



HAL
open science

AI Agents and the Future of Deliberation: Designing Human-AI Collaboration for Democratic Dialogue

Weiyu Zhang, Shunyi Yeo, Simon Perrault, Jiaxin Pei, Anna De Liddo, Francesco Veri, Maurice Flechtner, Horacio Saggion

► **To cite this version:**

Weiyu Zhang, Shunyi Yeo, Simon Perrault, Jiaxin Pei, Anna De Liddo, et al.. AI Agents and the Future of Deliberation: Designing Human-AI Collaboration for Democratic Dialogue. CHI EA '26: Extended Abstracts of the CHI Conference on Human Factors in Computing Systems, Apr 2026, Barcelone, Spain. pp.Article No.: 986, Pages 1 - 6, <10.1145/3772363.3790075>. <hal-05469979>

HAL Id: hal-05469979

<https://hal.science/hal-05469979v1>

Submitted on 13 Feb 2026

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

AI Agents and the Future of Deliberation: Designing Human–AI Collaboration for Democratic Dialogue

Weiyu Zhang
National University of Singapore
Singapore
weiyu.zhang@nus.edu.sg

ShunYi Yeo
National University of Singapore
Singapore
yeoshunyi.sutd@gmail.com

Simon Tangi Perrault
Télécom Paris, Institut Polytechnique
France
simon.perrault@telecom-paris.fr

Jiaxin Pei
Stanford University
Palo Alto, United States
pedropei@stanford.edu

Anna De Liddo
Knowledge Media Institute
The Open University
Milton Keynes, United Kingdom
anna.deliddo@open.ac.uk

Francesco Veri
Centre for Democracy Studies Aarau
University of Zurich
Aarau, Switzerland
francesco.veri@zda.uzh.ch

Maurice Flechtner
Centre for Democracy Studies Aarau
University of Zurich
Aarau, Switzerland
maurice.flechtner@zda.uzh.ch

Horacio Saggion
Universitat Pompeu Fabra
Barcelona, Spain
horacio.saggion@upf.edu

ABSTRACT

As societies grapple with increasing polarization and information complexity, the need for constructive, inclusive, and well-informed deliberation has reached an unparalleled level. At the same time, AI agents, ranging from LLMs and multi-agent simulations and systems to conversational assistants and reflective companions, are rapidly reshaping how people communicate, reason together, and form collective judgments. These technologies hold the potential to scale democratic participation, foster inclusivity by bridging linguistic and cultural barriers, and introduce new forms of collaborative reasoning. Yet they also pose epistemic challenges to established notions of authenticity, legitimacy, and human autonomy in civic dialogue. This panel brings together leading researchers from Asia, Europe and North America to examine how AI technologies are transforming deliberation as both a social process and a design problem. It interrogates AI’s role in shaping deliberative norms, influencing group dynamics, and redefining what it means to “reason together” in hybrid human–AI spaces. Through interactive polling, structured debates, and audience co-deliberation, the session invites CHI participants to collectively explore how we can design responsible, inclusive, and trustworthy deliberation interfaces that preserve the democratic values of deliberation while embracing the creative potential of AI.

CCS CONCEPTS

• **Human-centered computing**; • **HCI Design and Evaluation Methods**;

KEYWORDS

deliberation, civic engagement, civic dialogue, democratic participation, digital democracy, AI agent, human-AI, large language models

ACM Reference Format:

Weiyu Zhang, ShunYi Yeo, Simon Tangi Perrault, Jiaxin Pei, Anna De Liddo, Francesco Veri, Maurice Flechtner, and Horacio Saggion. 2018. AI Agents and the Future of Deliberation: Designing Human–AI Collaboration for Democratic Dialogue. In *Woodstock ’18: ACM Symposium on Neural Gaze Detection, June 03–05, 2018, Woodstock, NY*. ACM, New York, NY, USA, 6 pages. <https://doi.org/XXXXXXXX.XXXXXXX>

1 INTRODUCTION

Deliberation — reasoned, inclusive discussion among diverse individuals, has long been regarded as a foundation for democratic decision-making and civic participation [15, 19]. Rooted in Habermasian ideals of communicative rationality [19] and further developed through deliberative democracy theory [3, 9, 13, 15], deliberation emphasizes inclusive, reasoned dialogue among equals as a means of reaching mutual understanding and consensus [5, 12, 17, 21].

