Social Choice with Text: Collective Decision Making in the LLM Era

Umberto Grandi✉
IRIT, University Toulouse Capitole
Toulouse, France
umberto.grandi@irit.fr

Robert Bredereck
Technische Universität Clausthal
Clausthal-Zellerfeld, Germany
robert.bredereck@tu-clausthal.de

Théo Delemazure
University of Amsterdam
Amsterdam, The Netherlands
theo.delemazure@uva.nl

Ulle Endriss
University of Amsterdam
Amsterdam, The Netherlands
ulle.endriss@uva.nl

Jan Maly
WU Vienna University of Economics and Business
Wien, Austria
jan.maly@wu.ac.at

Nicholas Mattei✉
Tulane University
New Orleans, USA
nsmattei@tulane.edu

Nicolas Maudet
Sorbonne Université
Paris, France
nicolas.maudet@lip6.fr

Oliviero Nardi
WU Vienna University of Economics and Business
Wien, Austria
oliviero.nardi@wu.ac.at

Stanisław Szufa
University of Geneva
Geneva, Switzerland
stanislaw.szufa@unige.ch

Abstract

The recent rapid advances in natural language processing (NLP) and the development of large language models (LLMs) open up new research directions broadly across the field of AI, including the key area of decision making. In this position paper we frame a burgeoning research agenda by charting the significant range of opportunities and existing work for integrating text-processing algorithms into the design and analysis of mechanisms for collective decision making and in democratic practices such as citizen assemblies and large-scale participatory initiatives.

ACM Reference Format:
Umberto Grandi, Robert Bredereck, Théo Delemazure, Ulle Endriss, Jan Maly, Nicholas Mattei, Nicolas Maudet, Oliviero Nardi, and Stanisław Szufa. 2026. Social Choice with Text: Collective Decision Making in the LLM Era. In *ACM AI Summit 2026 (AI Summit '26), August 31-September 02, 2026, Atlanta, GA, USA.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3806096.3844825

1 Introduction

Artificial Intelligence (AI) and systems powered by large language models (LLMs) are being rapidly deployed and are already impacting numerous aspects of our societies, and collective decision making—from national elections to board committees, to offline and online deliberations—have not been spared [15, 16]. Large scale platforms are now in use all around the world. Polis [8] has been used by governments in more than 15 countries; project Jigsaw deployed in numerous cities including Bowling Green, KY, USA, where almost 8,000 participants proposed and discussed nearly 4,000 ideas regarding the future of their county [23]; and the Habermas Machine [44] which enabled large scale deliberation in the UK. These new platforms merge classical deliberative techniques with LLMs, where the LLMs are used to structure debate, aggregate insights, and produce reports. These systems originate both from academic consortia [38] and practitioners' associations [41]. The field of political science is debating and experimenting to identify the effects of LLMs in deliberative assemblies [20, 31] and in democratic practices [1].

Vision. The underlying ideology of these systems is that through larger, more inclusive deliberation, societies can arrive at better decisions [15]. However, beyond aggregating text, making a collective decision requires expressing a preference and making a (social) choice through a reasoned methodology, something that traditionally is done through a voting system [2, 5]. **In this paper we argue for more research at the intersection of social choice, which provides concrete, explainable, and principled methodologies for working with preferences, and natural language processing, whose recent advances, including LLMs, make it possible to summarize, explore, and generate textual information from stakeholders. We believe that this combination can bring about more informed and more representative collective decisions at scales never seen before**.

We see two primary features that make LLMs particularly fit for disrupting current democratic practices. First, they are easily accessible, configurable, and able to process text at a scale that previously required advanced expertise in traditional natural language processing (NLP) and data processing techniques. While we associate elections with the act of voting, a non-textual action, this is just the tip of the iceberg of discussions and deliberation in textual or speech form at web scale. Second, they revive the long-standing promise that digital technology could scale up the positives of small groups' deliberation and decision making [27].



