



**HAL**  
open science

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

Umberto Grandi, Robert Brederick, Théo Delemazure, Ulle Endriss, Jan Maly, Nick Mattei, Nicolas Maudet, Oliviero Nardi, Stanislaw Szufa

► **To cite this version:**

Umberto Grandi, Robert Brederick, Théo Delemazure, Ulle Endriss, Jan Maly, et al.. Social Choice with Text: Collective Decision Making in the LLM Era. 2026. hal-05548117

**HAL Id: hal-05548117**

**<https://hal.science/hal-05548117v1>**

Preprint submitted on 11 Mar 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.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

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

Umberto Grandi<sup>1</sup>, Robert Brederick<sup>2</sup>, Théo Delemazure<sup>3</sup>, Ulle Endriss<sup>3</sup>, Jan Maly<sup>5</sup>, Nick Mattei<sup>6</sup>, Nicolas Maudet<sup>7</sup>, Oliviero Nardi<sup>8</sup>, and Stanisław Szufa<sup>9</sup>

<sup>1</sup>IRIT, University Toulouse Capitole, France

<sup>2</sup>Technische Universität Clausthal, Germany

<sup>3</sup>University of Amsterdam, The Netherlands

<sup>5</sup>Wirtschaftsuniversität Wien, Austria

<sup>6</sup>Tulane University, USA

<sup>7</sup>Sorbonne Université, France

<sup>8</sup>TU Wien, Austria

<sup>9</sup>Université Paris Dauphine–PSL, France

3 March 2026

## Abstract

The recent rapid advances in natural language processing (NLP) and the development of large language models (LLMs) open up new research directions in other areas of AI, including the broad area of decision making. In this position paper we identify a novel research agenda for the field of computational social choice by charting the significant range of opportunities for integrating text-processing algorithms into the design and analysis of mechanisms for collective decision making.

## 1 Introduction

A collective decision involves a group of autonomous agents having to choose from a set of alternatives, with this choice affecting each individual decision-maker. From voting in political elections and deciding on municipal budgets, to collectively agreeing on binding texts of a contract or public statement, the applications of collective decision making are varied and complex. Recent developments in digital democracy—such as deliberation platforms and systems for

---

\*Preprint.

collective statement generation (Jigsaw, 2025)—introduce new challenges in this field of study due to the unstructured nature of the space of possible outcomes and the potentially large-scale participation enabled by modern technology. We argue that these challenges can be addressed through a careful integration of algorithms developed in computational social choice (COMSOC)—grounded in economics, discrete mathematics, and theoretical computer science—and natural language processing (NLP) techniques, including large language models (LLMs), which in turn rely on statistical analysis.

Collective decision making in multiagent systems, and in AI more broadly, has traditionally been analysed through the lens of social choice theory (Arrow et al., 2002). In essence, classical models consider a set of agents or voters who express preferences over a set of alternatives by means of various ballot types, such as approvals or rankings, and a voting rule or aggregation function tasked with selecting a collective decision from this input. COMSOC researchers have thoroughly studied computational problems in collective decision making, from computing outcomes with good properties to finding strategic ballots to influence results (Brandt et al., 2016).

Classical models of social choice typically treat the available alternatives as mathematical objects, denoted simply by letters, and assume that there are no *a priori* dependencies between alternatives—with the notable exception of voting in combinatorial domains (Lang and Xia, 2016)). This, of course, is an extreme abstraction from reality. City projects submitted to a participatory budgeting campaign are rich semantic objects containing, at the very least, a title and a description, and often also a location and a category (such as “education” or “environment”). Proposals on deliberation platforms typically are written in some recognisable style of language, often identifying the political affiliation of the proposer. In these domains, it is often easy for humans to recognise duplicate proposals, or to infer the opinion of other voters from just a few samples of their preferences on related issues, while these tasks are hard if not impossible when alternatives are represented by abstract mathematical symbols.