Within Human–Computer Interaction (HCI), these democratic ideals have inspired research on deliberation technologies: systems designed to scaffold civic reasoning, reflection and equality. Early studies explored online discourse design and deliberative platform architectures [29, 40, 45], argument visualization [14, 22, 23, 30], and civic engagement tools that support structured dialogue [11, 25, 26, 40]. Collectively, this body of work has shown how interface design can shape participation, fairness [2], and reflection in digital spaces [31, 41, 42]. Through careful design, HCI researchers have

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Conference acronym ’XX, June 03–05, 2018, Woodstock, NY

© 2018 Association for Computing Machinery.
ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00
<https://doi.org/XXXXXXXX.XXXXXXX>

sought to foster empathy, bridge divides [38], and enhance the quality of public reasoning.

However, the ongoing wave of AI-driven technologies, particularly large language models (LLMs) and multi-agent systems, is introducing new dynamics that are transforming how deliberation unfolds in practice. AI can now moderate discussions by filtering or flagging content [24, 27], mediate conversations through summarization or clustering [18, 20], generate arguments and narratives [28, 34, 43], and even simulate deliberation through synthetic agents modeling diverse social perspectives [35]. These developments promise to scale civic dialogue to unprecedented levels while offering tools for multilingual translation [1], content simplification [33], and deliberative reflection [38]. However, as these innovations challenge traditional boundaries between human participants and computational systems, they also introduce profound epistemic and normative challenges: If AI agents can emulate human reasoning, to what extent should they act as co-deliberators rather than mere tools [10]? How might reliance on AI and algorithmic facilitation affect trust, legitimacy, and autonomy in civic discourse [4]? How can designers maintain transparency and accountability when AI systems shape or even generate public arguments?

As the HCI community, situated at the nexus of technology, design, and society, is uniquely positioned to examine these tensions. In designing AI deliberative systems that influence civic reasoning, these questions have become both urgent and unavoidable.

The last CHI-wide engagement on deliberation took place during the 2015 SIG on Online Deliberation, predating the AI revolution that has since transformed deliberation and civic discourse. A decade later, the convergence of deliberation and AI demands renewed collective inquiry. This panel seeks to catalyze a global conversation among HCI, communication, and AI researchers by examining how intelligent deliberative systems mediate public reasoning, reconfigure human participation, and reshape the democratic values embedded in design. By bringing together empirical research, design experimentation, and ethical reflection from around the world, the panel offers a timely platform to define the contours of a new research frontier: AI and online deliberation.

2 GROUNDWORK FOR PANEL DISCUSSION

2.1 Early Foundations in Digital Deliberation

Foundational work on argumentation systems, including Compendium [11] and Deliberatorium [25], focused on structuring reasoning and visualizing arguments to support large-scale public discourse. These systems illustrated how computational scaffolds could make complex debates more navigable, highlight relationships among claims, and facilitate collaborative sense-making. In parallel, research on online forums and civic platforms examined the social and procedural dimensions of deliberation, including norms of interaction, moderation strategies, and design interventions that enhance deliberative quality [8, 37, 39, 42]. These studies collectively revealed that seemingly small design choices, ranging from interface layout (e.g., [29]) and visualization techniques (e.g., [14]) to reflection prompts (e.g., [42]) and feedback mechanisms, can significantly shape deliberative outcomes, influencing aspects such as equality of participation, quality of argumentation, and mutual respect among participants.

Together, this body of work established that deliberation is not purely a social process but is tightly coupled with the technological environments in which it unfolds. By demonstrating the profound impact of design on reasoning, participation, and the cultivation of shared understanding, these early efforts laid a critical foundation for subsequent innovations at the intersection of HCI and online deliberation.

2.2 The Rise of AI in Civic Contexts

The integration of AI into online deliberation technologies represents a new phase in the trajectory of digital deliberation research. Earlier systems primarily sought to structure or visualize human discourse (e.g., [14]); contemporary systems can now actively intervene in the deliberative process. AI-driven deliberative tools can moderate discussions [24], identify incivility or misinformation to preserve discourse integrity [7], and assist in synthesizing complex, large-scale conversations [20]. They are also employed to extract key arguments, cluster viewpoints to enhance accessibility and engagement [36], and even generate deliberative contributions that mimic or complement human reasoning [33, 35]. This shift marks a profound transformation, from tools that merely scaffold human reasoning to agents that participate in shaping the environment and direction of civic dialogue. The rise of LLMs and algorithmic facilitation opens both unprecedented opportunities and complex challenges for the quality, inclusiveness, and legitimacy of online civic discourse.

However, as recent reviews suggest [36], much of this literature remains oriented toward developing technical fixes for scale—optimizing algorithms to handle more users or streamline debate, while often overlooking the more nuanced social, cognitive, and ethical dimensions of high-quality deliberation. Emerging work has begun to address this gap by incorporating reflective prompts and self-reflection mechanisms designed to foster deeper reasoning and mutual understanding among participants [41, 42, 46]. These efforts signal a growing recognition within HCI that deliberation cannot be “scaled” through efficiency alone; it must be cultivated through design that values empathy, justification, and perspective-taking.

