Deliberative Muscles & AI

Claudia Chwalisz, Sammy McKinney, Jorim Theuns, and Eugene Yi | 9 July 2026

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 or strengthens 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, 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 civic 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.