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. LLMs have redefined the modern understanding of the term “AI”, not only in the media but also for many researchers. LLMs are being used to annotate huge corpora, fundamentally changing social science research (Grossmann et al., 2023), and to provide real-time transcription and analysis for both TV and surveys (Chang et al., 2024). LLMs are also being misused to simulate participants in social science data (Wang et al., 2025b) and to replace software evaluators (Wang et al., 2025a). These extraordinary abilities of general text processing have resulted from a steady march of innovation, investment, and software tools being developed; from Word2Vec (Mikolov et al., 2013), to transformers (Vaswani et al., 2017), to easy-to-use software packages and distribution methods, and LLMs being accessible as API services.

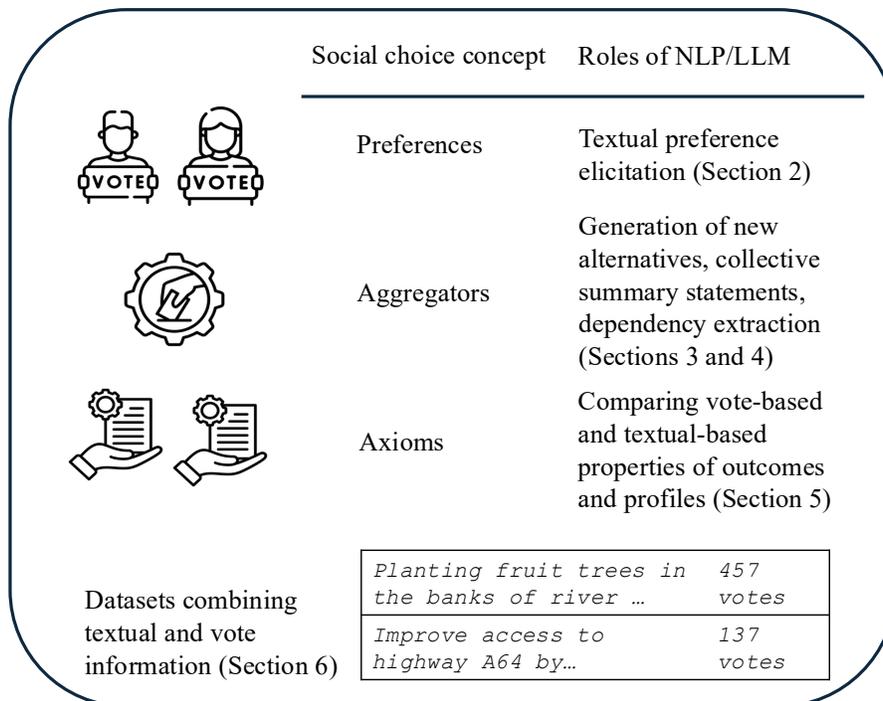


Figure 1: Schematic view of collective decision making as aggregation of preferences, elicited from voters, into outcomes, with potential contributions of NLP at each level.

Through these tools and APIs, LLMs are easily available to practitioners of digital democracy and policy makers, who are increasingly in search of safe and well-tested approaches to combine classical voting methods with the conversational capabilities of the latest LLM technology. In this position paper we argue that contemporary NLP techniques, including LLMs, provide researchers in computational social choice with powerful tools to explore the use of textual data in collective decision making. They further empower researchers to design innovative decision-making tools that are more likely to be adopted by large segments of the population and policy makers—with all the risks that this entails and which thus must be carefully assessed and managed.

The aim of this paper is twofold. First, by structuring and surveying the recent research efforts on developing algorithms for collective decisions with text, we are able to show that they do not form a set of isolated initiatives but an identifiable research agenda. Second, our paper is an invitation to researchers from fields such as computational social choice and algorithmic game theory to embrace textual data and view collective decisions in this new light. To present this research agenda, we use a schematic (and simplified) view of collective deci-

sion 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*) that constrain what would be an acceptable outcome in view of the preferences collected. Each of the next four sections details how NLP techniques, and LLMs in particular, can be implemented to advance research and development in this domain, as illustrated in Figure 1. Section 2 discusses the elicitation of voter preferences by means of text, and the use of NLP techniques to tackle the information overload resulting from high levels of participation in digital platforms. Sections 3 and 4 present innovative aggregation methods obtained by employing NLP techniques in the design of voting processes, with Section 4 focusing specifically on the problem of aggregating textual statements. Section 5 considers the application of classical social choice concepts, and notably the axiomatic method, on textual voting data. We conclude in Section 6 with a brief review of currently available datasets that collect user preferences over alternatives described or complemented by textual information, and in Section 7 with a critical assessment of the risks and limitations of our proposed research programme.

