of democratic practice: reason-giving, listening, self-reflection, curiosity, sitting with disagreement, updating one's views in light of new information, imagining new possibilities, moving from individual preference to shared judgement. They atrophy when they are not exercised. **Engaging deliberative muscles can help awaken or strengthen people's sense of agency** – the feeling that they can, indeed, shape their future and influence change.

One of the main goals of [institutionalising and scaling democratic innovations](#) is to enable a virtuous cycle that promotes a society of resilient citizens who can build and reimagine new democratic institutions at a time of rising authoritarianism and institutional incapacity to grapple with planetary-scale challenges (McKinney and Chwalisz, 2025). Arguably, institutional and governance imagination in the decades to come will be more important than ever.

From this angle, we reflect on the role of AI and emerging technologies in deliberative democratic innovations. Not "does this technology produce a recommendation faster?" but "does this technology leave people more or less capable of deliberating well?"

We use the language of 'muscle' deliberately. A muscle is not an idea one holds; it is a capacity one maintains through practice; it weakens when unused and it is improved by being actively trained and strained. Skills and dispositions need exercise, and **a key design question for deliberative processes and technologies is whether they offer participants more or less opportunities to flex their deliberative muscles.**

How and whether these muscles are strengthened or weakened is determined by the tools that we build as well as the cultures and institutions that we scaffold around them.



To evaluate which AI integrations strengthen deliberative muscles and which substitute for them, **we draw on David Krakauer's (2016) distinction between complementary and competitive cognitive artefacts.** The distinction is empirical: something is complementary if the underlying human capacity persists when the tool or technique is removed, as arithmetic does after training with an abacus, or physical strength after weightlifting. If capacity degrades when the tool or technique is removed, it is competitive. As many technologies have not been subject to this empirical scrutiny, we must reason deductively about which way a tool is likely to tilt, and do so throughout this paper.

With AI capabilities proliferating, we see the temptation of people turning to competitive cognitive tools that might make deliberation seem, or indeed be, faster, easier, cheaper – removing some friction and hard work. But at what cost? There is emerging evidence and growing discussion of ‘cognitive debt’ and other potential negative cognitive effects of relying on AI tools on reasoning and critical thinking (Kosmyna et al., 2025, Lee et al., 2025; Gerlich, 2025). As an emerging technology, the body of evidence is not yet fully established, but it is plausible. Much of the focus to date has been on AI contributing to cognitive decline in individuals, whereas in this paper we focus on both individual as well as collective and interpersonal democratic processes that rely also on collective cognitive capabilities.

Not every automation carries the same cost to the muscle, and it is worth distinguishing three cases. For some tasks, automation is close to free for the capacity at stake: transcription, translation, captioning, and scheduling remove time-consuming tasks without removing any deliberative work a participant would otherwise grow from. For others, the cost is real but may be worth paying under constraint: a first-pass summary or a filtered question list does some of the work participants could have done, but returns it to them to interrogate, trading a little muscle building for reach or speed. For a third class, automation destroys the good it claims to deliver: **outsourcing sensemaking, facilitation, or the affective work of encountering another person removes the very activity that makes the exercise deliberation.** The hard cases sit at the boundary. If a tool yields a measurably better recommendation, and participants still grasp why, but they could not have produced it themselves, under what conditions do we accept that trade off? We do not think there is a single answer, but the muscle lens helps a commissioner ask the question precisely.


The extreme example is platforms, such as [Artificial Societies](#), that treat deliberation as a machine for producing an opinion distribution: they build a synthetic population from behavioural data, query it, and recover what a public "would" think faster and cheaper, without anyone in a room. If getting to a recommendation quickly was the only thing of value, some people would see this as a strict improvement.

The muscle lens explains why it is not – a synthetic population learns nothing, listens to no one, and leaves no citizen more capable than before. The recommendation coming out of a deliberation is not simply an average of the perspectives people came in with, which is largely what would be expected from a synthetic population. Rather, it's the result of them gaining a more thorough understanding of each perspective, using that to address the weaknesses of various solutions, and coming up with an output greater than the sum of the inputs. Furthermore, deliberation arguably necessitates people actually using a minimum of some combination of their deliberative muscles; if no muscles are engaged, it ceases to even be a deliberation and becomes something else entirely.

Against this backdrop, we are asking ourselves: what kind of technology would be complementary rather than competitive to our deliberative muscles, and thereby support the virtuous cycle of democracy?

We also recognise that a tool may be cognitively competitive but at the same time open up new skills. For instance, a calculator might compete with mental arithmetic while unlocking higher-order mathematical reasoning. How might AI substitution of sensemaking, for instance, unlock other cognitive functions or democratic capability? We use the complementary-competitive distinction as a lens, recognising that a tool's impact is multifaceted, and encourage future research and practice to examine how outsourcing certain tasks may unlock other values that are democratically salient.

We also emphasise the central importance of context. When one looks at any particular capability of AI, the situation in which one is using it is very important to figuring out if that use of AI is a beneficial use; there will be different trade-offs and consequences for the same tool in different situations. This means that not all uses of AI that outsource our deliberative capacities are necessarily objectionable in all contexts. Given that one of the core goals of deliberative assemblies, however, is to cultivate these muscles, such uses should be approached with significant caution.



CHAPTER 03

Seven deliberative
muscles, and how AI
could help or hinder each



In this section, we expand upon seven deliberative muscles. For each, we offer a brief definition, identify the challenges that may make practitioners reach towards technology, and map various AI uses along the competitive-complementary spectrum. We provide real-world examples for competitive and complementary uses, when available, and otherwise we explain how AI could be used in those ways based on existing capabilities. The examples provided are not exhaustive; they are offered to illustrate the deliberative muscles framework and its relevance to practice.

The seven muscles are also **not exhaustive, and not always cleanly separable**; in practice, they reinforce one another and a single deliberative moment usually exercises several at once. We have chosen these seven because they generate concrete design questions and they map well to recognisable stages of deliberative practice, such as those outlined in DemocracyNext's [Assembly Guide](#) (2024), Mosaic Lab's [Facilitating Deliberation](#) book (2020), and deliberative facilitator training programmes like those run by We Do Democracy, Shared Future, and others. At the same time, **we hope that the deliberative muscles lens will open up to new process design explorations beyond what is typical without emerging technology.**

Some readers will be familiar with [Deb Roy](#) and [Audrey Tang's](#) references to civic muscles. Practitioners that know of Zakia Elvang's [Democracy Fitness](#) training will also recognise some overlaps (Life With Machines, 2025, Tang, 2026a, We Do Democracy, 2026). We share their perspectives that civic and democratic capacities are trainable. Our typology narrows the frame to deliberation specifically by emphasising the epistemic and judgement-forming dimensions that distinguish deliberative practice from democratic or civic participation more broadly, while folding several of the dispositional muscles (courage, compromise, activism) into deliberation-specific muscles. We came to these seven muscles drawing on years of practice in deliberative democracy, with feedback from an international and interdisciplinary group involving practitioners and academics.

Finally, we want to emphasise that **deliberation is a group activity. It entails individual and collective skills as well as cognitive capabilities**, and there is often an interplay between the two.



Overview of the seven deliberative muscles

01

Self-reflection

The capacity to examine one's own values and beliefs before, during, and after a deliberative exchange, with the openness to changing one's mind.

02

Reasoning

The capacity to engage critically with evidence and expertise.