This work is licensed under a Creative Commons Attribution 4.0 International License.
AI Summit '26, Atlanta, GA, USA
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2638-5/2026/08
https://doi.org/10.1145/3806096.3844825

However, collective decisions and democratic practices are often sensitive topics in which pluralism must be preserved [47] and where oversight and transparency is of the utmost importance [15]. And the use of LLMs and other large-scale text processing techniques is, at least for now, lacking on both accounts, leading to significant areas of research on explainability and causality for these methods [37]. Frontier models are often commercial and rarely open-source, and the science of obtaining formal safety bounds and explanations for LLM outputs is still evolving.

A promising pathway for positive and productive LLM usage in participatory practices and collective decision making lies in integrating the centuries-old theory of social choice [2, 5] with the most recent computational text-processing developments. Recent work has already begun to explore this approach. For example, Fish et al. [14] derive safety bounds from LLM errors when selecting representative statements; Tessler et al. [44] employ a clone-proof voting rule within an LLM pipeline to construct consensual statements in caucus deliberation; and Chooi et al. [7] use voting and LLMs to find novel, agreeable statements from textual preferences. By surveying recent advances and identifying future directions, we frame a new research agenda that bridges expertise in discrete structures, such as votes, and statistical algorithms in NLP.

Collective decision making in AI has traditionally been analyzed via social choice theory [2], where agents express preferences over alternatives using ballots (rankings, approvals), and voting rules aggregate them into collective decisions. Computational social choice (COMSOC) has thoroughly studied computational problems in this space, from outcome computation to strategic manipulation [5]. However, classical models treat alternatives as abstract, independent symbols, whereas real-world alternatives (e.g., in participatory budgeting or on deliberation platforms) are rich semantic objects with titles, descriptions, locations, categories, and political affiliations, enabling humans to recognize duplicates or infer opinions, tasks that are rendered hard or impossible in the classical abstract modeling. This calls for an enrichment of classical models of social choice by incorporating such textual information. While this would have seemed an unrealistic objective even just a few years ago, recent advances in AI have made powerful text-processing technologies accessible to both researchers and the public at large.

Roadmap. To structure this growing research agenda, we use a schematic (and simplified) view of collective decision making, starting with the elicitation of voter preferences over alternatives, followed by the use of a suitable aggregation function to produce outcomes, all of which is governed by a number of high-level normative principles (so-called *axioms*). Each of the next three sections details how NLP techniques, and LLMs in particular can be used to advance research and development in this domain (see Figure 1). Further, to make this research agenda concrete, we review the currently available datasets that combine user preferences with textual components, and we conclude with a critical assessment of the risks and limitations of this ambitious research program.

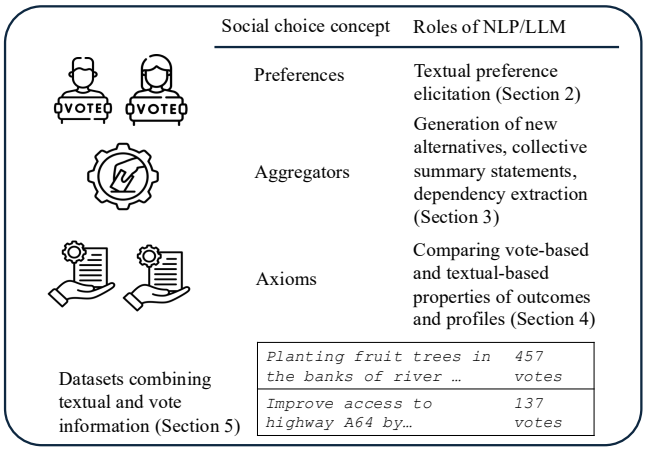


Figure 1: Schematic view of collective decision making as aggregation of preferences, elicited from voters, into outcomes, with potential contributions of NLP/LLM at each level.

2 Textual Preference Elicitation

The elicitation of preferences is usually the first direct contact between a stakeholder (voter) and a collective decision making problem. During elicitation, voters discover the alternatives, form their opinion, and eventually express that opinion. In many real-world applications, this may be the hardest part of the process: voters may not discover relevant options, may not be able to make meaningful comparisons, or misunderstand constraints.