## 2 Textual Preference Elicitation

The elicitation of preferences is usually the first direct contact between a 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, Soumalias et al. (2025) and Huang et al. (2025) 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 4), most approaches train and assess a specific LLM for each user to be able to infer the user’s preferences over new statements (Fish et al., 2024; Tessler et al., 2024). Mohsin et al. (2021) 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 (Peri et al., 2021) or to generate dependencies between alternatives using text descriptions (Vallam et al., 2019). Finally, reconstructing voter preferences from demographic data has been tested in participatory budgeting (Yang et al., 2024) and pairwise voting (Gudiño-Rosero et al., 2024).

A number of research directions can be further explored in the elicitation stage. First, LLM agents can help users identify relevant alternatives. In partic-

ipatory 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 neighbourhood, 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. Third, a suitably designed LLM interface can support the elicitation of preferences, for example, by asking comparative questions, reducing the cognitive load associated with casting a ballot; we are beginning to see work in this direction (Montazerlghaem et al., 2025; Soumalias et al., 2025). 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 (Kephart et al., 2019).

### 3 Text-Aware Aggregation

After we have elicited voters’ preferences, we need to aggregate these preferences. Unfortunately, the information provided by the collected ballots might not suffice to find a satisfactory outcome, perhaps because preference information is (very) incomplete or because logical dependencies between alternatives are ignored. If additional textual data on the alternatives is available, NLP techniques can be leveraged to overcome these problems, e.g., by finding similar proposal or filling in preferences (Vallam et al., 2019). A paradigmatic application would be a system for online discussions where users can post comments and up- or down-vote those of others. Indeed, this might be one of the most common forms of social choice happening today, both on deliberation and on social media platforms.

In these cases, it is possible to make decisions purely based on the votes, for example by interpreting up-votes as approvals and selecting the most popular or representative comments using approval-based, multi-winner voting rules (Lackner and Skowron, 2023), as demonstrated by the website PROPORTIONALITY PRESS (Rey, 2025). However, in many applications, the vote data is highly incomplete, as the number of comments is often too large for any user to view even a significant percentage of the contributions.<sup>1</sup> More fundamentally, by ignoring the text, we might select comments that are contradictory or involuntarily ignore important sub-discussions and topics.

The textual information present in these datasets in the form of posted comments can be leveraged to extract structural information about the alternatives. For example, we can use NLP techniques to extract logical connections between

---

<sup>1</sup>This can be checked in conversations extracted from Pol.is (Computational Democracy Project, 2024), particularly the larger ones. For example, in the `march-on.operation-marchin-orders` conversation (2,161 comments), the average user reacted to 85 comments, while in the `austria-climate.5tzfrp5eaa.2022-07-07` conversation (2,138 comments), the average user reacted to 95. Other Pol.is conversations exhibit similarly sparse votes: on average, across all available data, users react to about 9% of the comments.

comments, such as whether one argument attacks another—a problem extensively studied by the argumentation mining community (Lawrence and Reed, 2019; Gorur et al., 2025). Based on such logical structures, we can provide logically consistent recommendations using methods from judgment aggregation (Baumeister et al., 2024; Grossi and Pigozzi, 2014) or solve the problem of incomplete vote information by expanding voters’ partially expressed preferences into maximally consistent viewpoints (Bernreiter et al., 2024). Similarly, we can use LLMs to provide topical clusterings of the comments (Liu et al., 2025), which we can then use to constrain outcomes to ensure full coverage of the topics discussed (Masariik et al., 2024). This approach ensures 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.

To succeed, many challenges need to be faced. While voting with constraints on outcomes has been studied intensively, the literature often assumes rather precise and complex constraints expressed in, e.g., propositional logic. Such constraints are much harder to extract from text than simple relations such as argument attacks. Further work is therefore necessary to identify the best trade-offs in terms of the expressivity of the extracted structures. Additionally, there are complex questions related to how to represent and use conflicting or contentious information, and these issues are gaining traction in the broader NLP community (Kim and Diaz, 2025; Wang et al., 2025c). Social choice research may offer new approaches to address these problems.