03

Dialogue

The capacity to listen attentively, respond, and provide reasons for one's claims.

04

Vulnerability

The capacity to courageously share one's feelings and reflections, tolerate conflict, and feel the weight of others' experiences.

05

Collaboration

The capacity to move from individual reasoning to shared judgement, experiencing collective ownership of an outcome.

06

Imagination

The capacity to envision futures, counterfactuals, and alternative arrangements concretely enough to deliberate about them.

07

Facilitation

The capacity to guide small group deliberations productively and inclusively.

Table 1: Seven deliberative muscles & AI impacts — Overview

Muscle	Description	Competitive AI use	Complementary AI use
1. Self-reflection	Examining one's own values and beliefs before, during, and after a deliberative exchange, with openness to changing one's mind.	Chatbots that infer a participant's position from demographic data, or "AI agent" systems trained on stated preferences to represent them in simulated deliberations.	Structured pre-deliberation prompts (e.g. Plural Reality's Cartographer, Revel and Pénigaud's AI Reflector) that ask participants to articulate their own values and surface tensions in their own reasoning.
2. Reasoning	Engaging critically with evidence and expertise; evaluating competing claims; updating beliefs when warranted.	AI that presents pre-digested conclusions, unilaterally identifies knowledge gaps, or weighs competing perspectives on citizens' behalf.	AI that transforms dense information packs into multi-modal formats (podcasts, videos, images), or that filters and prioritises citizen-generated questions for experts.
3. Dialogue	Listening attentively, responding, taking turns, and providing reasons for one's claims.	AI that drafts a participant's contribution for them, producing polished statements presented as their own; "AI agents" that exchange arguments in place of participants.	Real-time multilingual translation and live captioning that let participants engage directly with the substance and texture of others' contributions; private prompts that ask participants what value or experience they are drawing on.
4. Vulnerability	Courageously sharing one's feelings and reflections; tolerating conflict; feeling the weight of others' experiences.	"F-bomb translators" that smooth harsh language; AI-simulated stakeholder personas replacing lived-experience testimony; consensus generators that synthesise across views without human encounter.	AI that surfaces perspectives the group has passed over, flags where an emerging consensus has not engaged with a strong objection, or helps participants articulate the values beneath their disagreement.
5. Collaboration	Moving from individual reasoning to shared judgement; experiencing collective ownership of an outcome through sensemaking together.	Handing transcripts to a language model and presenting the resulting thematic summary back to the group as "what you said" – creating what practitioners describe as a "layer of plastic wrap" between people and their own outputs.	Facilitator-reviewed first-pass summaries and translations that give the group a shared starting point to revise together; argument-embedding visualisations that let a group explore the shape of its own thinking.
6. Imagination	Envisioning futures, counterfactuals, and alternative arrangements concretely enough to deliberate about them.	Asking a model what the future could look like and deliberating about the answer; using AI as an oracle that narrows the space of imagined futures while presenting the narrowing as breadth.	AI as a queryable index of prior attempts, precedents, and analogous projects; complex-systems dynamics modelling so participants can see potential impacts of a proposed change; Socratic tools that return questions rather than answers.
7. Facilitation	Guiding small-group deliberations productively and inclusively; navigating conflict; building the capacity of others to do the same.	AI moderators that decide who speaks when and enforce rigid speaking-time allocations regardless of context, emotion, or communication style.	Conversation-guide drafting tools; AI transcription that frees a facilitator to attend to the group; private nudges to novice or citizen facilitators that build their capacity over time.

The self-reflection muscle

The capacity to examine one's own values and beliefs before, during, and after a deliberative exchange – to know what one thinks and why, to hold one's views provisionally enough to be moved by what one hears, to be open to changing one's mind, and to allow oneself to leave the room as a different person than who entered it.

The challenge is that people often arrive at a deliberative process having given little structured thought to the topic. They may have strong feelings without articulated reasons, or articulated positions that are inherited from a narrow information environment. Without a chance to surface and examine their own starting point, they enter the room less able to contribute substantively and recognise when their thinking has shifted.

A competitive use of AI here would be to have a chatbot tell participants what they probably think, or hand them a position generated from their demographics or other available data, effectively skipping the work of self-examination entirely. Some [“AI agent”](#) proposals, similar to the synthetic population example discussed above, do exactly this: training a model on a participant's stated preferences and then having it “represent” them in a simulated deliberation (Hidalgo, 2025; Gudiño, J.F., Grandi U. and Hidalgo, C., 2024). The participant never has to do the work of understanding, articulating and re-evaluating their own values and judgement; the work is outsourced to a system that approximates it. **Such “AI agent” proposals take away humans' possibility to flex any of the seven muscles, removing their agency and democratic skill-building potential, including that of self-reflection.** We won't repeat the example in every section, making the point once here.

A complementary use of AI is well-illustrated by tools like [Plural Reality's pre-deliberation prompts in their Cartographer app](#) (Aoyama, 2026), or features of Manon Revel and Théophile Pénigaud (2025) notion of an [“AI Reflector”](#): a system that does not tell the participant what they think but asks them, in a structured way, to articulate their own values, surface tensions in their own reasoning, and arrive at the deliberative room knowing where they stand and where they are uncertain. The technology is creating the conditions for the participant to reflect more rigorously than they would alone.

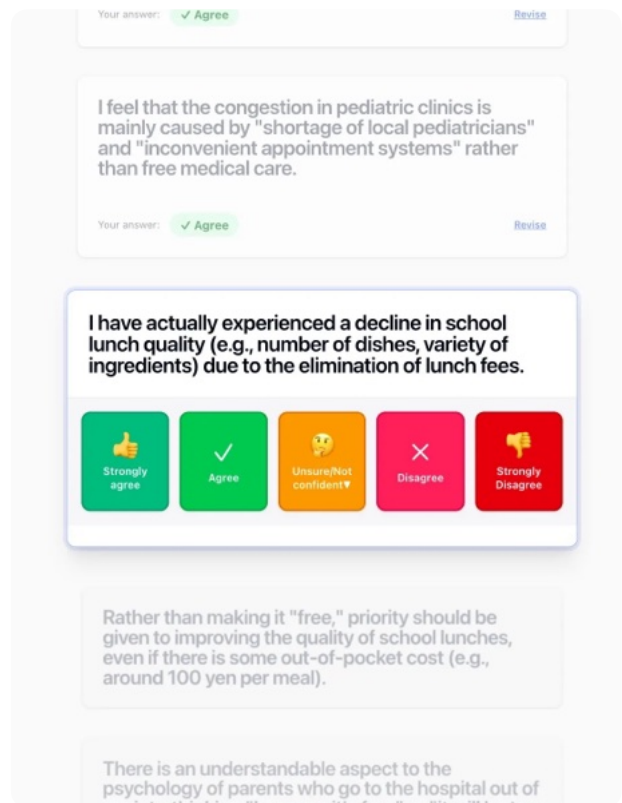
Self-reflection is more than just a deliberative warm-up, but is relevant throughout the whole deliberative process. For instance, Plural Reality's tool has been used at various stages through deliberative processes in Japan. After an expert speaker presents and answers questions, before assembly members deliberate in their small groups, they first individually interact with the Cartographer to reflect on what they think and what they do not yet understand. Those outputs are then shared collectively and reviewed as a group on a big screen at each table before beginning their conversation. This helps them begin with each individual having a more thought-through position and set of questions, but also helps to overcome another challenge – more closely related to the dialogue and facilitation muscles – of ensuring that everyone's voices can come into conversation equally and the most confident and outspoken perspectives do not end up dominating.