LLMs can provide a powerful interface layer that mediates between voters and formal mechanisms to tackle such problems. To cite some examples, Huang et al. [21] experiment with textual preference elicitation in the related field of combinatorial auctions, comparing classical ML-based query generation approaches to LLM agents that can interact with bidders using text. When aggregating textual statements (a problem we will discuss in Section 3), most approaches train and assess a specific LLM for each user to be able to infer the user's preferences over new statements [14, 44]. Mohsin et al. [32] experiment with a random-utility model to learn from users' textual comments over alternatives before preference aggregation. Other works have used NLP techniques, including transformers, to directly model preferences for auctions and other settings [39] or to generate dependencies between alternatives using text descriptions [46]. Finally, reconstructing voter preferences from demographic data has been tested in participatory budgeting [49] and pairwise voting [19].

A number of research directions can be further explored in the elicitation stage. First, LLM agents can help users identify relevant alternatives. In participatory budgeting, for example, they can help users in dealing with the overflow of information, by presenting them with key facts about certain projects, by providing concise summaries of projects in a neighborhood, and by catching duplicates (i.e., near-identical proposals). Second, LLMs can help voters understand the implications of selecting particular alternatives. For instance, by clarifying who benefits and what are the consequences of a given choice, allowing for more informed and conscious decisions, e.g., CivicChats.org. Third, a suitably designed LLM interface can support the elicitation of preferences, for example, by asking comparative questions, reducing cognitive load; we are beginning to see work in this direction [33]. This can take an iterative form,

with the interface presenting to the user their final ballot highlighting trade-offs and constraints, allowing them to update their vote before it is submitted [24].

3 Text-Aware Aggregation

The design and study of aggregation methods is at the core of collective decision making studies. This ranges from fair participatory budgeting rules, that aggregate approval votes over projects into a set of projects to be funded, to the design of processes merging LLM outputs and votes to obtain consensual statements on questions of public policy. While aggregation is typically associated with a collective decision, it is also widely used in support of deliberative processes, charting the opinions of a group or extracting salient or conflicting points in a debate.

Incomplete preferences are the first problem that can be tackled with a clever mixture of preference (vote) and textual data. In many applications the preference data is highly incomplete, as the number of comments (and ordering thereof) is often too large for any user to view even a significant percentage of the contributions. Textual descriptions on alternatives can be used to find similar proposal or filling in preferences [46], or to extract logical connections between comments, such as whether one argument attacks another—a problem extensively studied by the argumentation mining community [17, 28]. Based on such logical structures, researchers could go further by proposing logically consistent recommendations using methods from judgment aggregation [18] or solve the problem of incomplete vote information by expanding voters' partially expressed preferences into maximally consistent viewpoints [3]. Finally, LLMs can provide topical clusterings of the comments [7, 29], which we can then use to constrain outcomes to ensure full coverage of the topics discussed. The added value of combining NLP and social choice techniques is to obtain explainable and auditable outcomes, as the extracted structure between the comments can be made understandable and interpretable for humans, while the social choice mechanisms using these structures can remain white-box and deterministic. The remaining challenge is the representation and use of conflicting or contentious information, issues that are gaining traction in the NLP community [25, 48].

One could also directly aggregate pieces of text. This is a prominent direction that is likely to impact democratic practices in the coming years. Such LLM-generated consensus statements can play an important role in deliberative exercises, where the perceived neutrality of an algorithm-generated statement with respect to a human-generated one makes them particularly fitting for heavily polarized crowds, contributing to efforts to develop technology that allows deliberation to scale. Work in this area initiated with a large-scale experiment testing the LLM effect in so-called "caucus deliberations" [44]. In these experiments, users submit textual opinions to an LLM that summarizes them into LLM-generated consensus statements. They show that this process improves the acceptance of consensual statements relative to human mediators. Further, Konya et al. [26] report on how LLM-generated consensual statements can be useful in finding common ground between two populations in conflict. Closer to social choice, Fish et al. [14]—starting from a set of statements proposed by users—tested the use of LLMs to generate new representative statements to be submitted to a vote among users, complementing their empirical findings with provable mathematical bounds. More recently, LLM-generated summaries have been tested in a formal coalition formation model [6], discussed as a potential improvement for the Pol.is platform [41], and implemented in a civic planning context [23]. Going further, formal guarantees on the aggregation can be studied by borrowing and adapting concepts from social choice theory, in line with the initial work of Fish et al. [14].