Contemporary research has also begun to distinguish between three forms of AI involvement in deliberation: (1) AI-facilitated deliberation, (2) AI-augmented deliberation and (3) AI-substituted deliberation. In AI-facilitated deliberation, systems such as Deliberation.io¹ aim to identify common ground and facilitate deeper conversations by employing automated argument analysis to support large-scale participation and enhance coherence across discussions. In AI-augmented deliberation, initiatives like AI4Deliberation Europe² employ AI tools for summarization, fact-checking, moderation and hate speech detection to ensure safe, informative, and manipulation-free discussions. Finally, AI-substituted deliberation, as demonstrated in multi-agent simulations [35] uses AI stand-ins to emulate deliberative dynamics in silico, offering a testbed for understanding collective reasoning but also raising deep questions about authenticity, agency, and democratic legitimacy.

Taken together, these developments mark a paradigm shift in the design and study of civic dialogue. AI no longer serves merely

¹<https://deliberation.io/>

²<https://www.ai4project.eu/>

as an assistive layer that structures human input; it now acts as an epistemic agent that can shape the contours of deliberation itself. This evolution calls for renewed attention from the HCI community to interrogate how AI-mediated systems can sustain, not supplant, the human values at the core of democratic deliberation.

2.3 Global Efforts and Research Gaps

Researchers are experimenting with how AI can expand and enrich democratic participation. Across Europe, projects such as ORBIS³, MultiPoD⁴, and iDEM⁵ investigate how AI can scale participatory democracy, facilitate multilingual translation, and simplify complex policy language to promote inclusivity. In the United States and Asia, parallel initiatives explore AI personas [6, 16, 44], digital twins [32, 35], and reflective agents [41, 42] that support civic engagement and foster deliberative reflection in online spaces. Together, these efforts signal a growing global investment in AI-assisted civic dialogue.

Yet, despite this momentum, the field remains fragmented across disciplines and geographies. Few cross-regional venues exist for comparing findings, aligning methodologies, or articulating shared design principles. This fragmentation has resulted in persistent conceptual and methodological gaps that hinder the formation of a coherent research agenda for AI-mediated deliberation. Three key gaps stand out:

- **Ethical and normative ambiguity:** Scholars and practitioners lack consensus on the legitimate role of AI in public reasoning, whether these systems should act as facilitators that guide discussion, co-thinkers that contribute arguments, or proxies that represent collective or synthetic perspectives.
- **Design and evaluation standards:** Existing HCI frameworks have yet to adapt to the realities of human–AI co-deliberation, including the presence of synthetic participants, algorithmic mediation, and the need for transparency and explainability in deliberative processes.
- **Scalability–authenticity tension:** While AI offers unprecedented potential to scale civic participation, it also risks diluting deliberative virtues such as empathy, equality, and self-reflection that underpin meaningful democratic discourse.

Addressing these challenges requires global coordination and intellectual synthesis. By convening leading initiatives across continents, this panel seeks to lay the groundwork for an integrative research agenda, one that treats deliberation not as a fixed process but as a living design challenge in the age of AI. Specifically, the panel will open space for dialogue on foundational questions that cut across disciplinary and regional boundaries:

- How should we design for human–AI collaboration in deliberation?
- What constitutes authenticity when AI-generated arguments enter public reasoning?
- How can HCI frameworks uphold ethical legitimacy while leveraging AI's potential for inclusion and scale?

In bringing together researchers from Asia, Europe, and North America, this panel bridges emerging work on AI-mediated deliberation with broader HCI research in human–AI interaction, participatory design, and democratic technology. Moving beyond the view of AI as merely a facilitator, augments, or substitute for human participants, the discussion will probe how AI systems reshape civic reasoning, redistribute epistemic authority, and reconfigure the democratic values embedded in design.

Ultimately, this panel aims to chart a unified, global research trajectory for understanding and critically shaping the evolving relationship between AI systems and human deliberation.

3 PANEL

3.1 Panel Format and Interactivity

The panel is designed as an interactive *deliberation* about deliberation, modeling the very participatory and reflective practices under discussion. The format combines short provocations, structured debates, and audience participation, supported by live polling and optional AI-facilitated interaction. Through this design, panelists and attendees collaboratively explore the opportunities and risks of integrating AI into civic dialogue.