[Differences] Points of disagreement that should be explored through dialogue

While there is common ground, there are also differences in opinion and perspective on several important points. These differences bring a variety of viewpoints into focus and form the core issues for further discussion.

Difference	Analysis / Insight
1. To install air conditioners, should 'other budgets' be cut?	Regarding the urgency of installing air conditioners in the gymnasium (#4), opinions are divided between 'it should be implemented even if it means cutting budgets for other child-rearing programs' ([redacted] , [redacted]), and 'we should first consider how to secure funding' ([redacted]), and 'not that far...' ([redacted]). While there is shared recognition of the need to solve the problem, there is a gap in awareness regarding the realistic cost question of what should be sacrificed for it (trade-offs) .
2. Is responsibility for ensuring safety '100% government' or 'shared responsibility'?	In response to the question, "Should the responsibility for ensuring children's safety be borne entirely by the government?" (#12), everyone answered "strongly agree," yet their free-response comments differed significantly. <ul style="list-style-type: none"> [redacted]: "Entirely the government." [redacted]: "The government should bear the primary responsibility, but a 'shared responsibility' in which residents also cooperate is desirable." This is a very important difference. There is agreement that "the ultimate responsibility lies with the government," but it shows a difference in the degree to which residents are expected to be involved in the process leading to that point—how much they see it as their own issue .
3. Is 'behavioral change' by residents themselves necessary?	Regarding the idea that "solving the problem requires behavioral change by residents themselves" (#15), one redacted respondent strongly disagrees . This can be interpreted as a strong challenge to the idea that, before easily shifting to arguments of 'self-help' and 'mutual help,' we should first clarify the responsibilities that "public support" must fulfill. While the other three respondents acknowledge the need for behavioral change, this point of conflict is an essential perspective for deepening the discussion.
4. How incentives for participation in 'mutual help' are understood	Regarding incentives for park-cleaning activities (#9), everyone agrees that they "make participation easier," but they assign different weight to them. <ul style="list-style-type: none"> [redacted]: "Time and frequency" are more important than incentives. [redacted]: "If the situation is dangerous, I would want to help even without compensation." [redacted]: "They can provide an initial trigger, but different factors are needed for sustained participation." This suggests that not only simple financial motivation, but also more fundamental drivers—such as ease of participation (systems), a sense of urgency, and the meaning of the activity —are considered important.

Screenshot from Plural Reality's Cartographer tool. Credit: Shutaro Aoyama



Right-hand side: Deliberation in Japan, Photo credit: Shutaro Aoyama; Left-hand side: Screenshot from Plural Reality’s Cartographer app, Credit: Shutaro Aoyama.

The contextual caveat is that even the complementary version could potentially become competitive if it is over-scripted. A reflection tool that nudges users towards the same set of positions would elicit ‘self-examination’ that is just compliance with the tool’s frame. Aoyama, one of the developers of Plural Reality’s tool, has written a humble [reflection piece](#) (2026) assessing its limitations:

“Transparency addresses intentional bias; it doesn’t solve emergent framing effects. Question order matters. Summary language matters. Any system that mediates through language will shape thought to some degree, because language is not a neutral vessel.

The best name I have for the failure mode is this: ventriloquism with good UX. The system produces coherence on your behalf and hands it back to you as if it were self-expression. That is not empowerment. It is something much more dangerous, dressed in empowerment’s clothes.”

The reasoning muscle

The capacity to engage critically with evidence and expertise – to distinguish reliable knowledge from noise, to evaluate competing claims, to handle incomplete or contested information, to hold beliefs provisionally, and update them when warranted.

The challenge is that deliberations tend to address complex problems that necessitate engaging with expertise, considering multiple perspectives, questioning stakeholders, and digesting large amounts of information. Learning and grappling with complexity and trade-offs in a deliberative process is hard. It demands critical thinking and curiosity.

In current practice, information packs are often long, text-heavy, and more likely to be read by those with more time and higher levels of education. Citizens arrive with very different levels of background knowledge. Relevant experts and stakeholders are not always in the room, and when they are, they may struggle to make their contributions accessible and engaging to a diverse audience.

A competitive use of AI would be to outsource judgement, problem solving, and expertise to an AI system: AI presents the information, unilaterally identifies knowledge gaps, weighs up competing perspectives, and bypasses human stakeholder involvement. Using AI in these ways follows broader trends of ‘cognitive offloading’ to AI systems: rather than doing the hard work of engaging with complexity ourselves, we increasingly outsource it to AI. There is evidence that this atrophies our reasoning skills.

For example, it was found that using AI even for ten minutes to support tasks in mathematics and reading comprehension reduced performance when AI was taken away (Liu et al., 2026). If citizens participating in deliberative processes are not doing the work of asking questions, reflecting on knowledge gaps, or weighing competing perspectives themselves, then these skills at the heart of democratic citizenship are likely to weaken.

A complementary use of AI treats it as something that helps citizens to engage more deeply with evidence and expertise. For example, AI can help make information packs more engaging through transforming them into multi-modal formats, such as images, podcasts, and/or videos. By enabling the evidence base to be communicated in a way that accommodates a diversity of learning styles, citizens could be better equipped to learn about and reason through the complexity of the issue under consideration. We note that this is a different task to developing the evidence framework. Deliberation process designers should keep human subject matter expertise in-the-loop during the evidence framework creation and first draft materials to ensure information is accurate and robust. We note that this is a different task to developing the evidence framework. Deliberation process designers should keep human subject matter expertise in-the-loop during the evidence framework creation and first draft materials to ensure information is accurate and robust.

Another example of using AI to complement reasoning is inspired by De et al. (2026). They note that a core part of deliberative processes is when citizens identify knowledge gaps and pose questions to experts. Especially in deliberative processes with a large number of participants, it can be very challenging to determine which questions should be prioritised for an expert to respond to. To address this, De et al. propose using AI to filter the questions that citizens propose into a digestible number that can then be reviewed and posed to experts. Citizens still get to exercise their reasoning skills - the work of generating and prioritising critical questions remains with them - but AI helps with distilling it into actionable outputs.



The dialogue muscle

The capacity to listen attentively, respond, take turns, and provide reasons for one's claims.

The challenge is that people have little practice with dialogue in this definition of the term, especially in the political realm. In many realms of life, we are exposed to performance and bite-size signalling that reduces complexity. Point scoring and winning debates are more frequent goals than trying to arrive at deep understanding across divides. Many arrive at deliberative forums shaped by these experiences and habits.

Deliberation requires, however, that we exercise our capacity to listen openly, to respond charitably to others, and to offer reasons that could be acceptable by someone who does not share our view. This demands attentiveness, generosity, and a willingness to slow down.


A competitive use of AI would be for it to take on the listening and responding work of the participant. The most extreme version of this we mentioned above when discussing "AI agent" deliberation in the section about self-reflection muscles. The participant would in that case gain the "efficiency" of never having had to listen or respond to anyone, nor of sitting with the discomfort of articulating a perspective with which others might disagree.



A milder but also competitive use of AI draws on similar logic to draft a participant's contribution for them – taking an initial rough input and producing a polished, reasoned argument to be shared as if it were one's own. In this case as well, the dialogic work is not actually being done by the human; their agency is relegated to the AI and their skills in dialogue (as well as self-reflection and reasoning) are not being strengthened.