4 Axiomatic and Textual Analysis

Central to classical social choice theory is the axiomatic method [2], where we encode desirable properties of mechanisms, such as fairness, in precise mathematical terms to then classify mechanisms in view of the properties they satisfy. Modern NLP techniques provide opportunities for complementing and enriching this axiomatic analysis. Thus, while in earlier sections our focus has been on opportunities for innovations aimed at enhancing various stages of the process of collective decision making, here we instead identify opportunities for enhancing our capabilities for evaluating such mechanisms and the outcomes they produce.

The axiomatic method sometimes produces theoretical results that arguably are too abstract to provide true insights regarding real-world decision making, due to alternatives being treated as abstract mathematical objects. Textual analysis offers significant opportunities for modelling both the internal structure of alternatives and the broader context in which a decision is to be taken.

First, voting data that includes textual descriptions of alternatives could be used to investigate whether the normative principles encoded by certain axioms are also reflected at the level of these textual descriptions. For instance, it has been observed that sophisticated methods for participatory budgeting that guarantee outcomes to be *proportional* relative to the preferences reported by voters also tend to achieve a balanced spread of funds between different categories of projects [36, 40]. A category label is a very simple form of providing textual information. Does this kind of effect generalise to richer textual descriptions?

Second, it can be desirable to refine classical axioms by referencing the textual descriptions of alternatives. For instance, the axiom of *neutrality* stipulates that we should treat all alternatives the same, which certainly is a reasonable baseline to start from. But when alternatives come with textual descriptions, we might want to relax this requirement and treat two alternatives the same only if they have been described at a similar level of detail.

Third, textual analysis might allow us to refine algorithms for the detection of so-called *clones* [45]. Two alternatives are clones if they are ranked similarly by all voters. At the axiomatic level, introducing a clone into an election should not have a negative impact on the alternative being cloned. If alternatives are associated with textual information, we might want to extend the notion of clones to alternatives with similar descriptions.

Fourth, to date the formulation of new axioms has been the prerogative of experts. Can we design systems that would allow laypeople to formulate requirements relevant to them in natural language—to then be translated into formal axioms that can be used to recommend or reject a given voting rule?

Fifth, while there has been recent work on automatically generating explanations for why a given election outcome is an adequate outcome, with those explanations being grounded in axioms [4], the ability to generate textual descriptions of existing axioms would further enhance the accessibility of such explanations.

Finally, while our focus here is on the potential of using algorithms for processing text to enhance algorithms for aggregating preferences, this really is a two-way street. Indeed, as argued elsewhere [9], the axiomatic method can be an important tool for guiding aggregation processes taking place—explicitly or implicitly—in AI systems powered by NLP techniques, with the clearest example being the aggregation processes occurring during the fine-tuning of LLMs using reinforcement learning with human feedback (RLHF).

5 Datasets Combining Text and Votes

The public library of preference data Preflib.org [30] enabled major advances in the field of computational social choice by providing standardized benchmark datasets of voter preferences, typically stored as rankings over alternatives. It was recently complemented by Pabulib.org [12], a library of preference data specific to participatory budgeting. Datasets from both libraries are well suited for testing the computational efficiency of algorithms, or for comparing statistical preference generation models with real-world data [42]. Such tests do not require any textual information. But current technology enables the processing of enriched datasets also containing textual descriptions of the alternatives.

