




TESTING DELIBERATIVE TECHNOLOGIES TO IDENTIFY OPTIMAL USE

AN INFORMAL EXPERIMENT COMPARING DIGITAL DIALOGUE TOOLS

Davis Smith

About the Author



DAVIS SMITH

Davis Smith is a Senior Digital Collaboration Associate for Common Good AI. He graduated from the University of North Carolina (UNC) at Chapel Hill in 2023 with a degree in Computer Science. Before transferring to UNC, Davis attended Samford University in Birmingham, AL where he completed the University Fellows honours program, a “Great Books” curriculum focused on primary source readings and Socratic seminar discussion. Davis is also an avid sailor; he was a sailing instructor for six summers in North Carolina and most recently taught offshore sailing in the Lesser Antilles.

Cover image: TeraVector / shutterstock.com

The views expressed in this report are those of the author(s) alone. They do not necessarily reflect the views of the Toda Peace Institute. An online edition of this and related reports and policy briefs can be downloaded on our website: toda.org/policy-briefs-and-resources.html



Toda Peace Institute
Samon Eleven Bldg. 5F,
3-1 Samon-cho, Shinjuku-ku
Tokyo 160-0017, Japan

Tel. +81-3-3356-5481
Fax. +81-3-3356-5482
Email: contact@toda.org

© 2025 Toda Peace Institute

Abstract

Deliberative technologies are software tools that help create large-scale dialogue among participants. This article outlines the experience of testing four of these tools—Crowdsmart, Pol.is, Talk to the City, and Deliberation.io—with a group of student volunteers to understand their function and effectiveness, and to identify digital facilitation strategies. The paper concludes with recommendations to make deliberative tools more accessible in the future to enable collective decision-making. Common Good AI, a US-based nonprofit organization, created this programme to support its mission to foster inclusive civic engagement and social cohesion. The organization aims to transform how communities find common ground and solve problems together.

Introduction

Deliberative technology refers to a wide category of software tools aimed at harnessing collective intelligence in decision-making.^[1] They do this by allowing participants to engage in dialogue, written or verbal, often mediated through algorithms designed to elevate shared values to identify common ground. These tools come in many forms; some use AI^[2] to analyse user inputs, while others use bridging algorithms.^[3] Most deliberative tools allow participants to directly interact with each other's ideas, while others try to synthesise common positions based on participant input. In the past decade, there was a proliferation of deliberative tech applications globally; examples range from the Sunflower Movement in Taiwan^[4] to draft national policies, to California^[5] to inform the 2025 Los Angeles wildfires response, from community-driven participatory budgeting in Iceland^[6] to creating consensus statements between Israeli and Palestinians.^[7] There is a growing body of evidence that deliberative technologies increase participation in civic engagement and reduce toxic polarization.^[8]

To better understand the types of technologies available, Common Good AI (CGAI) launched a student ambassador programme to test and validate multiple deliberative tools. The programme recruited college students across the US to participate in multiple conversations about the current state of higher education, using deliberative tools. While not a rigorous, scientific study, the testing with the students yielded important learning on tool effectiveness and facilitation strategies. The programme allowed students to engage in eight deliberative sessions between May and July 2025.

^[1] Lisa Schirch, "Deliberative Technology: Designing AI and Computational Democracy for Peacebuilding in Highly-Polarized Contexts," Toda Peace Institute, October 14, 2024. <https://toda.org/policy-briefs-and-resources/policy-briefs/report-201-full-text.html>.

^[2] Goldberg, Beth, et al. "AI and the Future of Digital Public Squares." Last modified 2024. arXiv:2412.09988.

^[3] Felix Sieker, "Web Publication - Bridge-Building Instead of Polarization: How Algorithms Can Improve Digital Discourse," Startseite, January 22, 2024. <https://www.bertelsmann-stiftung.de/en/our-projects/reframetech-algorithmen-fuers-gemeinwohl/project-news/bridge-building-instead-of-polarization-how-algorithms-can-improve-digital-discourse>.

^[4] "Rethinking Democracy," vTaiwan, accessed October 3, 2025, <https://info.vtaiwan.tw/>.

^[5] Office of Data and Innovation - State of California, "Los Angeles Fires Recovery," Engaged California, accessed October 3, 2025, <https://engaged.ca.gov/lafires-recovery/>.

^[6] 1. "Better Reykjavik," Observatory of Public Sector Innovation, February 3, 2020, <https://oecd-opsi.org/innovations/better-reykjavik/>.