A complementary use of AI would support the participant in becoming a better listener and help them to articulate their own views more clearly. This could take the form of participants receiving reason-giving prompts during an exchange. For instance, a participant struggling to articulate why they hold a view could be helped by a prompt that asks them what experience or value(s) they are drawing upon, without being told what to say. Arguably, this is the role that a skilled facilitator tends to play.

Another example is that real-time multilingual translation can help in relevant contexts. When a participant can hear another participant's reasons in their own language – including the texture and tone of voice from their original contribution – they can listen and respond from a place of deeper understanding. The same principle applies to other accessibility measures such as live captioning.



Assembly members working in small groups with Dembrane running on a facilitator's phone. Photo: Gustav Kjær Vad Nielsen.

The vulnerability muscle

The capacity to courageously share one’s feelings and reflections, to sit with discomfort, to tolerate genuine conflict, to experience love, and to feel the weight of others’ experiences. Vulnerability entails friction, fragility, and felt accountability.

The challenge is that deliberation is more than a rational exercise involving evidence and facts; it also entails engaging with our and others’ emotions, both positive and negative. To be capable of doing the work of deliberation – of grappling with complexity, being open to changing one’s mind, and willing to compromise on a shared decision – research shows that building trust, bonding with others, and even feeling love are crucial factors that enable deliberation (Landemore, 2026; Niemeyer et al., 2024).

On the other hand, encountering perspectives that challenge one’s identity, sitting with information that conflicts with one’s worldview, feeling the force of a stranger’s personal testimony – these can be uncomfortable. The affective work of deliberation is one of the key things that distinguishes it from preference aggregation. In deliberation, we exercise empathy, build trust, connect with our fellow humans, feel seen and heard. This echoes Martha Nussbaum’s wider work on the fragility of goodness, building upon Aristotle’s argument that the good life depends upon what is fragile: love, friendship, civic life, the body (Nussbaum, 2001). These conditions enable us to be open to new perspectives. As Landemore has suggested, “You will only get the epistemic benefits, all the collective intelligence, if you trigger the love part first.” (DemocracyNext, 2026).

Furthermore, a deliberative process should result in shared decisions backed by an accountability relationship. Gillian Stamp has articulated in a compelling way why this also necessitates vulnerability in her writing on [felt accountability](#) – the felt state of dependency on the people we rely on in a piece of work: our colleagues, our partners, “the people the work runs between”. As she puts it: “accountability is felt as dependency on the knowledge, expertise, competence, and judgement of those people” (Stamp, 2007). In this view of accountability, it is not an inner state, nor an external sanction, but it exists in the felt mutuality between embodied subjects.

A competitive use of AI would be to smooth the discomfort and remove the friction. For instance, in some online-only settings, there are moderation guidelines implemented whereby AI systems translate harsh language into softer formulations before they reach the recipient. There may be contexts where this is helpful and could prevent escalation in volatile online forums. But in a deliberative setting, where one goal is to develop the capacity to encounter difficulty directly, an “F-bomb translator” may train participants out of the tolerance for friction that the process is meant to cultivate.

There is also increasing interest in using AI to simulate absent stakeholder perspectives in deliberative assemblies (Fulay, Dimitrakopoulou, and Roy, 2025). Whilst using personas has been a long-time practice to help in crafting or testing recommendations later in a process, this is not a substitute for primary key lived experience input. Whilst using personas has been a long-time practice to help in crafting or testing recommendations later in a process, this is not a substitute for primary key lived experience input. Replacing it removes the friction of real-world human interaction and abstracts from the conditions that enable people to exercise vulnerability and deep connection across divides. During the Irish Citizens’ Assembly on abortion, for example, one of the most consequential moments was hearing testimony of those whose lives had been impacted by abortion laws. It is hard to imagine a similar transformation taking place if lived experience was replaced by a generative AI model.

Similarly, AI consensus generators like the [Habermas Machine](#), which produces synthesised statements designed to bridge divergent views without any human ever directly interacting with another human, remove the friction, fragility, and the resulting felt accountability that make deliberation valuable (Tessler et al., 2024). This is another example where one would be hard-pressed to call it deliberation at all.

Asynchronous text-based tools, such as Pol.is, might be a helpful complement to deliberative forums depending on their design and purpose. From the vulnerability angle, whilst these are not competitive artefacts, using these technologies in isolation is likely to flatten the experience: all communication that results from tone of voice, facial expressions, and body language is lost, meaning the emotions shaping people's understanding of the facts are also absent.

A complementary use of AI does the opposite: it intentionally helps to surface friction. For example, AI tools that identify perspectives the conversation may have passed over without justification, or that flag when an emerging consensus has not really engaged with a strong objection can deepen the affective work rather than bypass it. Tools that help people to articulate more clearly why they disagree and levels of disagreement, surfacing their underlying values, can sharpen disagreement into something meaningful and productive.

A frictionless ideal may seem desirable, but it actually makes learning, deciding, and any other meaningful activity difficult, if not impossible. We are reminded of Wittgenstein's famous quote about returning to the rough ground:

“We have got onto slippery ice where there is no friction and so in a certain sense the conditions are ideal, but also, just because of that, we are unable to walk. We want to walk: so we need friction. Back to the rough ground.”

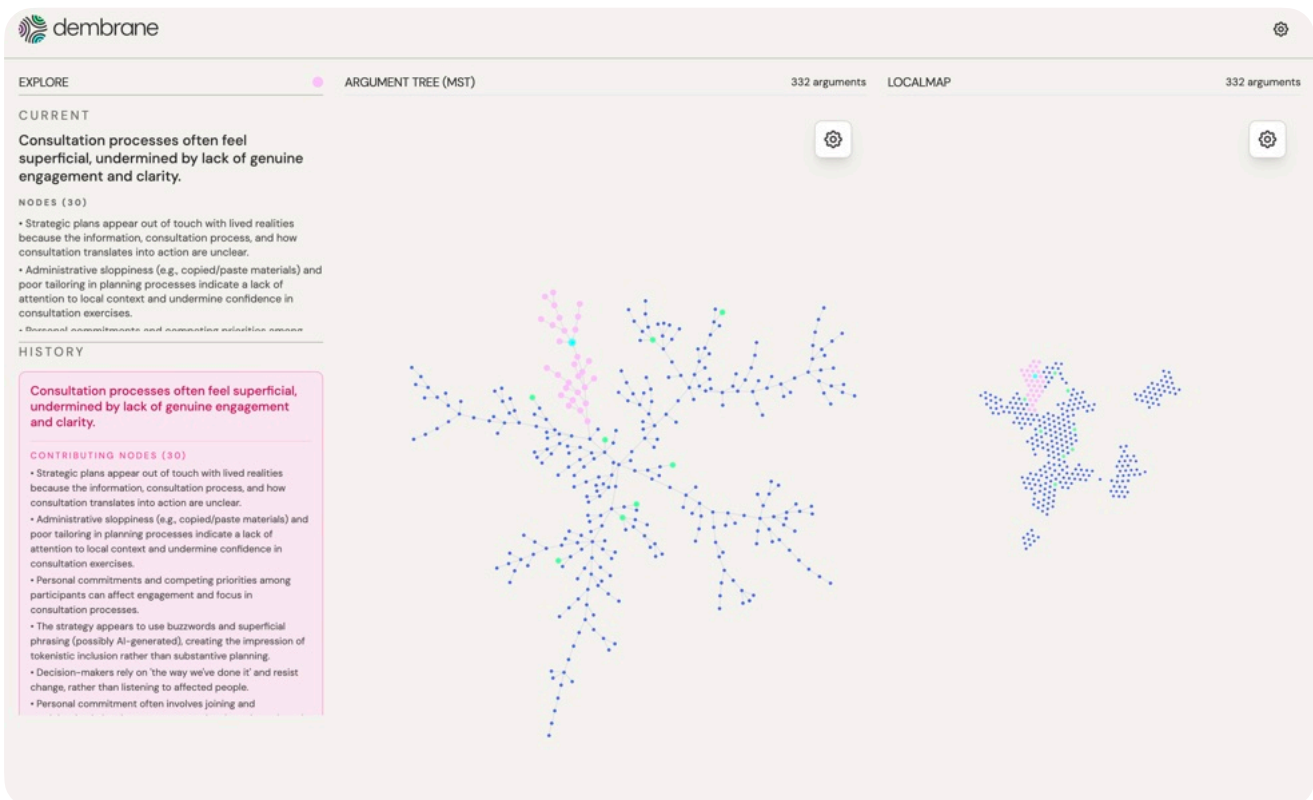
Friction is not an obstacle. It is the condition. When it comes to the elements of vulnerability captured in fragility and felt accountability, there may not be a complementary use of AI.