## 4 Large Language Models as Aggregators

Consider a question under discussion by a community, such as “Should LLM reviews be integrated into the paper acceptance decisions at an international AI conference?”. Ask each community member—in this example the members of the conference program committee—to write a short paragraph summarising their opinion, then feed all this to an LLM with a carefully designed prompt to generate one or more candidates to be used as consensual statements representing the opinion of the community. 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 make them particularly fitting for heavily polarised crowds, contributing to efforts to develop technology that allows deliberation to scale. The process of generating such statements can be analysed as a social choice problem: if we can aggregate rankings into a collective view with a voting rule, why not do so directly with textual preferences? Note that here we are not interested in knowing how the preferences of an individual can be correctly and verifiably elicited using text (as discussed in Section 2), but how to aggregate statements so that each voter feels represented.

To give a (partial) overview of existing work in this direction, we can start from the work of Fish et al. (2024), who initiated the study of LLMs in a social

choice context. Starting from a set of statements proposed by users, they tested the use of LLMs to generate new representative statements to be submitted to a vote amongst users, complementing their empirical findings with provable mathematical bounds. Their empirical framework borrows ideas from earlier work by Bakker et al. (2022), which developed into a large-scale experiment testing the LLM effect in so-called “caucus deliberations” (Tessler et al., 2024). In their experimental setting, a group of five participants submit textual opinions to a central LLM mediator, and then vote in a two-stage process on the set of resulting LLM-generated consensus statements. They show that this process improves the acceptance of consensual statements relative to human mediators. Further, Konya et al. (2025) report on how LLM-generated consensual statements can be useful in finding common ground between two populations in conflict. More recently, LLM-generated summaries have been tested in a formal coalition formation model (Briman et al., 2025), discussed as a potential improvement for the `Pol.is` platform (Small et al., 2023), and implemented in a civic planning context (Jigsaw, 2025).

To identify research challenges for social choice theorists, we can simplify this process to the bone: by fixing an LLM and a textual prompt that instructs the LLM to generate  $k$  consensual statements from an input set of voter statements, we obtain the equivalent of a multi-winner voting rule that works on textual input (with some variability potentially induced by the LLM temperature). This opens up the way to formal studies of their properties, in line with the analysis of Fish et al. (2024) on the proportionality of query protocols. For instance, in line with the axiom of *anonymity*, one can check the sensitivity to the order of presentation of the opinions. Or, building on *strategyproofness*, the robustness of LLM-aggregators to prompt-attacks by participants can be assessed (Gudiño-Rosero et al., 2025). While most existing work focuses on including LLMs in a more complex democratic process, here the study could be purely on getting the aggregation right without the need for multiple layers of LLMs.

## 5 Axiomatic and Textual Analysis

Central to classical social choice theory is the axiomatic method (Arrow et al., 2002), 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 (Nelissen, 2023; Peters and Skowron, 2025). 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* (Tideman, 1987; Elkind et al., 2011). 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.<sup>2</sup>

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 (Boixel et al., 2022), 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 (Conitzer et al., 2024), 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).

---

<sup>2</sup>A famous example illustrating the relevance of this idea is the public vote to choose a name for the new city arising from the amalgamation of Port Arthur and Fort William in Ontario, Canada, in 1969. Citizens were asked to choose between three names: *Lakehead*, *The Lakehead*, and *Thunder Bay*. *Thunder Bay* won by a small margin (City of Thunder Bay Archives, 2018).

## 6 Datasets Combining Text and Votes

The public library of preference data `Preflib.org` (Mattei and Walsh, 2013) enabled major advances in the field of computational social choice by providing standardised benchmark datasets of voter preferences, typically stored as rankings over alternatives. It was recently complemented by `Pabulib.org` (Faliszewski et al., 2023), 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 (Szufa et al., 2025). 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 `Po1.is` (Small et al., 2021; Computational Democracy Project, 2024) and `Decidim` (Serramia et al., 2019; `Decidim.Barcelona`, 2024) enable discussion of divisive topics, comment posting, and voting on comments (approve, disapprove, neutral, abstain). Large datasets from Reddit, Twitter, and other online fora could be leveraged for this purpose (Leskovec and Sosič, 2016; Kelly et al., 2023).
- Collective selection of consensual statements (Fish et al., 2024; Tessler et al., 2024; Fish et al., 2025; DeepMind Technologies Limited, 2024): 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 (Navarrete et al., 2024, 2023). 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 (Zambrano et al., 2025).