3.1.1 Panel Structure (90 minutes).

- (1) Opening Remarks (5 min): Moderators frame the session's core topic, introduce the panelists, and conduct the first live poll to capture audience expectations and baseline attitudes toward AI in deliberation.
- (2) Lightning Statements (25 min): Each panelist delivers a 5-minute provocation outlining a unique opportunity, ethical dilemma, or design challenge arising from AI-facilitated deliberation.
- (3) Moderated Debates (30 min): Moderators organize a structured debate among panelists and volunteer audience members, optionally supported by AI agents. Participants are divided into groups representing contrasting positions: No AI at all, AI as facilitators, and AI as augmenters or substitutes. Roles may include AI developers, researchers, policymakers, regulators, civil society actors, and citizens. Pre-designed debate questions guide each group, prompting discussion on legitimacy, transparency, and agency.
- (4) Audience Statements (15 min): Attendees break into small groups to discuss emerging themes and share short reflections. A second live poll captures shifts in opinion and highlights consensus or polarization.
- (5) Panel Reflections (10 min): Panelists respond to audience contributions, drawing connections to broader research and practice in HCI, civic tech, and AI ethics.
- (6) Closing Synthesis (5 min): Moderators summarize key tensions, emerging principles, and open questions for the CHI community. A final poll assesses attitude changes and captures commitments for future collaboration.

3.1.2 *Audience Engagement Methods.* Audience participation is central to the panel's design. Throughout the 90-minute session, *interactive polls* will be used at key moments during the opening, mid-discussion, and closing, to capture shifts in audience perspectives and foster collective reflection. *Structured debates* will invite

³<https://orbis-project.eu/>

⁴<https://multipod-project.eu/>

⁵<https://idemproject.eu/>

both panelists and attendees to deliberate from assigned stakeholder roles, such as AI developers, policymakers, researchers, and citizens, allowing participants to experience the tensions and trade-offs of human–AI collaboration firsthand. To further model the affordances of AI in civic contexts, *AI-assisted facilitation* may optionally be employed to demonstrate how agents structure, moderate or summarize real-time discussions, showcasing how automation can support or reshape deliberative dynamics. Audience members will also engage in *mini-reflection activities* within small groups, sharing concise insights or takeaways through moderators or digital platforms. *Hybrid Q & A participation* will be encouraged through both microphones and online question submissions, ensuring accessibility and inclusivity for attendees in diverse settings. Collectively, these methods transform the session from a traditional panel into a live, participatory experiment in deliberative engagement.

3.1.3 Logistical and Accessibility Considerations. The panel requires a standard audiovisual setup, including microphones for all speakers, a projector for shared visuals, and live polling software (e.g., Slido) to facilitate real-time audience interaction. Accessibility is a key priority: live captioning will be provided to support multilingual and hearing-impaired participants, and all visual materials will be designed for readability and clarity. Moderators will coordinate both on-site and online engagement, ensuring that audience questions and reflections from hybrid attendees are equally represented. Optional AI agents may be used to assist in managing discussion flow or summarizing audience input. These logistical and accessibility measures aim to create an inclusive, participatory environment that embodies the panel’s commitment to equitable and transparent deliberation.

3.2 Panelist

There will be 5 panelists. Each panelist brings a distinct perspective on the design, ethics, and implementation of AI in deliberative systems, representing expertise across academia, civic technology, and industry. The organizer will introduce each panelist, who will then discuss their experiences addressing the multifaceted challenges in AI-facilitated deliberation, from conceptual framing and technical development to deployment, evaluation, and long-term sustainability. Collectively, the panelists’ insights illuminate emerging intersections between human–AI interaction, participatory design, and democratic technology, shaping an integrative research agenda for deliberation in the age of AI.

There will be 3 moderators (2 Full Professors and 1 Postdoctoral Fellow). All three bring extensive experience in facilitating focus groups and stakeholder interviews on topics related to AI, civic engagement, and deliberation.

Professor Weiyu Zhang (Organizer, Moderator). is a Provost’s Chair Full Professor at National University of Singapore, Department of Communications and New Media. She leads the Civic Tech Lab ⁶. Zhang’s research focuses on civic engagement, Information and Communication Technologies and AI-augmented collective intelligence, with an emphasis on Asia. She has pioneered empirical studies on multi-agent deliberation simulations, AI-supported facilitation and digital democracy in Asia. She also led multi-nation

projects on youth engagement, online deliberation, and civic tech in Asia. Her current interest is to develop and examine civic tech applications to facilitate citizen deliberation on science topics such as artificial intelligence, Covid-19 vaccines, climate change, and novel food. Zhang is the author of the book “The Internet and New Social Formation in China: Fandom Publics in the Making”. Her published articles have appeared in *Journal of Communication*, *Communication Theory*, *Communication Research*, *Journal of Computer-Mediated Communication*, *Information, Communication, & Society*, *New Media & Society*, *Computers & Education*, *Public Understanding of Science*, the *ACM CHI Conference on Human Factors in Computing Systems*, and many others.