The collaboration muscle

The capacity to move from individual reasoning to shared judgement – to experience collective ownership of an outcome by doing the work of sensemaking together. It also involves working through potential conflict, debating differences, and figuring out compromise and common ground.

The challenge is that synthesising a group’s thinking is slow, demanding work. Clustering ideas, identifying themes, deciding which threads to prioritise, working out where the group actually stands – this typically takes many hours of facilitated effort and risks losing nuance or minority perspectives along the way. As the group size grows, this becomes even more challenging.

A competitive use of AI is to take the work away. It is tempting to hand the transcripts to a language model, get back a clean thematic summary, and present it to the group as “what you said.” But, as practitioners have repeatedly observed, the work of making sense of a conversation together is one of the most generative parts of the process. As one facilitator told us at a DelibTech Network meeting in May 2026, clustering ideas into themes “is the heart of the process – if we took out that heart and gave it to somebody else, it feels like the end product would be a lesser version.” Others note an “instant lacklustre” when AI sensemaking is introduced – energy dissipates from the room as people sense their work being taken away. As Zakia Elvang from We Do Democracy aptly put it recently: the danger is a “layer of plastic wrap” between the participants and the realness of the encounter. There is something that seems transparent but ends up creating a layer of distance between people and the outputs that are supposed to be theirs.



Screenshot of Minimal Spanning Trees, Credit: Dembrane.

A complementary use of AI does not outsource the work of sensemaking but helps scaffold it. AI can produce first-pass summaries that participants can then interrogate, revise, and own – particularly in larger or multilingual settings. For instance, in Esch-sur-Alzette, [a 40-person citizens’ assembly deliberated in seven languages](#) using the [Dembrane](#) tool to give each participant a facilitator-reviewed summary in their own language at the end of the day (MacDonald-Nelson and Terry, 2026). AI summarisation and translation here enables the group to have a common starting point to collaboratively reflect and revise.

Beyond this, different kinds of AI can also help visualise the relationships between ideas the group is generating. In a research preview, Dembrane is using argument embeddings and a century-old technique called “minimal spanning trees” to grow networks of arguments based on how semantically similar they are. The emergent result is that more unique arguments end up on the edges of the graph, while arguments that combine topics become central nodes. The trees are presented for people to explore, rather than in a report-like format. In each case, AI tools can create artefacts that support the participants to work on further, rather than being a conclusion they receive. Such visual heuristics hold great potential for holding the overall emerging picture in sight while retaining argumentative nuance. Such visual heuristics hold great potential for holding the overall emerging picture in sight while retaining argumentative nuance.

The imagination muscle

The capacity to envision futures, counterfactuals, and alternative arrangements concretely enough to deliberate about them; to give form to what does not yet exist so it can be weighed alongside what does. It includes the curiosity to pursue questions rather than foreclose them, and to remain genuinely interested in what one does not yet know, in why others see things differently, and in possibilities one has not yet considered.

The challenge is that imagination is unevenly distributed and easily foreclosed. Participants are often asked to deliberate about long time horizons, systemic change, or technologies that don't yet exist. But most adults are out of practice at sustained imaginative work, and the dominant cultural frames offer a narrow set of pre-imagined futures: techno-utopia; collapse, or more of the same. Without active scaffolding, deliberation about the future can collapse into deliberation about the present, just projected forward.

A competitive use of AI for imagination is to outsource the imagining itself – ask the models what the future could look like, accept what comes back, and deliberate about that. Participants skip the imaginative work. Those who arrive at the room having explored the topic with a chatbot may arrive with answers rather than questions. Using AI as an imagination oracle forecloses the exercising of our latent creativity and may actively narrow the space of futures available for deliberation while presenting that narrowing as breadth.

A complementary use of AI may treat it not as the imaginer but as a tool that supports human imagining. One mode might be analogous to using a big thesaurus. AI could help us explore a wide and queryable index of what has already been tried. It can surface real-world projects that have grappled with similar challenges, prior thinking from outside the participants' usual reference set, or precedents from other domains and historical periods. The imaginative work stays with the participants; the AI extends their reach into the existing space without performing the synthesis for them.

A further complementary mode might treat AI as a way to model and visualise complex systems, so that participants can reason about how a proposed change might actually unfold. The questions deliberation tackles are almost always questions about complex systems, our institutions, economies, and ecologies, each with their own internal dynamics, feedback loops, and friction points. Used well, AI can act less like an oracle and more like a guide map: letting a group simulate the diversity of ways in which a proposed change might ripple outward, helping surface their second-order effects, and locating the points where a change might be pivotal, or where it might precipitate a black swan scenario they had not considered. The imaginative and evaluative work stays in the room; what the tool adds is a navigable picture of terrain that is otherwise invisible, including broad concessions of the unknown unknowns. To do this well, we might need systems that search against an objective rather than predict the next token: if used at all the tool must demonstrate that it can extend the space of futures under consideration rather than flatten it.

A third mode is the Socratic one: AI that returns questions rather than answers, surfacing gaps in the evidence, perspectives not yet heard, and assumptions the participant has not tested. Tools like Revel and Pénigaud's "AI Reflector" (2025) and [Dembrane's "Explore" functionality](#) (where the language model can be prompted to ask biting, socratic questions, tailored to challenge participant thinking) gesture in this direction and could be sharpened specifically to keep the imaginative space open rather than to fill it up with noise.

A contextual caveat is that these complementary modes are not without their issues. AI as thesaurus or librarian does actively compete with our own abilities in source finding and attribution, and fact checking authoritative sounding sources (current AI models still make up citations) is easier said than done, especially when an unrelenting AI generates seemingly endless reams of secondary source material at the click of a button.

Another caveat is that the line between surfacing existing alternatives for the participants to consider, and handing the participants a pre-formed future is thinner than it looks. Even the complementary modes can tip into competitiveness if they are over-curated. A tool that returns three "scenario options" is closer to an oracle than a thesaurus – however distinct the options may seem, they will be a far cry from the truly transformative option that might be proposed in the absence of AI, or when those resources dedicated to the AI would go to bringing creative professionals or completely novel facilitation techniques into the assembly instead.