Without claiming to be exhaustive, we highlight some publicly available datasets that might be of interest to researchers wanting to take up some of the research challenges outlined in this paper:

- Deliberative platforms like Pol.is [8] and Decidim [10] enable discussion of divisive topics, comment posting, and voting on comments (approve, disapprove, neutral, abstain).
- Collective selection of consensual statements [11, 13, 14, 44]: A group of users give their textual opinions on a societal question, and these opinions are merged using an LLM into a slate of consensual statements. These statements are then ranked or evaluated by the users.
- Pairwise comparisons of political proposals [34, 35]: These proposals have been extracted from political manifestos of presidential candidates in France and Brazil and then submitted to the population by means of a voting platform.

The various datasets of debates in national and supra-national parliaments, and of the votes of members of parliament on bills and amendments also constitute important sources of votes over textual statements. We also note that classical participatory budgeting datasets can easily be augmented by adding project titles and project descriptions to Pabulib.org datasets [50].

6 Risks, Limitations, and Opportunities

We have argued that there are many opportunities for incorporating NLP techniques and LLMs into research on social choice, with some exciting initial work already taking place today. However, as many of these problems lie at the intersection of AI and democratic studies, the research program we propose necessitates substantial critical reflection. Given the high-stakes nature of many collective decision-making tasks such as electing leaders and funding projects, many of the real-world applications of this line of work fall under the highest-risk category of the European AI Act, requiring strict regulatory oversight. While combining textual information with votes could offer insights on how to obtain formal guarantees regarding the quality of decision making [14, 21], these rely on the accuracy of NLP and LLM-based techniques for preference extraction and representation, which remains difficult to ensure.

LLMs can exhibit (sometimes trivial) biases that may compromise participatory or deliberative processes. There also is a significant over-representation of English-language data in the training sets of most LLMs, which can decrease accuracy and introduce discriminatory effects for certain groups [22]. Moreover, even in the simplest settings, LLMs will unavoidably make mistakes when extracting preferences, or other relevant structures, from textual input [43].

In adopting a responsible approach, researchers need to design systems that are efficient both in terms of the number of LLM calls and the computational tractability of the applied social-choice method. It is too easy to depend on extensive LLM queries without considering the computational resources required to support them.

In summary, we have identified multiple opportunities where techniques from the broad area of NLP can be used to enhance the scope and scale of research in collective decision making. Instead of naively "throwing a bunch of data at an LLM", hoping that the LLM will correctly "infer" the collective will, we see these techniques as an opportunity to greatly expand how we elicit and reason with preferences. Contemporary NLP techniques have the potential to make it easier for people to communicate their preferences on complex questions, bring more stakeholders into democratic deliberation, and to enrich the design space available to researchers in social choice and democratic innovations. Overcoming the aforementioned challenges will enable collective decision making research to expand its role in society, reach more communities, and allow this research to have broad impact for social good.

Acknowledgments

This work was initiated at the Dagstuhl Seminar on "Societal Impact of Computational Social Choice" (https://www.dagstuhl.de/25401) in 2025. UG, TD, and UE acknowledge ERC Synergy Grant ADDI (http://doi.org/10.3030/101166894). RB acknowledges DFG project COMSOC-MPMS (grant 465371386). JM acknowledges Austrian Science Fund (FWF) grants 10.55776/PAT7221724 and 10.55776/COE12. N. Mattei acknowledges NSF grants IIS-RI-2134857, IIS-RI-2339880, and CNS-SCC-242723. N. Maudet acknowledges French National Research Agency grant ANR-23-IACL-0007 (AI Cluster PostGenIA).

References

[1] Goshi Aoki. 2024. Large language models in politics and democracy: A comprehensive survey. *arXiv preprint arXiv:2412.04498* (2024).
[2] Kenneth J. Arrow, Amartya K. Sen, and Kotaro Suzumura (Eds.). 2002. *Handbook of Social Choice and Welfare.* North-Holland.
[3] Michael Bernreiter, Jan Maly, Oliviero Nardi, and Stefan Woltran. 2024. Combining Voting and Abstract Argumentation to Understand Online Discussions. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS).*
[4] Arthur Boixel, Ulle Endriss, and Oliviero Nardi. 2022. Displaying Justifications for Collective Decisions. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI).*
[5] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (Eds.). 2016. *Handbook of Computational Social Choice.* Cambridge Univ. Press.