Professor Simon Tangi Perrault (Organizer, Moderator). is a Full Professor at Telecom ParisTech, France. He was previously an Assistant Professor at Singapore University of Technology and Design, where he led the SHUAILAB, and has also served as a visiting professor at KAIST, Korea, and as an Assistant Professor at Yale-NUS College, Singapore. Perrault’s research focuses on Human–Computer Interaction and Human–AI Collaboration, with a particular interest in how users engage with online content, including social media. He study ways to help users identify potentially harmful content, such as misinformation, while also supporting self-reflection when they create their own content. As a computer scientist, he designed systems that combine Natural Language Processing, Large Language Models, Multimodal Models, and interface design, emphasizing assistance that empowers users without taking control away from them. He is also interested in information visualization for spatial data, including cartogram, and in developing AI-powered tools to support collaborative qualitative analysis.

Dr. Jiaxin Pei (Panelist). is a Postdoctoral Fellow at Stanford Institute for Human-Centered Artificial Intelligence, where he is also affiliated with the Stanford Digital Economy Lab, the Stanford NLP group and the Institute for Decentralized AI. He obtained his Ph.D. from School of Information at the University of Michigan. In Fall 2026, Pei will join the School of Information at UT-Austin as an Assistant Professor. His research focuses on LLMs, HCI, and Computational Social Science. He develops systems including *Deliberation.io* to facilitate large-scale, open dialogues between citizens and strengthen collective decision-making. His current research directions include designing interfaces and systems to enable more interactive, collaborative, and trustworthy AI agents; developing standards and protocols for AI-driven digital infrastructure; building effective data pipelines for LLM training and evaluation; creating collective intelligence systems to support group collaboration and decision-making; and investigating the impact of AI and technology on individuals, communities, and society. Pei received the best student paper award at ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO), best demo paper award at AAAI Conference on Human Computation and Crowdsourcing (HCOMP), honorable mention award at International Conference on Computational Social Science (IC2S2) and best paper award at the Workshop on Social Influence in Conversations.

Professor Anna De Liddo (Panelist). is a Professor of Human-Computer Interaction at Knowledge Media Institute in The Open University (UK). Her research focuses on the study of Collective

⁶www.civictechlab.org

Intelligence facilitated through new forms of dialogue and deliberation. She leads the Knowledge Media Institute's Intelligent Deliberation Group IDea ⁷ which investigates theories methods and tools accounting for the centrality of social interaction and discourse in public engagement, urban informatics, e-democracy and social innovation contexts. Anna's research focuses on models of dialogue and argumentation; models of crowdsourcing and participatory representation; and the design, implementation and uptake of online systems that seek to increase collective environmental awareness, and collective capacity to make sense of complex issues, such as social justice and environmental sustainability. In 2019, Anna was awarded a prestigious individual fellowship from the Global Office of Naval Research (ONR Global) to research and devise the next generation of technological solutions for discussion-enhanced collective decision-making. Anna has also lead the design and development of 7 different Collective Intelligence technologies, including ORBIS and MultiPoD, which focus on creating co-designed, AI-enhanced tools for deliberative participation. In this panel, she will explore a central question: Can deliberative democratic processes scale across people, cultures, places, and environments, or are the concepts of democracy and scale inherently in tension?

Professor Efthimios (Themis) Tambouris (Panelist). is a Full Professor of Information Systems and eGovernment at the Department of Applied Informatics, University of Macedonia, Thessaloniki, Greece. Prior to that, he served at research centers NCSR 'Demokritos' and CERTH/ITI as well as the IT industry. His research interests include eGovernment, eParticipation, eLearning, eHealth and Information Systems. In the last 20 years, Tambouris has initiated, coordinated and participated in numerous research projects, including AI4Deliberation Europe, EURO-CITI, eGOV, SmartGov, OntoGov, OneStopGov, LD-Cast, DEMO-Net, eGovPolliNet, EATrain2, WAVE, OpenCube, #ask, OpenGovIntelligence, inGov and more. He has authored more than 170 research publications in eGovernment, eParticipation, eLearning. He is co-chairing the International Conference on eParticipation (ePart) and is associate editor of Government Information Quarterly (GIQ) and Digital Government (DGOV) journals.