The facilitation muscle

The capacity to surface different perspectives, guide discussions productively and inclusively, encourage contributions from everyone to navigate conflict and ultimately, to improve the capacity of others to do the same.

The challenge is that facilitation is a high-capability skill and there are not enough people with these skills today to match the demand and need of spreading deliberation more widely. A competent facilitator does many things at once – managing time, drawing out quieter voices, surfacing power asymmetries, holding space for silence, watching for emotional dynamics, holding the thread of the conversation (Fide and We Do Democracy, 2026). Even well-trained facilitators miss some elements, and the lack of capacity is a real bottleneck to scaling deliberative practice.

A competitive use of AI here is to replace the facilitator. This is an increasingly pursued endeavour amongst technologists and academics. An “AI moderator” could decide who speaks when, steer the conversation, and manage the dynamics of the group. Stanford’s deliberation platform allocates the same amount of speaking time to all participants in a small group, for instance, and people have to stop speaking when their time is up. While this may carve up speaking time equally, it is also not how conversations tend to flow.

Some people may need more time than others because rhetorical skills are unequally distributed in society; some people can do more with less time; some people may need more time than others because rhetorical skills are unequally distributed in society; some people can do more with less time; sometimes one point needs greater elaboration or someone shares a personal story that speaks to it; sometimes two people ping-pong short questions and answers to clarify a point before a conversation moves on; sometimes an introvert is more likely to listen to a conversation intently, taking time to form their thoughts to come in towards the end of a session, rather than needing to make a short point consistently throughout.

As one facilitator in the DelibTech Network noted, applying this approach rigidly could make some people feel anxious and withdraw from a process, or feel silenced. It takes pre-engagement for some people to even have the confidence to step into the room, let alone feel ready to share a view out loud. How would an AI handle someone bursting into tears and being given the spaciousnesses they need to recuperate? This is a relatively frequent experience in deliberative settings. One might question whether with time, AI facilitators could be better at adapting to the nuance described above rather than following pre-set specifications.

However, even setting aside whether AI systems could do this well, the deeper concern is that democratic societies would benefit from more citizens being skilled in the craft of facilitation, not less. This is already a tricky skill to develop and being able to do it live, in the open, is beneficial in more challenging and low-trust contexts. Automating it transfers a valuable human skill to a machine. In increasing times of polarisation, the more people with the capacity to navigate complex and fraught topics in a civic-minded fashion, the better. Facilitation is one part of what helps to enable our collaboration muscles to be activated in full force as well, supported by the scaffolding that enables a group to produce a collective output.

Indeed, this is also the muscle where the competitive and complementary lens turns on the human doing the work, not only AI. A facilitator of either kind who carries all of the facilitative work can leave a group of citizens no more able to guide their own deliberation than when they arrived; the capacity stays with the facilitator rather than the citizens. **The point is therefore not only to defend professional facilitation against automation, but to ask of any facilitation, human or machine, whether it builds a group's capacity to facilitate itself.**




Facilitators worked closely with Dembrane co-founder Evelien Nieuwenburg throughout the Esch assembly to ensure the platform was capturing deliberations correctly and integrating into the process. Photo credit: Gustav Kjær Vad Nielsen

On this reading, the highest expression of facilitation is the kind that gradually withdraws as a group grows able to hold its own, and the role of a lead facilitator shifts from running conversations toward growing facilitators: helping co-facilitators and participants take on more of the work over time.

AI-supported facilitation is complementary when it serves this same end: conversation guides and private nudges that help a novice or citizen facilitator find their feet, transcription that frees whoever is facilitating to attend to the group. It becomes competitive when it installs itself as the facilitator a group comes to depend on. Tools like [Coalesce](#) and [Frankly's Agenda Builder](#) can help facilitators draft conversation guides drawing on past templates of good examples, and helping the facilitators to think more deeply about what they might like to surface from the group and how to best help them move from idea generation to proposal formation, for instance (Overney, 2025; Harvard Berkman Klein Center, 2025).

In the room, AI transcription can free a facilitator from note-taking so they can attend to the group. In an online setting, private prompts to facilitators – gentle nudges about who has not spoken yet, or which themes haven't yet been covered – may support skilful navigation of power asymmetries and enable productive deliberative exchange. Some online platforms, like Stanford's and Frankly, also include a feature of private prompts being sent directly to participants themselves if they haven't spoken for a while. These functions are precisely what could be used to support citizens themselves to self-facilitate small group deliberations, rather than relying on outsourcing facilitation to external consultants. This expansion of facilitation capacity beyond expert facilitators, and the reframing of expert human facilitators as "facilitation trainers" will be necessary for disseminating this muscle more widely.



CHAPTER 04

Investing in civic infrastructure and civic AI: A precondition for democratic resilience

The vision is that an increasing proliferation of deliberative assemblies, in which citizens have an opportunity to exercise and cultivate their deliberative muscles, creates a virtuous cycle of capacity building. More citizens participating in meaningful processes, more civil servants getting exposure to deliberative forums, more stakeholders contributing their stories and expertise. As these practices proliferate, we have the potential to bolster social cohesion, channel collective intelligence and - crucially - strengthen our deliberative capacities beyond the assembly.

This matters for how we think about the technological choices at stake. The case against over-automation for the sake of efficiency is not just about an individual leaving a deliberative forum not having exercised their deliberative muscles.

It is that the societal capacity that democratic societies need – the rich, civic competence to run deliberative processes, sustain civic associations, train new facilitators, integrate citizens’ recommendations into formal decision making structures, and defend processes against elite capture – is built through many individual exercises of self-reflection, reasoning, dialogue, vulnerability, collaboration, imagination, and facilitation.

There is no shortcut to building up that capacity at a societal level. It needs to be exercised, widely and repeatedly in many different institutions of government and organisations of daily life, over time.

There are three key considerations related to societal capacity.

First, for this virtuous cycle to materialise, **we need to invest in civic infrastructure that can support and sustain the structures that enable people to strengthen their deliberative muscles (e.g. training, standards, and ecosystem building bodies)**. Chwalisz and McKinney (2026) highlight numerous features of effective civic infrastructure for scaling democratic innovations. They argue that scaling quality deliberative practices can never just be a technological question, but one that requires attention to legal frameworks, education, community building, public communication, and more.

Second, there is a resilience argument that goes beyond civic capacity. **If deliberative practice becomes deeply dependent on AI infrastructure – on foundation models from a small number of providers, on data centres, on supply chains that span hostile jurisdictions – then shocks to that infrastructure become shocks to democratic practice.** Model access costs could spike. Geopolitical disruption could remove access to particular models. By buying the leading AI models from the large oligopolies, and outsourcing our technical skills to large AI companies, we are also foreclosing a future where we invest in these capacities ourselves. DIY AI, open models, or models being developed by governments or public institutions (such as Spain’s ALIA Public AI or the Swiss Public AI coop) might be expensive and the results might not be as good at first, but they enable a civic AI future. A deliberative practice that has built deep human capacity can continue when the technology falters; one that has substituted for that capacity cannot. The risk compounds when public sector jobs are being cut at the same time that AI dependency is growing. The risk compounds when public sector jobs are being cut at the same time that AI dependency is growing.