[6] Eyal Briman, Ehud Shapiro, and Nimrod Talmon. 2025. AI-Generated Compromises for Coalition Formation: Modeling, Simulation, and a Textual Case Study. In *Proceedings of the 20th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*.

[7] Jay Chooi, Paul Gölz, Ariel D Procaccia, Benjamin Schiffer, and Shirley Zhang. 2026. Finding Common Ground in a Sea of Alternatives. In *Proceedings of the 27th ACM Conference on Economics and Computation (EC)*.

[8] Computational Democracy Project. 2024. Open Pol.is Data. Available at https://github.com/compdemocracy/openData. Last accessed on 2026-09-28.

[9] Vincent Conitzer, Rachel Freedman, Jobst Heitzig, Wesley H. Holliday, Bob M. Jacobs, Nathan Lambert, Milan Mossé, Eric Pacuit, Stuart Russell, Hailey Schoelkopf, Emanuel Tewolde, and William S. Zwicker. 2024. Position: Social Choice Should Guide AI Alignment in Dealing with Diverse Human Feedback. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*.

[10] Decidim.Barcelona. 2024. MetaDecidim Data. Available at https://github.com/elaragon/metadecidim/tree/master. Last accessed on 2026-09-28.

[11] DeepMind Technologies Limited. 2024. Habermas Machine Dataset. Available at https://github.com/google-deepmind/habermas_machine. Last accessed on 2026-09-28.

[12] Piotr Faliszewski, Jarosław Flis, Dominik Peters, Grzegorz Pierczyński, Piotr Skowron, Dariusz Stolicki, Stanisław Szufa, and Nimrod Talmon. 2023. Participatory Budgeting: Data, Tools and Analysis. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*.

[13] Sara Fish and Paul Gölz et al. 2025. Generative Social Choice Open Datasets. Available at https://github.com/generative-social-choice/survey_data/. Last accessed on 2025-11-20.

[14] Sara Fish, Paul Gölz, David C. Parkes, Ariel D. Procaccia, Gili Rusak, Itai Shapira, and Manuel Wuthrich. 2026. Generative Social Choice. *J. ACM* 73, 2 (2026), 11:1–11:52.

[15] James S Fishkin. 2025. *Can deliberation cure the ills of democracy?* Oxford University Press.

[16] Beth Goldberg, Diana Acosta-Navas, Michiel Bakker, Ian Beacock, Matt Botvinick, Prateek Buch, Renée DiResta, Nandika Donthi, Nathanael Fast, Ravi Iyer, Zaria Jalan, Andrew Konya, Grace Kwak Danciu, Hélène Landemore, Alice Marwick, Carl Miller, Aviv Ovadya, Emily Saltz, Lisa Schirch, Dalit Shalom, Divya Siddarth, Felix Sieker, Christopher Small, Jonathan Stray, Audrey Tang, Michael Henry Tessler, and Amy Zhang. 2026. AI and the future of digital public squares. *Collective Intelligence* 5, 2 (2026).

[17] Deniz Gorur, Antonio Rago, and Francesca Toni. 2025. Can Large Language Models Perform Relation-Based Argument Mining?. In *Proceedings of the 31st International Conference on Computational Linguistics (COLING)*. Association for Computational Linguistics.

[18] Davide Grossi and Gabriella Pigozzi. 2014. *Judgment Aggregation: A Primer*. Morgan & Claypool.

[19] Jairo Gudiño-Rosero, Umberto Grandi, and César Hidalgo. 2024. Large Language Models (LLMs) as Agents for Augmented Democracy. *Philosophical Transactions A* 382, 2285 (2024), 20240100.

[20] Damith Herath, Hans Asenbaum, Janie Busby Grant, Maleen Jayasuriya, Emanuela Savini, and Harshith Ghanta. 2026. Robots in the Room: A Qualitative Exploration of Democratic Agency and AI-Augmented Deliberations. In *Proceedings of the 21st ACM/IEEE International Conference on Human-Robot Interaction*.