Dr. Francesco Veri (Panelist). is a Senior Researcher at the Centre for Democracy Studies at the University of Zurich, a Senior Researcher at the Department of Political Science and International Relations at the University of Geneva, and an Associate at the Centre for Deliberative Democracy and Global Governance at the University of Canberra. Francesco's work bridges political behaviour, democratic participation, and deliberative democracy, with a growing focus on the intersection of AI and democratic governance. He combines deep knowledge of representative and direct democratic systems with methodological innovation, including discourse mapping, cultural consensus theory, and computational approaches to reasoning. Francesco also leads work on the EU Horizon project AI4Deliberation, advancing institutionalised, multimodal, gamified mass deliberation through AI tools. Francesco major contribution is the development of the Deliberative Reason Index (DRI), the sole benchmark for deliberative reasoning, published in the American Political Science Review.

⁷<http://idea.kmi.open.ac.uk/>

Professor Horacio Saggion (Panelist). is a Full Professor and Chair in Computer Science and Artificial Intelligence at the Department of Information & Communication Technologies, Universitat Pompeu Fabra (UPF), Barcelona, Spain. Horacio currently leads TALN Research Group and is the director of the Large Scale Text Understanding Systems Lab of TALN where he works in several areas of Natural Language Processing (NLP), automatic text summarization, text simplification, NLP for Sign Languages, information extraction, figurative language, sentiment analysis and related topics. His work combines symbolic/linguistic approaches and novel deep learning techniques. Horacio has published over 70 works in leading scientific journals, conferences, and books in the field of human language technology. Horacio is also the project coordinator for iDEM ⁸, which develops AI language tools to support marginalized groups, such as people with disabilities, older people and migrants in the democratic processes. In this panel, Horacio seeks to discuss the methodological approach for a more inclusive democracy.

4 EXPECTED OUTCOMES & CONTRIBUTIONS TO THE CHI COMMUNITY

This panel will make several contributions to CHI 2026 and the broader HCI community:

- (1) **Conceptual Integration:** It will synthesize fragmented research on AI-mediated deliberation across HCI, communication, and political science, establishing a unified vocabulary for human-AI civic interaction.
- (2) **Global Perspective:** By bringing together Asian, European, and North American teams, the session will foster international dialogue on regional priorities, ethical frameworks, and governance models for AI deliberation.
- (3) **Design Principles:** Through debate and reflection, the panel aims to articulate emerging principles for trustworthy deliberation interfaces, including transparency, inclusivity, reflection, and co-agency.
- (4) **Practical Outcomes:** Panelists will co-develop a post-conference summary brief and an online repository of tools, datasets, and design guidelines to support AI deliberation research.
- (5) **Community Building:** The panel will lay the foundation for a SIGCHI special interest group or recurring workshop series on AI and deliberation, sustaining long-term engagement beyond CHI 2026.
- (6) **Pedagogical Impact:** By modeling participatory deliberation during the session, it offers attendees hands-on insight into designing inclusive and reflective human-AI interactions.

Ultimately, this panel seeks to position CHI as a central venue for shaping the ethical, technical, and social foundations of AI-augmented collective intelligence, ensuring that deliberation remains human-centered while being AI-enabled.

REFERENCES

- [1] 2025. How AI Translation Can Increase Civic Engagement – wordly.ai. <https://www.wordly.ai/blog/civic-engagement/>
- [2] Aadesh Bagmar, Kevin Hogan, Dalia Shalaby, and James Purtilo. 2022. Analyzing the Effectiveness of an Extensible Virtual Moderator. 6 (2022), 18:1–18:16. Issue GROUP. doi:10.1145/3492837