As Audrey Tang has argued, [there's a difference between treating data as oil \(drilled\) and as soil \(tilled\)](#) (Tang, 2026b). **If we build civic infrastructure for data, training, and evaluations that any community can use, the tacit knowledge of facilitators, care workers, and the people a community already trusts becomes the soil models grow in, rather than something extracted and lost. Civic AI isn't just a hedge against shock – it's a precondition for the virtuous cycle.** If data is soil then the communities, practitioners and local organisations are not just end-users of DelibTech, but central to the conditions that make it work at all. Human capacity isn't just a hedge against geopolitical risk; it's what makes the soil fertile in the first place.

At a moment when public sector capacity is being cut and community organisations are chronically underfunded, the relational and civic knowledge that deliberative technology depends upon is being eroded at precisely the same time as that technology is being introduced. This tension sits at the heart of whether the deliberative muscles framework can work in practice. If data is soil then the communities, practitioners and local organisations are not just end-users of DelibTech, but central to the conditions that make it work at all. Human capacity isn't just a hedge against geopolitical risk; it's what makes the soil fertile in the first place. At a moment when public sector capacity is being cut and community organisations are chronically underfunded, the relational and civic knowledge that deliberative technology depends upon is being eroded at precisely the same time as that technology is being introduced. This tension sits at the heart of whether the deliberative muscles framework can work in practice.

Third, this reframes what "good" deliberative technology looks like to practitioners and funders. **The lens of efficiency or scale alone is too narrow, and we must ask which tools leave the participants, the facilitators, and the wider institutional ecosystem more capable at the end of the process than at the start.** That is a harder thing to measure, but we argue that it is essential.



Examples of 'scaling catalyst' hubs that embody the civic infrastructure discussed above. Top row: We Do Democracy, Denmark. Source: Claudia Chwalisz; Demokrati Garage, Denmark. Source: Claudia Chwalisz. Bottom row: MMA's space, The Hague, The Netherlands. Source: <https://www.emma.nl/artikelen/over-generaties-heen-jongerenparticipatie-met-meer-impact>; Impact House Brussels. Source: <https://impacthouse.be/offer/event-space/>



CHAPTER 05

Three future trajectories to advance the deliberative muscles approach

Three future trajectories to advance the deliberative muscles approach

The aim of this paper has been to establish the seven deliberative muscles, motivate their importance to practice, and reflect on how this can guide design choices relating to AI's role in deliberative spaces. But this is not all that can or should be said on this topic. We highlight three particularly important areas to develop the deliberative muscles approach in future research and practice.

This list is not exhaustive and we welcome other ideas for research and outputs that would serve the field.

01

Exploring and measuring how different process designs affect the cultivation of deliberative muscles.

02

Communication guidance about using AI in deliberative processes.

03

Exploring the political economy dimensions of AI in deliberation and how this affects deliberative muscle cultivation.

Exploring & measuring how different process designs affect the cultivation of deliberative muscles

When designing a deliberative process, process designers and stakeholders make a variety of design choices depending on goals and context. This includes, for example, how many people to include in the process, how long the process will last, and what kind of expertise to include. There is no 'ideal' process design that can be exported across contexts.

Different designs will activate muscles in different ways: we can expect a small group of fifty people meeting over many weekends to more thoroughly exercise a breadth of deliberative muscles than one-thousand people meeting for a short one-off online deliberation, for example. These dynamics should be mapped out more holistically: understanding how process designs affect muscle activation and trade-off with other goals of deliberative processes, such as policy impact, inclusion, and efficiency, will support more informed decision-making.

One trajectory that will support this is to **develop metrics for measuring the extent to which muscles are strengthened or weakened in deliberative processes**. This would entail fleshing out a diverse methodological tool-kit of longitudinal study, including interviews, surveys, and participant observation, to understand how process designs shape the cultivation of each of the muscles during and beyond the deliberative process.

Communication guidance about using AI in deliberative processes

When AI is being used in a deliberative assembly, practitioners often struggle to convey how and why in a transparent and effective way. The deliberative muscles lens could help with this. Future research should develop a citizen-facing output for deliberation participants about the use of AI in deliberative processes.

This would show how the use cases aim to support, instead of replace, the cultivation of deliberative muscles. Such guidance should be grounded in evidence on the technical efficacy of specific design choices: which features complement participants' deliberative work and which compete with it (see [Focus on Features](#) by the Integrity Institute for one such feature-level approach).

The aim of this is not to coerce citizens into passively accepting AI use, but rather to enable informed judgements on its role in the process.

Exploring the political economy dimensions of AI in deliberation and how this affects deliberative muscle cultivation

The particular kinds of AI that get deployed in deliberative processes are not just shaped by attempts to advance democratic goals, but also by economic imperatives to win tenders and reduce costs, as well as a commissioning culture that emphasises short term benefits.

Future research should consider the extent to which, in practice, these market incentives and procurement practices cut against developing technologies that support deliberative muscle building in democratic innovations, and what political economic interventions may help mitigate these challenges.

If the deliberative muscles approach is to be channeled to its fullest potential, it must be put in closer dialogue with questions of AI ownership and infrastructure, expanding on the considerations identified in section 4.

This also points to questions of model provenance and ownership. A deliberative process is only as trustworthy as the chain of provenance behind the cognitive tools embedded in it, so clearer provenance for training data and model behaviour is not a separate infrastructural concern but part of what makes a tool complementary to democratic capacity in the first place.



CHAPTER 06

Conclusion

We live in a democratic climate increasingly defined by polarisation, manipulation, hostility, and decreasing institutional capacity to act in the face of epochal challenges. Yet we also have a tremendous opportunity to harness the power of deliberation and new technologies to shape a better future. We have argued that strong deliberative muscles are not a nice-to-have for a resilient democratic society; they are a necessary part of a holistic response to building and sustaining it.

Strong deliberative muscles are essential for navigating nuance and complexity, productively engaging with pluralism and divergent views, and cultivating empathy and civic love – qualities which are largely missing in electoral politics and many decision-making contexts today.

Deliberative assemblies are valuable not only for their outputs, but for what they enhance in the people who take part in them – individually and collectively. AI integration that prioritises efficiency over capacity building risks hollowing out a key ingredient that makes deliberation worthwhile.

An essential design question for technologists and practitioners is whether and how to use AI such that people leave with stronger deliberative muscles. Across seven dimensions - self-reflection, reasoning, dialogue, vulnerability, collaboration, imagination, and facilitation - we have identified pathways to support the navigation of this complex terrain.

Whilst we have focused predominantly on deliberative forums like citizens' assemblies, **the considerations in this paper are relevant to all places where deliberation takes place** – workplaces, schools and universities, museums, financial institutions, corporations and cooperatives, membership-based associations, and other organisations. The deliberative muscles framework and how technology affects their cultivation can act as a guiding resource for evaluating the integration of emerging technology into these spaces.

The stakes are high: will we develop and deploy technologies that strengthen our individual and collective capacity to enable democratic flourishing, or will we hollow out these capacities through technocratic shortcuts? Which pathway is taken will depend on various forces, including market imperatives towards competition and efficiency.

But that is not the whole picture, for **there is also the role of human agency**:

- **commissioners and practitioners** should resist easy efficiency gains that bypass hard work and human capacity building;
- **public actors** should invest in and develop digital public infrastructure and civic AI that bolsters trust and distributes power;
- **technologists** should reflect upon how the technologies they develop could strengthen and not weaken people's deliberative capacities;
- **scholars** should support the mapping of deliberative muscles and how technology affects their cultivation;
- **communities and civil society** should demand and design deliberative processes that sustain and strengthen people's deliberative muscles.