^[7] Konya, Andrew, Luke Thorburn, Wasim Almasri, Oded Adomi Leshem, Ariel D. Procaccia, Lisa Schirch, Michiel A. Bakker. "Using Collective Dialogues and AI to Find Common Ground Between Israeli and Palestinian Peacebuilders." Last modified 2025. arXiv:2503.01769.

^[8] Lisa Schirch, "Deliberative Technology: Designing AI and Computational Democracy for Peacebuilding in Highly-Polarized Contexts"

This article outlines each of the deployed technologies, with observations and suggestions on how they can best contribute to deliberative processes. It offers lessons on how to create digital engagements that inspire participation and spark engagements that encourage participants to find common ground.

Methodology

Common Good AI recruited a cohort of 25 students via LinkedIn from a diverse set of academic institutions, majors, and backgrounds. Over the course of eight weeks, students utilized a deliberative tool and participated in a virtual weekly discussion. The goal of these discussions was twofold: 1) review the results of the previous week's collaboration, and 2) hear what participants thought of the deliberative tool used that week. The decision to recruit students stemmed from the expectation that their interest in innovative digital tools and their common background as college attendees would foster meaningful participation. Topics were selected on a weekly basis, guided by prior conversations and a sense of which topics would most engage the group. The students were not paid for their time but served as volunteers. There was a final evaluation via a virtual focus group in the final week, as well as anonymous online surveys to get student feedback to the overall programme.

“CrowdSmart is most effective later in a deliberative process, when a group is asked to make a specific decision about what solution should be implemented

CrowdSmart

The programme started by testing CrowdSmart, the primary tool in Common Good AI's deliberative tech suite. CrowdSmart is a deliberative tool that relies on a ranking mechanism. Users are encouraged to answer a question and then prioritise the responses of other users. The system determined, over several iterations, the group's overall priority. Participants viewed and responded to each other's answers in real time, which allowed for direct interaction with other ideas. The built-in analytics dashboard creates a sense of transparency, as everyone can see how the group's thinking evolves throughout the discussion. Unlike other tools, CrowdSmart allows participants to actively monitor the state of the collaboration and reflect on its ability to build consensus.

Early in the programme, it became clear that CrowdSmart requires more effort than most other tools. Not only does it ask the user to rank seven ideas from best to worst for each question, it also needs participants to engage two or three times throughout the collaboration period. This correlated with relatively low engagement in terms of users returning to the platform multiple times. The students said using CrowdSmart felt like more effort compared to other tools.

Based on students' feedback and the team's own analysis, CrowdSmart is most effective later in a deliberative process, when a group is asked to make a specific decision about what solution should be implemented or what solutions should be prioritised.

Pol.is

“Pol.is was most powerful and useful when trying to do “broad listening” to understand the wider opinion landscape of the group

Pol.is allows users to upvote and downvote ideas submitted by their peers, as well as submit their own ideas. The Pol.is’ algorithm groups users into subgroups and highlights not just the most popular ideas, but ideas popular across different opinion groups. Pol.is produces easy-to-read summaries on areas of consensus and division, as well as a visual map of the different opinion groups. That being said, participants found the final “report” Pol.is generated challenging to understand.

In the experiment, the statements that represented the group as a whole and even the subtraction were a consensus view: safe, conventional, and unlikely to stir either enthusiasm or dissent. Without the proper facilitation and framing, this leads to a lot of broadly agreeable, inconsequential ideas. It’s easy to see where participants agree and disagree, but hard to understand the views of the opinion groups and where they differ.

In testing, Pol.is was most powerful and useful when trying to do “broad listening” to understand the wider opinion landscape of the group. Pol.is was less effective at generating and highlighting consequential, actionable ideas.

Talk to the City

“Students responded positively to chatting with the bot, but said they wanted it to challenge their ideas rather than simply affirm them

Talk to the City (T3C) is an LLM-powered tool built by the AI Objectives Institute that helps to categorise and group participants’ statements. Unlike the other tools tested, participants do not interact directly with T3C. Instead, T3C is used in the backend to group and understand statements from responses provided through a chatbot, or computer-generated conversation, on WhatsApp. This tool is increasingly popular in areas with limited data accessibility because users can chat with an AI-powered bot using communications apps they’re already familiar with, like WhatsApp or Telegram.