⁸<https://idemproject.eu/>

- [3] John Carlo Bertot, Paul T. Jaeger, Sean Munson, and Tom Glaisyer. 2010. Social Media Technology and Government Transparency. *Computer* 43, 11 (2010), 53–59. doi:10.1109/MC.2010.325
- [4] Reuben Binns. 2018. Fairness in machine learning: Lessons from political philosophy. In *Conference on fairness, accountability and transparency*. PMLR, 149–159.
- [5] James Bohman. 2000. *Public deliberation: Pluralism, complexity, and democracy*. MIT press.
- [6] Michael A Bumpus. 2025. From prototype to persona: AI agents for decision support and cognitive extension. *Issues in Information Systems* 26, 1 (2025), 338–351.
- [7] Eshwar Chandrasekharan, Shagun Jhaver, Amy Bruckman, and Eric Gilbert. 2022. Quarantined! Examining the effects of a community-wide moderation intervention on Reddit. *ACM Transactions on Computer-Human Interaction (TOCHI)* 29, 4 (2022), 1–26.
- [8] Stephen Coleman and Giles Moss. 2012. Under construction: The field of online deliberation research. 15 pages.
- [9] Todd Davies and Reid Chandler. 2013. Online deliberation design: Choices, criteria, and evidence. *arXiv preprint arXiv:1302.5177* (2013).
- [10] Karl de Fine Licht and Jenny de Fine Licht. 2020. Artificial intelligence, transparency, and public decision-making: Why explanations are key when trying to produce perceived legitimacy. *AI & society* 35, 4 (2020), 917–926.
- [11] Anna De Liddo, Simon Buckingham Shum, Ivana Quinto, Michelle Bachler, and Lorella Cannavacciuolo. 2011. Discourse-centric learning analytics. In *Proceedings of the 1st international conference on learning analytics and knowledge*. 23–33.
- [12] Rianne Dekker and Victor Bekkers. 2015. The contingency of governments' responsiveness to the virtual public sphere: A systematic literature review and meta-synthesis. *Government Information Quarterly* 32, 4 (2015), 496–505.
- [13] John S Dryzek. 2012. *Foundations and frontiers of deliberative governance*. Oxford University Press.
- [14] Siamak Faridani, Ephrat Bitton, Kimiko Ryokai, and Ken Goldberg. 2010. Opinion space: a scalable tool for browsing online comments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1175–1184.
- [15] James Fishkin. 2009. *When the people speak: Deliberative democracy and public consultation*. Oup Oxford.
- [16] Suyash Fulay and Deb Roy. 2025. The Empty Chair: Using LLMs to Raise Missing Perspectives in Policy Deliberations. *arXiv preprint arXiv:2503.13812* (2025).
- [17] Robert E Goodin. 2000. Democratic deliberation within. *Phil. & Pub. Aff.* 29 (2000), 81.
- [18] Jarod Govers, Eduardo Velloso, Vassilis Kostakos, and Jorge Goncalves. 2024. AI-Driven Mediation Strategies for Audience Depolarisation in Online Debates. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (New York, NY, USA) (CHI '24)*. Association for Computing Machinery, 1–18. doi:10.1145/3613904.3642322
- [19] Jurgen Habermas. 1984. Communicative Action: Reason and the Rationalization of Society.
- [20] Yilun Hua, Nicholas Chernogor, Yuzhe Gu, Seoyeon Julie Jeong, Miranda Luo, and Cristian Danescu-Niculescu-Mizil. 2024. How did we get here? summarizing conversation dynamics. *arXiv preprint arXiv:2404.19007* (2024).
- [21] Lawrence R Jacobs, Fay Lomax Cook, and Michael X Delli Carpini. 2009. *Talking together: Public deliberation and political participation in America*. University of Chicago Press.
- [22] Hyunwoo Kim, Haesoo Kim, Kyung Je Jo, and Juho Kim. 2021. StarryThoughts: Facilitating Diverse Opinion Exploration on Social Issues. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW1 (2021), 1–29.
- [23] Nam Wook Kim, Jonghyuk Jung, Eun-Young Ko, Songyi Han, Chang Won Lee, Juho Kim, and Jihee Kim. 2016. Budgetmap: Engaging taxpayers in the issue-driven classification of a government budget. In *Proceedings of the 19th ACM conference on computer-supported cooperative work & social computing*. 1028–1039.
- [24] Soomin Kim, Jinsu Eun, Joseph Seering, and Joonhwan Lee. 2021. Moderator Chatbot for Deliberative Discussion: Effects of Discussion Structure and Discussant Facilitation. 5 (2021), 87:1–87:26. Issue CSCW1. doi:10.1145/3449161
- [25] Mark Klein. 2012. Enabling large-scale deliberation using attention-mediation metrics. *Computer Supported Cooperative Work (CSCW)* 21, 4 (2012), 449–473.
- [26] Travis Kriplean, Jonathan Morgan, Deen Freelon, Alan Borning, and Lance Bennett. 2012. Supporting reflective public thought with considerit. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work*. 265–274.
- [27] Frank L. 2025. Understanding AI Content Moderation: Types & How it Works – getstream.io. <https://getstream.io/blog/ai-content-moderation/>.