## 7 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 programme we propose necessitates substantial critical reflection. Given the high-stakes nature of many

COMSOC-related 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 (Council of European Union, 2024), 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 (Fish et al., 2024; Huang et al., 2025), these rely on the accuracy of NLP and LLM-based techniques for preference extraction and representation, which remain difficult to ensure (Schroeder et al., 2025).

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 (Jadhav et al., 2025). Moreover, even in the simplest settings, LLMs will unavoidably make mistakes when extracting preferences, or other relevant structures, from textual input (Tan et al., 2024).

In adopting a responsible approach (Kenthapadi et al., 2023), 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 substantially 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 and to enrich the design space available to research in social choice. Overcoming the aforementioned challenges will enable collective decision making research to expand its role in society (Abebe et al., 2020), reach more communities (Birhane et al., 2022), and allow this research to have broad impact for social good (Hager et al., 2019).

## Acknowledgments

This work reports and expands the discussion held in a working group at the Dagstuhl Seminar 25401 on “Societal Impact of Computational Social Choice” (<https://www.dagstuhl.de/25401>) held in October 2025.

Nicolas Maudet acknowledges the support of the ANR-23-IACL-0007 grant (AI Cluster PostGenIA) under the France 2030 program. Umberto Grandi, Ulle Endriss, and Théo Delemazure acknowledge the support of ERC Synergy Grant ADDI, ID: 101166894 (<http://doi.org/10.3030/101166894>)



*Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.*

## References

- Rediet Abebe, Solon Barocas, Jon Kleinberg, Karen Levy, Manish Raghavan, and David G Robinson. 2020. Roles for Computing in Social Change. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAccT)*.
- Kenneth J. Arrow, Amartya K. Sen, and Kotaro Suzumura (Eds.). 2002. *Handbook of Social Choice and Welfare*. North-Holland.
- Michiel Bakker, Martin Chadwick, Hannah Sheahan, Michael Tessler, Lucy Campbell-Gillingham, Jan Balaguer, Nat McAleese, Amelia Glaese, John Aslanides, Matt Botvinick, and Christopher Summerfield. 2022. Fine-Tuning Language Models to Find Agreement among Humans with Diverse Preferences. In *Proceedings of the 36th International Conference on Neural Information Processing Systems (NeurIPS)*.
- Dorothea Baumeister, Gábor Erdélyi, Ronald de Haan, and Jörg Rothe. 2024. Judgment Aggregation. In *Economics and Computation: An Introduction to Algorithmic Game Theory, Computational Social Choice, and Fair Division*, Jörg Rothe (Ed.). Springer, 467–504.
- 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)*.
- Abeba Birhane, William Isaac, Vinodkumar Prabhakaran, Mark Diaz, Madeleine Clare Elish, Iason Gabriel, and Shakir Mohamed. 2022. Power to the People? Opportunities and Challenges for Participatory AI. In *Proceedings of the 2nd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*.
- 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)*.

- Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (Eds.). 2016. *Handbook of Computational Social Choice*. Cambridge University Press.
- 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)*.
- Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, Wei Ye, Yue Zhang, Yi Chang, Philip S. Yu, Qiang Yang, and Xing Xie. 2024. A Survey on Evaluation of Large Language Models. *ACM Transactions on Intelligent Systems and Technology* 15, 3 (2024), 1–45.
- City of Thunder Bay Archives. 2018. Amalgamation. Available at <https://www.thunderbay.ca/en/city-hall/amalgamation.aspx>. Last accessed on 2025-11-20.
- Computational Democracy Project. 2024. Open Polis Data. Available at <https://github.com/compdemocracy/openData>. Last accessed on 2025-11-20.
- 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)*.
- Council of European Union. 2024. Regulation (EU) 2024/1689.
- Decidim.Barcelona. 2024. MetaDecidim Data. Available at <https://github.com/elaragon/metadecidim/tree/master>. Last accessed on 2025-11-26.
- DeepMind Technologies Limited. 2024. Habermas Machine Dataset. Available at [https://github.com/google-deepmind/habermas\\_machine](https://github.com/google-deepmind/habermas_machine). Last accessed on 2025-11-20.
- Edith Elkind, Piotr Faliszewski, and Arkadii Slinko. 2011. Cloning in Elections: Finding the Possible Winners. *Journal of Artificial Intelligence Research* 42 (2011), 529–573.
- Piotr Faliszewski, Jaroslaw Flis, Dominik Peters, Grzegorz Pierczynski, Piotr Skowron, Dariusz Stolicki, Stanislaw Szufa, and Nimrod Talmon. 2023. Participatory Budgeting: Data, Tools and Analysis. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*.

- Sara Fish, Paul Gözl, David C. Parkes, Ariel D. Procaccia, Gili Rusak, Itai Shapira, and Manuel Wüthrich. 2024. Generative Social Choice. In *Proceedings of the 25th ACM Conference on Economics and Computation (EC)*.
- Sara Fish, Paul Gözl, David C. Parkes, Ariel D. Procaccia, Gili Rusak, Itai Shapira, and Manuel Wüthrich. 2025. Generative Social Choice Open Datasets. Available at [https://github.com/generative-social-choice/survey\\_data/](https://github.com/generative-social-choice/survey_data/). Last accessed on 2025-11-20.
- 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.
- Davide Grossi and Gabriella Pigozzi. 2014. *Judgment Aggregation: A Primer*. Morgan & Claypool.
- Igor Grossmann, Matthew Feinberg, Dawn C. Parker, Nicholas A. Christakis, Philip E. Tetlock, and William A. Cunningham. 2023. AI and the Transformation of Social Science Research. *Science* 380, 6650 (2023), 1108–1109.
- 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.
- Jairo Gudiño-Rosero, Clément Contet, Umberto Grandi, and César Hidalgo. 2025. Prompt Injection Vulnerability of Consensus Generating Applications in Digital Democracy. *arXiv preprint arXiv:2508.04281* (2025).
- Gregory D. Hager, Ann Drobnis, Fei Fang, Rayid Ghani, Amy Greenwald, Terah Lyons, David C. Parkes, Jason Schultz, Suchi Saria, Stephen F. Smith, and Milind Tambe. 2019. Artificial Intelligence for Social Good. *arXiv preprint arXiv:1901.05406* (2019).
- David Huang, Francisco Marmolejo-Cossío, Edwin Lock, and David C. Parkes. 2025. Accelerated Preference Elicitation with LLM-Based Proxies. *arXiv preprint arXiv:2501.14625* (2025).
- 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)*.
- Jigsaw. 2025. How one of the Fastest-Growing Cities in Kentucky used AI to Plan for the next 25 Years. Available at <https://jigsaw.google/>. Last accessed on 2025-12-08.
- Markelle Kelly, Rachel Longjohn, and Kolby Nottingham. 2023. The UCI Machine Learning Repository. Available at <https://archive.ics.uci.edu>. Last accessed on 2025-12-08.

- Krishnaram Kenthapadi, Himabindu Lakkaraju, and Nazneen Rajani. 2023. Generative AI meets Responsible AI: Practical Challenges and Opportunities. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*.
- Jeffrey O. Kephart, Nicholas Mattei, and Francesca Rossi. 2019. Collective Decision Making by Consensus in Cognitive Environments. US Patent Application 15/845,509.
- 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)*.
- 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)*.
- Martin Lackner and Piotr Skowron. 2023. *Multiwinner Voting with Approval Preferences*. Springer.
- Jérôme Lang and Lirong Xia. 2016. Voting in Combinatorial Domains. In *Handbook of Computational Social Choice*, F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia (Eds.). Cambridge University Press, Chapter 9.
- John Lawrence and Chris Reed. 2019. Argument Mining: A Survey. *Computational Linguistics* 45, 4 (2019), 765–818.
- Jure Leskovec and Rok Sosič. 2016. SNAP: A General-Purpose Network Analysis and Graph-Mining Library. *ACM Transactions on Intelligent Systems and Technology (TIST)* 8, 1–20 (2016), 1.
- 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)*.
- Tomás Masarík, Grzegorz Pierczynski, and Piotr Skowron. 2024. A Generalised Theory of Proportionality in Collective Decision Making. In *Proceedings of the 25th ACM Conference on Economics and Computation (EC)*.
- Nicholas Mattei and Toby Walsh. 2013. PrefLib: A Library of Preference Data ([preflib.org](http://preflib.org)). In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT)*. Springer.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient Estimation of Word Representations in Vector Space. *arXiv preprint arXiv:1301.3781* (2013).