For the test, the bot was interactive and followed up on user comments to make the engagement deliberative, as opposed to just asking a series of pre-defined questions, like in a survey. Again, the discussion focused on the use of AI in college, specifically asking participants how and why they used AI for their assignments. The exact behaviour of the bot was controlled through a system prompt. For this experiment, we designed the bot to follow up on what participants had to say and encouraged them to expand.

T3C’s was a low barrier to entry because students engaged with it on WhatsApp, a familiar texting platform. It was easy to use and required only a brief time commitment, which made it easy to gather a broad sense of participants’ perspectives. Students responded positively to chatting with the bot, but said they wanted it to challenge their ideas rather than simply affirm them. They found the bot too agreeable. Also, the bot continuously asked questions which created confusion because there was no definitive endpoint to the conversation.

Participants noted that they preferred tools which allowed them to directly interact with others’ ideas. In this tool, students were not interacting with others’ comments; rather, they were interacting with the bot. With no analytics page or insight into other participants’ ideas, the students felt like they were putting their ideas into a black box. After using the other tools, like CrowdSmart and Pol.is, it was clear that the students wanted to interact with others’ ideas.

T3C offers the opportunity to analyse the sentiment of numerous users. It is most effective for a deep listening exercise where the goal is to broadly understand participant sentiment. Above all, in the backend, T3C excelled at grouping and clustering responses, proving more effective than other tools in analysing and representing the collective opinion of the group. T3C was effective at providing deeper analysis than Pol.is and even CrowdSmart data, effectively grouping ideas and creating synthesis statements.

Deliberation.io

“Students liked that their opinion was both affirmed and challenged. They liked that they could reconsider their initial idea

Deliberation.io is a platform co-developed by the Massachusetts Institute of Technology's (MIT) GovLab and Stanford University. It has a modular design, meaning it is a complex system created through smaller, distinct parts. Depending on the engagement, it allows facilitators to drop in different modes of participation. One of the primary goals of the tool is research: figuring out what modalities of deliberation are most effective in which circumstances.

For this use of Deliberation.io, participants were asked to rank an opinion on a 1–5 scale and provide the reason for their answer. Next, students chatted with a Socratic AI bot that pushed back on their responses and encouraged them to refine their answers. Then, participants were given an opportunity to see and vote on responses from other students. Finally, it asked them again about their opinion on the issue and whether their mind had changed after interacting with other ideas.

There was an overwhelmingly positive response to the use of the Socratic AI bot because students liked that their opinion was both affirmed and challenged. They liked that they could reconsider their initial idea. Also, Deliberation.io requires users to engage with the platform only once, more similar to a traditional poll, unlike Pol.is and CrowdSmart that work best when participants engage multiple times. It can be challenging to get people to return to a platform multiple times. At the same time, having only one interaction with the tool limits its deliberative potential compared to Pol.is or CrowdSmart, where participants have the opportunity to develop their ideas over time through conversation. Deliberation.io collects and synthesises participants' ideas, but is not set up for deeper continuous conversation.

Combining: Pol.is + CrowdSmart

While these tools identified consensus areas, it was important to understand how to identify and explore more contentious issues. First, Pol.is was used to surface reactions to the topic: if students should use AI in college, what are the productivity gains versus risks of misuse? Pol.is highlighted the group's disagreements. To probe deeper, students were asked to use CrowdSmart to weigh in on a statement and rank each other's reasoning. This approach yielded a more nuanced picture of the group's sentiment, with the prevailing view that while AI can efficiently handle routine tasks, its use may also erode students' critical thinking by offloading more complex work. The experiment affirmed the value of combining Pol.is and CrowdSmart: the former maps consensus and division, while the latter deepens understanding and reveals group priorities. The initial findings suggest that pairing these tools can enrich deliberation and foster more effective group decision-making.

Observations on digital engagement

Common Good AI's Student Ambassador Programme was a test bed to learn about deliberative tech tools, not a rigorous scientific evaluation. The takeaways from this programme, while valuable, are subject to some potential biases, but ultimately offered early lessons on how to best engage diverse groups to discuss topics using asynchronous tools.

Deliberative tech is most effective when used with real people in real situations to make real decisions. While each week was designed to offer an engaging topic for discussion, the tools were not informing a 'real' deliberative process. This led to lower engagement in both number of participants and depth of the conversation. It took significant effort to motivate students to engage in the online dialogues to generate substantive discussions. Also, many of these tools were designed to engage large groups of upwards of 50 or more people, and this group was under that threshold.

“The broader questions did not yield high participation in the digital tools; rather, it was the more divisive topics that stimulated engagement.”