[21] David Huang, Edwin Lock, Francisco Marmolejo-Cossío, and David C. Parkes. 2025. Accelerated Preference Elicitation with LLM-Based Proxies. In *Proceedings of the 21st Conference on Web and Internet Economics (WINE)*.

[22] Suramya Jadhav, Abhay Shanbhag, Amogh Thakurdesai, Ridhima Sinare, and Raviraj Joshi. 2025. On Limitations of LLM as Annotator for Low Resource Languages. In *Proceedings of the 8th International Conference on Natural Language and Speech Processing (ICNLSP)*.

[23] Jigsaw-Google. 2025. Reimagining the Town Hall Meeting for the AI Age. https://jigsaw.google/our-work/reimagining-the-town-hall-meeting/. Accessed: 2026-09-28.

[24] Jeffrey O. Kephart, Nicholas Mattei, and Francesca Rossi. 2019. Collective Decision Making by Consensus in Cognitive Environments. US Patent Application 15/845,509.

[25] To Eun Kim and Fernando Diaz. 2025. Towards Fair RAG: On the Impact of Fair Ranking in Retrieval-Augmented Generation. In *Proceedings of the 2025 International ACM SIGIR Conference on Innovative Concepts and Theories in Information Retrieval (ICTIR)*.

[26] Andrew Konya, Luke Thorburn, Wasim Almasri, Oded Adomi Leshem, Ariel Procaccia, Lisa Schirch, and Michiel Bakker. 2025. Using Collective Dialogues and AI to Find Common Ground between Israeli and Palestinian Peacebuilders. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency (FAccT)*.

[27] Hélène Landemore. 2022. Can AI bring deliberative democracy to the masses?. In *HAI Weekly Seminar, Working Paper*. 166–191.

[28] John Lawrence and Chris Reed. 2019. Argument Mining: A Survey. *Computational Linguistics* 45, 4 (2019), 765–818.

[29] Jianghan Liu, Ziyu Shang, Wenjun Ke, Peng Wang, Zhizhao Luo, Jiajun Liu, Guozheng Li, and Yining Li. 2025. LLM-Guided Semantic-Aware Clustering for Topic Modeling. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (ACL)*.

[30] Nicholas Mattei and Toby Walsh. 2013. PrefLib: A Library of Preference Data (preflib.org). In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT)*. Springer.

[31] Sammy McKinney. 2024. Integrating artificial intelligence into citizens' assemblies: Benefits, concerns and future pathways. *Journal of Deliberative Democracy* 20, 1 (2024).

[32] Farhad Mohsin, Lei Luo, Wufei Ma, Inwon Kang, Zhibing Zhao, Ao Liu, Rohit Vaish, and Lirong Xia. 2021. Making Group Decisions from Natural Language-Based Preferences. In *Proceedings of the 8th International Workshop on Computational Social Choice (COMSOC)*.

[33] Ali Montazeralghaem, Guy Tennenholtz, Craig Boutilier, and Ofer Meshi. 2025. Asking Clarifying Questions for Preference Elicitation With Large Language Models. *arXiv preprint arXiv:2510.12015* (2025).

[34] Carlos Navarrete, Mariana Macedo, Rachael Colley, and Jingling Zhang et al. 2023. Data Collection of two Real-World Digital Democracy Systems. Available at https://doi.org/10.7910/DVN/8E0EA4. Last accessed on 2026-09-28.

[35] Carlos Navarrete, Mariana Macedo, Rachael Colley, Jingling Zhang, Nicole Ferrada, Maria Eduarda Mello, Rodrigo Lira, Carmelo Bastos-Filho, Umberto Grandi, Jérôme Lang, and César A. Hidalgo. 2024. Understanding Political Divisiveness Using Online Participation Data from the 2022 French and Brazilian Presidential Elections. *Nature Human Behaviour* 8, 1 (2024), 137–148.