- [28] Mina Lee, Percy Liang, and Qian Yang. 2022. Coauthor: Designing a human-ai collaborative writing dataset for exploring language model capabilities. In *Proceedings of the 2022 CHI conference on human factors in computing systems*. 1–19.
- [29] Sanju Menon, Weiyu Zhang, and Simon T Perrault. 2020. Nudge for deliberativeness: How interface features influence online discourse. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [30] Sean A Munson and Paul Resnick. 2010. Presenting diverse political opinions: how and how much. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 1457–1466.
- [31] Tom Murray, Lynn Stephens, Beverly Park Woolf, Leah Wing, Xiaoxi Xu, and Natasha Shrikant. 2013. Supporting social deliberative skills online: the effects of reflective scaffolding tools. In *Online Communities and Social Computing: 5th International conference, OCSC 2013, Held as Part of HCI International 2013, Las Vegas, NV, USA, July 21–26, 2013. Proceedings* 5. Springer, 313–322.
- [32] Claudio Novelli, Javier Argota Sánchez-Vaquero, Dirk Helbing, Antonino Rotolo, and Luciano Floridi. 2025. A replica for our democracies? On using digital twins to enhance deliberative democracy. *AI & SOCIETY* (2025), 1–19.
- [33] Cassandra Overney. 2025. Designing for Constructive Civic Communication: A Framework for Human-AI Collaboration in Community Engagement Processes. *arXiv preprint arXiv:2505.11684* (2025).
- [34] Cassandra Overney, Hang Jiang, Urooj Haider, Cassandra Moe, Jasmine Mangat, Frank Pantano, Effie G McMillian, Paul Riggins, and Nabeel Gillani. 2025. Human-AI Narrative Synthesis to Foster Shared Understanding in Civic Decision-Making. *arXiv preprint arXiv:2509.19643* (2025).
- [35] Joon Sung Park, Carolyn Q Zou, Aaron Shaw, Benjamin Mako Hill, Carrie Cai, Meredith Ringel Morris, Robb Willer, Percy Liang, and Michael S Bernstein. 2024. Generative agent simulations of 1,000 people. *arXiv preprint arXiv:2411.10109* (2024).
- [36] Ruth Shortall, Anatol Itten, Michiel van der Meer, Pradeep Murukanaiah, and Catholijn Jonker. 2022. Reason against the machine? Future directions for mass online deliberation. *Frontiers in Political Science* 4 (2022), 946589.
- [37] Maria A Simone. 2010. Deliberative democracy online: Bridging networks with digital technologies. *The Communication Review* 13, 2 (2010), 120–139.
- [38] Michael Henry Tessler, Michiel A Bakker, Daniel Jarrett, Hannah Sheahan, Martin J Chadwick, Raphael Koster, Georgina Evans, Lucy Campbell-Gillingham, Tantum Collins, David C Parkes, et al. 2024. AI can help humans find common ground in democratic deliberation. *Science* 386, 6719 (2024), eadq2852.
- [39] Scott Wright and John Street. 2007. Democracy, deliberation and design: the case of online discussion forums. *New media & society* 9, 5 (2007), 849–869.
- [40] Lu Xiao, Weiyu Zhang, Anna Przybylska, Anna De Liddo, Gregorio Convertino, Todd Davies, and Mark Klein. 2015. Design for online deliberative processes and technologies: Towards a multidisciplinary research agenda. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems*. 865–868.
- [41] ShunYi Yeo, Zhuoqun Jiang, Anthony Tang, and Simon Tangi Perrault. 2025. Enhancing Deliberativeness: Evaluating the Impact of Multimodal Reflection Nudges. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–26.
- [42] ShunYi Yeo, Gionnieve Lim, Jie Gao, Weiyu Zhang, and Simon Tangi Perrault. 2024. Help me reflect: Leveraging self-reflection interface nudges to enhance deliberativeness on online deliberation platforms. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–32.
- [43] Ann Yuan, Andy Coenen, Emily Reif, and Daphne Ippolito. 2022. Wordcraft: story writing with large language models. In *Proceedings of the 27th International Conference on Intelligent User Interfaces*. 841–852.
- [44] Angie Zhang, Olympia Walker, Kaci Nguyen, Jiajun Dai, Anqing Chen, and Min Kyung Lee. 2023. Deliberating with AI: improving decision-making for the future through participatory AI design and stakeholder deliberation. *Proceedings of the ACM on Human-Computer Interaction* 7, CSCW1 (2023), 1–32.
- [45] Weiyu Zhang, Xiaoxia Cao, and Minh Ngoc Tran. 2013. The structural features and the deliberative quality of online discussions. *Telematics and informatics* 30, 2 (2013), 74–86.
- [46] Yuanhao Zhang, Xiaoyu Wang, Jiaxiong Hu, Ziqi Pan, Zhenhui Peng, and Xiaojuan Ma. 2025. Designing AI-Infused Interactive Systems for Online Communities: A Systematic Literature Review. *arXiv preprint arXiv:2509.23309* (2025).