In addition, there are two other factors that informed engagement: questions must be compelling and relevant to stimulate response, and participation is stimulated by idea exchange.

Many of the early questions posed to the students in the deliberative tools were open-ended and broad, like “What is the value of college?” and “How should we use AI in schools?” These general topics were less engaging than exploring more controversial issues like “Agree or disagree: Generative AI tools should be banned in K–12 education.” When participants were asked about a topic that elicited divergent opinions, their engagement and depth of response increased significantly. The broader questions did not yield high participation in the digital tools; rather, it was the more divisive topics that stimulated engagement.

The cohort was intentionally diverse; their only broad commonalities were (1) being students and (2) having an interest in technology. Initially, students were asked to identify the topics they wanted to discuss. Engagement peaked when students were invited to speak about their own challenges rather than broad, abstract issues. In the final feedback session, they said the most compelling topics were those tied directly to their lives, such as post-graduation unemployment. It suggests that digital engagements need to draw on personal experiences to compel people to share their ideas.

The ‘velocity of ideation’ was another important factor that informed participation. This term refers to how quickly people work off other ideas to stimulate engagement. In live conversations, like virtual calls, participants build on one another’s ideas in real-time, allowing their thinking to evolve quickly. As highlighted above, the students preferred to engage and react to others’ ideas. While some digital tools included this feature, others did not. It meant the weekly live discussions took on new significance. They shared ideas built on each other’s thoughts; compared to the deliberative tools, some lacked this velocity dynamic, making it less engaging and resulting in lower participation.

Next steps

This series of experiments with deliberative technologies sparked a variety of issues that need further exploration.

DEVELOP SHARED METRICS AND A ‘DELIBERATION INDEX’

Since different tools have different modalities of engagement and elicit different behaviours and reactions, it is difficult to accurately compare these tools. Understanding what they do requires shared metrics that qualify ‘depth of deliberation’, beyond only participation counts. Metrics could also capture exposure to opposing views, reasoning quality, view change, idea iteration, and return rates. Across the deliberative tech eco-system, it would be a valuable investment to create a lightweight ‘deliberation index’ to better understand what these tools can achieve and how to optimise them. As tools are designed, they can capture these metrics to measure effectiveness.

EMPLOY A ‘DELIBERATIVE STACK’ METHODOLOGY

Given that each tool offered a different function, this programme highlighted the need and value of a ‘deliberative stack’, meaning combining and deploying different tools to serve specific purposes within a deliberation. Deliberation.io and Talk to the City were useful for broad listening, helping to chart the contours of group opinion using an easy, accessible chatbot function. Pol.is offered opinion mapping to foster ideation. CrowdSmart excelled in the final stages of prioritisation and decision-making, rounding out the deliberative process. Each tool provided a distinct value and when combined, with the right questions and engagement strategies, they can be a powerful mechanism for deliberation.

DEVELOP UNIVERSAL METRICS AND INTEROPERABLE TECHNOLOGIES

Together, the development of universal metrics and interoperable technologies can make deliberative tools more accessible for future use. It would help demystify their functions and enable users to make informed design decisions to maximise their purpose. Informed by the growing research on deliberative technologies, there is tremendous potential for these tools to engage more individuals in policy dialogue. Greater collaboration between technologists, academics, and practitioners can help align with the shared goal of using new technologies to build stronger societies and governance structures that value collective decision-making.



THE TODA PEACE INSTITUTE

The Toda Peace Institute is an independent, nonpartisan institute committed to advancing a more just and peaceful world through policy-oriented peace research and practice. The Institute commissions evidence-based research, convenes multi-track and multi-disciplinary problem-solving workshops and seminars, and promotes dialogue across ethnic, cultural, religious and political divides. It catalyses practical, policy-oriented conversations between theoretical experts, practitioners, policymakers and civil society leaders in order to discern innovative and creative solutions to the major problems confronting the world in the twenty-first century (see www.toda.org for more information).

CONTACT US

Toda Peace Institute

Samon Eleven Bldg. 5 th Floor
3-1 Samon-cho, Shinjuku-ku, Tokyo 160-0017, Japan

Email

contact@toda.org

Sign up for the Toda Peace Institute mailing list

<https://toda.org/policy-briefs-and-resources/email-newsletter.html>

Connect with us on the following media.

YouTube: [@todapeaceinstitute3917](https://www.youtube.com/@todapeaceinstitute3917)

X (Twitter): <https://twitter.com/TodaInstitute>

Facebook: <https://www.facebook.com/TodaInstitute>