- 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)*.
- 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).
- Carlos Navarrete, Mariana Macedo, Rachael Colley, Jingling Zhang, Nicole Ferrada, Maria Mello, Rodrigo Lira, Carmelo Bastos-Filho, Umberto Grandi, Jérôme Lang, and César A. Hidalgo. 2023. Data Collection of two Real-World Digital Democracy Systems. Available at <https://doi.org/10.7910/DVN/8E0EA4>. Last accessed on 2025-11-30.
- 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.
- Pelle Nelissen. 2023. An Empirical Analysis of Participatory Budgeting in Amsterdam. *arXiv preprint arXiv:2310.18033* (2023).
- 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)*.
- Domink 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.
- Simon Rey. 2025. Proportionality Press. Available at <https://simon-rey.github.io/ProportionalityPress/>. Last accessed on 2025-12-03.
- Hope Schroeder, Marianne Aubin Le Quéré, Casey Randazzo, David Mimno, and Sarita Schoenebeck. 2025. Large Language Models in Qualitative Research: Uses, Tensions, and Intentions. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*.
- Marc Serramia, Jordi Ganzer, Maite López-Sánchez, Juan A Rodríguez-Aguilar, Natalia Criado, Simon Parsons, Patricio Escobar, and Marc Fernández. 2019. Citizen Support Aggregation Methods for Participatory Platforms. In *Artificial Intelligence Research and Development*. IOS Press, 9–18.
- Christopher T. Small, Michael Bjorkegren, Timo Erkkilä, Lynette Shaw, and Colin Megill. 2021. Polis: Scaling Deliberation by Mapping High Dimensional Opinion Spaces. *Recerca: Revista de Pensament i Anàlisi* 26, 2 (2021).

- 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).
- Ermis Soumalias, Yanchen Jiang, Kehang Zhu, Michael Curry, Sven Seuken, and David C. Parkes. 2025. LLM-Powered Preference Elicitation in Combinatorial Assignment. *arXiv preprint arXiv:2502.10308* (2025).
- 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.
- 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)*.
- 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).
- T. Nicolaus Tideman. 1987. Independence of Clones as a Criterion for Voting Rules. *Social Choice and Welfare* 4, 3 (1987), 185–206.
- 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 (AA-MAS)*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention Is All You Need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS)*.
- Angelina Wang, Jamie Morgenstern, and John P. Dickerson. 2025b. Large Language Models that Replace Human Participants Can Harmfully Misportray and Flatten Identity Groups. *Nature Machine Intelligence* 7 (2025), 400–411.
- Han Wang, Archiki Prasad, Elias Stengel-Eskin, and Mohit Bansal. 2025c. Retrieval-Augmented Generation with Conflicting Evidence. *arXiv preprint arXiv:2504.13079* (2025).

- Ruiqi Wang, Jiyu Guo, Cuiyun Gao, Guodong Fan, Chun Yong Chong, and Xin Xia. 2025a. Can LLMs Replace Human Evaluators? An Empirical Study of LLM-as-a-judge in Software Engineering. *Proceedings of the ACM on Software Engineering* 2, ISSTA (2025), 1955–1977.
- 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)*.
- Juan Zambrano, Clément Contet, Jairo Gudiño-Rosero, Felipe Garrido-Lucero, Umberto Grandi, and César A. Hidalgo. 2025. Leveraging LLMs for Privacy-Aware Predictions in Participatory Budgeting. *arXiv preprint arXiv:2508.06577* (2025).