These structural pressures, however, require systemic change and support, which is why we also focused on the investment needed in civic infrastructure and civic AI if we are to enable conditions that make deliberative technology worth using at all.

Efforts across these multifaceted domains will support us in navigating one of the most pressing challenges facing humanity: building a world where technological innovation and democratic innovation work in harmony.

References

Aoyama, S. (2026). People Don't Know What They Want. <https://blu3mo.substack.com/p/people-dont-know-what-they-want?>

Buergerrat (2026). Citizens' Assembly Database. <https://www.buergerrat.de/en/citizens-assemblies/citizens-assembly-database/>.

Chwalisz, C. and McKinney, S. (2026). Scaling Democratic Innovation: Features of Effective Catalyst Organisations & Future Frontiers. DemocracyNext. <https://www.demnext.org/action/scaling-democratic-innovations>.

De, S., Gelauff, L., Goel, A., Milli, S., Procaccia, A.D. and Siu, A. (2026). Question the questions: Auditing representation in online deliberative processes. In Proceedings of the ACM Web Conference 2026 (pp. 1640-1650).

Democracy Fitness. The 10 democracy fitness muscles. <https://democracyfitness.org/the-10-muscles-of-democracy/>.

DemocracyNext (forthcoming). DemocracyNext Database of Deliberative Assemblies.

DemocracyNext (2026). Politics Without Politicians: A conversation with Hélène Landemore. Another Democratic Future. <https://demnext.substack.com/p/politics-without-politicians-a-conversation>.


FIDE Europe and We Do Democracy (2026). What a Lead Facilitator Needs to Know and Do. Available at: <https://drive.google.com/file/d/1V38zlodHQF70yNjOWq4ODZaPu45IIPHp/view>.

Fishkin, J., Bolotnyy, V., Lerner, J., Siu, A. and Bradburn, N. (2024). Can deliberation have lasting effects? American Political Science Review, 118(4), pp.2000-2020.

Fulay, S., Dimitrakopoulou, D. and Roy, D. (2025). The Empty Chair: Using LLMs to Raise Missing Perspectives in Policy Deliberations. arXiv preprint arXiv:2503.13812.

Gerlich, M. (2025). From Offloading to Engagement: An Experimental Study on Structured Prompting and Critical Reasoning with Generative AI. Data, 10(11), 172. <https://doi.org/10.3390/data10110172>.

Gudiño, J.F., Grandi U. and Hidalgo, C. (2024). Large language models (LLMs) as agents for augmented democracy. Phil. Trans. R. Soc. A 382: 20240100. <https://doi.org/10.1098/rsta.2024.0100>.



Harvard Berkman Klein Center (2025). Frankly Agenda Builder. <https://github.com/berkmancenter/frankly/issues/198>.

Hidalgo, C. (2025). Reimagining Democracy with AI Representatives. TED Talk: <https://www.youtube.com/watch?v=B2hRN8o6Xul>.

Integrity Institute (2026). Focus on Features. <https://features.integrityinstitute.org/>.

Ipsos (2025). The Ipsos AI Monitor 2025: A 30-Country Ipsos Global Advisor Survey. <https://www.ipsos.com/sites/default/files/ct/publication/documents/2025-06/ipsos-AI-Monitor-2025.pdf>.

Kosmyna, N., Hauptmann, E., Yuan, Y.T., Situ, J., Liao, X.H., Beresnitzky, A.V., Braunstein, I. and Maes, P. (2025). Your brain on ChatGPT: Accumulation of cognitive debt when using an AI assistant for essay writing task. arXiv preprint arXiv:2506.08872, 4.

Krakauer, D. (2016). Will A.I. Harm Us? Better to Ask How We'll Reckon With Our Hybrid Nature. Nautilus. <https://nautil.us/will-ai-harm-us-better-to-ask-how-well-reckon-with-our-hybrid-nature-236098?ref=journal.artificialityinstitute.org>.

Landemore, H. (2026). Politics Without Politicians: The Case for Citizen Rule. London: Penguin Random House.


Lee, H.P., A. Sarkar, L. Tankelevitch, I. Drosos, S. Rintel, R. Banks, N. Wilson (2025). The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers. CHI '25: Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems, Article No.: 1121, Pages 1 - 22. <https://doi.org/10.1145/3706598.3713778>.

Life With Machines. (2025). From Division to Dialogue—How AI Can Save Democracy | Deb Roy (Ep. 16). <https://www.youtube.com/watch?v=qKTS0mfXRck>.

MacDonald-Nelson, J. and Terry, H. (2026). Deliberation in Europe's Most Multilingual Country. DemocracyNext. <https://www.demnext.org/action/deliberation-in-europes-most-multilingual-country>.

McKinney, S. and Chwalisz, C. (2025). Five dimensions of scaling democratic deliberation: With and beyond AI. DemocracyNext. <https://www.demnext.org/action/five-dimensions-of-scaling-democratic-deliberation-with-and-beyond-ai>.

Niemeyer S., Veri F., Dryzek J.S., Bachtiger A. (2024). How Deliberation Happens: Enabling Deliberative Reason. American Political Science Review 118(1): 345-362. doi:10.1017/S0003055423000023.



Nussbaum, M. (2001). *The Fragility of Goodness*. Cambridge: Cambridge University Press.

OECD (2020). *Innovative Citizen Participation and New Democratic Institutions: Catching the Deliberative Wave.*, OECD Publishing, Paris, <https://doi.org/10.1787/339306da-en>.

Overney, C. (2025). *Coalesce: An Accessible Mixed-Initiative System for Designing Community-Centric Questionnaires with Generative AI*. MIT Center for Constructive Communication. <https://www.ccc.mit.edu/project/coalesce/>.

Van der Does, R. and Jacquet V. (2021). *Small-Scale Deliberation and Mass Democracy: A Systematic Review of the Spillover Effects of Deliberative Minipublics*. *Political Studies* 71, no. 1: 218-237.

Revel, M. and Pénigaud, T. (2025). *AI-Facilitated Collective Judgements*. Working paper. <https://arxiv.org/pdf/2503.05830v1>.

Reuchamps, M., Vrydagh, J. and Welp, Y. eds. (2023). *De Gruyter Handbook of Citizens' Assemblies*. de Gruyter.

Stamp, G. (2007). *Trust and Judgement in Decision-Making*. Published in: *Transformation*, 5th Edition: *Decision Making* (pages 35-39). Available at: <https://www.bioss.com/gillian-stamp/trust-and-judgement-in-decision-making/>.

Tang, A. (2026a). *Outrage to Overlap*. Audrey Tang. <https://au.civic.ai/p/outrage-to-overlap>.

Tang, A. (2026b). *Till Data Soil, Don't Drill Data Oil*. Audrey Tang. <https://substack.com/home/post/p-199529451>.

Tessler, M.H., Bakker, M.A., Jarrett, D., Sheahan, H., Chadwick, M.J., Koster, R., Evans, G., Campbell-Gillingham, L., Collins, T., Parkes, D.C., and Summerfield, C. 2024. *AI can help humans find common ground in democratic deliberation*. *Science* 386 eadq2852. DOI:10.1126/science.adq2852.