[36] Pelle Nelissen. 2023. An Empirical Analysis of Participatory Budgeting in Amsterdam. *arXiv preprint arXiv:2310.18033* (2023).

[37] Shangrui Nie, Kian Omoomi, Lucie Flek, Zhixue Zhao, and Charles Welch. 2026. PERSPECTRA: A Scalable and Configurable Pluralist Benchmark of Perspectives from Arguments. *arXiv preprint arXiv:2602.08716* (2026).

[38] Jiaxin Pei, José Ramón Enríquez, Umar Patel, Alia Braley, Nuole Chen, Lily Tsai, and Alex Pentland. 2025. *DELIBERATION.IO: Facilitating Democratic and Civil Engagement at Scale with Open-Source and Open-Science*. Technical Report. Working Paper.

[39] Neehar Peri, Michael Curry, Samuel Dooley, and John Dickerson. 2021. PreferenceNet: Encoding Human Preferences in Auction Design with Deep Learning. In *Proceedings of the 35th International Conference on Neural Information Processing Systems (NeurIPS)*.

[40] Dominik Peters and Piotr Skowron. 2025. Website on the Method of Equal Shares for Participatory Budgeting. Available at https://equalshares.net/. Last accessed on 2025-11-28.

[41] Christopher T. Small, Ivan Vendrov, Esin Durmus, Hadjar Homaei, Elizabeth Barry, Julien Cornebise, Ted Suzman, Deep Ganguli, and Colin Megill. 2023. Opportunities and Risks of LLMs for Scalable Deliberation with Polis. *arXiv preprint arXiv:2306.11932* (2023).

[42] Stanisław Szufa, Niclas Boehmer, Robert Bredereck, Piotr Faliszewski, Rolf Niedermeier, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. 2025. Drawing a Map of Elections. *Artificial Intelligence* 343 (2025), 104332.

[43] Zhen Tan, Dawei Li, Song Wang, Alimohammad Beigi, Bohan Jiang, Amrita Bhattacharjee, Mansooreh Karami, Jundong Li, Lu Cheng, and Huan Liu. 2024. Large Language Models for Data Annotation and Synthesis: A Survey. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.

[44] 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, Matthew Botvinick, and Christopher Summerfield. 2024. AI can help humans find common ground in democratic deliberation. *Science* 386, 6719 (2024), eadq2852.

[45] T. Nicolaus Tideman. 1987. Independence of Clones as a Criterion for Voting Rules. *Social Choice and Welfare* 4, 3 (1987), 185–206.

[46] Rohith Dwarakanath Vallam, Ramasuri Narayanam, Srikanth G. Tamilselvam, Nicholas Mattei, Sudhanshu S. Singh, Shweta Garg, and Gyana R. Parija. 2019. DeepAggregation: A New Approach for Aggregating Incomplete Ranked Lists using Multi-Layer Graph Embedding. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*.

[47] Angelina Wang, Jamie Morgenstern, and John P Dickerson. 2025. Large language models that replace human participants can harmfully misportray and flatten identity groups. *Nature Machine Intelligence* 7, 3 (2025), 400–411.

[48] Han Wang, Archiki Prasad, Elias Stengel-Eskin, and Mohit Bansal. 2025. Retrieval-Augmented Generation with Conflicting Evidence. *arXiv preprint arXiv:2504.13079* (2025).

[49] Joshua C. Yang, Damian Dalisan, Marcin Korecki, Carina I. Hausladen, and Dirk Helbing. 2024. LLM Voting: Human Choices and AI Collective Decision-Making. In *Proceedings of the 7th AAAI/ACM Conference on AI, Ethics, and Society (AIES)*.

[50] Juan Zambrano, Clément Contet, Jairo Gudiño-Rosero, Felipe Garrido-Lucero, Umberto Grandi, and César A. Hidalgo. 2026. Algorithmic Shortlisting in Participatory Budgeting. In *Proceedings of the 9th AAAI/ACM Conference on AI, Ethics, and Society (AIES)*.