

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

Umberto Grandi, Robert Bredereck, Théo Delemazure, Ulle Endriss, Jan Maly,
Nicholas Mattei, Nicolas Maudet, Oliviero Nardi, Stanisław Szufa

The Vision

Artificial Intelligence and Large Language Models (LLMs) are rapidly impacting democratic practices, from national elections to massive online deliberations.

By integrating classical social choice theory with advanced natural language processing (NLP), we can scale the benefits of small-group deliberation to society at large while maintaining principled and explainable outcomes.


How LLMs Change Collective Decision Making

Textual Preference Elicitation: LLMs can act as an interactive interface to help voters discover relevant alternatives, understand the implications of their choices, and reduce cognitive load by asking comparative questions.

Text-Aware Aggregation: In large-scale debates, preference data is often sparse and highly incomplete. NLP can fill these gaps by extracting logical connections between arguments and grouping similar topics to ensure diverse viewpoints are represented.

Consensus Generation: Instead of just aggregating votes, LLMs can directly merge individual textual opinions into unified, neutral consensus statements that can bridge the gap in heavily polarized crowds.

Classical Social Choice Theory




Impact of NLP/LLM

Textual preference elicitation, chatbots for preference formation and updates

Generation of new alternatives, collective summary statements, dependency extraction

Datasets combining votes with textual data on alternatives (Pol.is, generative social choice, collective statements...)



Read the Paper Here!
hal.science/hal-05548117

Social Choice and Text

Rethinking Axioms: Traditional social choice axioms treat choices as abstract symbols, but textual analysis allows us to refine these rules—such as identifying duplicate "clone" proposals based on semantic similarity.

Accessibility for Everyone: NLP could translate complex mathematical voting rules into natural language, empowering everyday citizens to formulate and audit their own fairness requirements.

High-Stakes Oversight: Applying these systems to real-world democratic tasks often places them in high-risk categories under frameworks like the European AI Act, demanding strict regulatory oversight.

Mitigating Biases: Practitioners must actively guard against LLM hallucinations, baked-in model biases, and the over-representation of English data, all of which can severely compromise fair participatory processes.