Incomplete Opinions in Collective Decision Making

Doctoral Consortium

Zoi Terzopoulou
Institute for Logic, Language and Computation
Amsterdam, The Netherlands
z.terzopoulou@uva.nl

ABSTRACT

I study settings of collective decision making where the members of a group may report intrinsically incomplete opinions. In such contexts, we need to design aggregation mechanisms that satisfy normatively desirable properties, are effective in discovering a ground truth, incentivise the agents to be truthful, or several of the above.

KEYWORDS

Social Choice; Incomplete Preferences; Incomplete Judgments

ACM Reference Format:
Zoi Terzopoulou. 2020. Incomplete Opinions in Collective Decision Making. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), Auckland, New Zealand, May 9–13, 2020,* IFAAMAS, 3 pages.

1 INTRODUCTION

Incomplete opinions arise naturally in many scenarios of collective decision making. Agents (whether human or artificial), when having to make decisions together, may not be able to form their own individual opinions about *all* issues in question.[1]

Example 1. On travel websites, users are asked to rank different hotels according to their experience and preferences. But most users will not have an opinion about all 28 million accommodation listings that `booking.com`, for instance, offers. Moreover, some hotels may be comparable (e.g., because they are located within close distance), while some others may not. If a, b, c, d, e denote different hotels, an individual preference might thus take the following form:

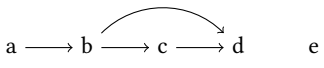

where an arrow from a to b means that a is preferred to b. Note that in general, the preferences reported by the agents may be non-transitive as well as incomplete—however, we assume that cyclic preferences are not admissible. △

We thus need aggregation mechanisms that allow for incomplete inputs and produce reasonable collective outcomes. Of course, what "reasonable" means depends on the context. We distinguish three such contexts, where we (from the viewpoint of the mechanism designer) care about one (or a combination) of the following: (*i*) reaching a compromise between the diverse opinions of the group members for issues to be decided upon that are subjective in nature (like in Example 1); (*ii*) discovering the ground truth with respect to a number of objective issues (like in crowdsourcing experiments); (*iii*) avoiding the strategic behaviour of agents that may report untruthful opinions in order to obtain a preferable outcome for themselves (like in election problems) [1].

The literature to date in computational social choice has been handling incomplete opinions mainly by applying the popular solution concept of *possible winners* [4]: all complete extensions of an incomplete preference profile are considered and the possible winners are those alternatives winning under *some* of those extensions. This approach is indeed sensible when the opinions of the agents are actually complete, but our information about them (or the information of the agents about the opinions of their peers) is incomplete [11, 13]. Instead, I am interested in *intrinsically* incomplete opinions. Some previous work has also explored the aggregation of such opinions, either in terms of incomplete preferences in the formal framework of *voting* [3, 7], or in terms of incomplete judgments in *judgment aggregation* [2].

Existing relevant literature largely focuses on impossibility results, that is, results showing that—also under the more general setting that allows for truly incomplete opinions—there do not exist aggregation rules that satisfy a collection of desirable properties simultaneously; this is for instance the topic addressed in the work of Pini et al. [7] and of Dietrich and List [2].

However, what is missing from the literature in the area is a detailed account of where the incompleteness of the agents' opinions comes from, what it is caused by—these assumptions play a significant role in the way our formal models are constructed, and consequently, in the results we obtain. By aiming attention at more specific aggregation contexts, we can see that most of the classical impossibility results are not so alarming as they may have originally seemed. In Sections 2, 3, and 4 I describe how my thesis project contributes to this discussion.

2 REACHING A COMPROMISE

How can we determine whether a sufficiently good compromise is reachable for a group of agents with diverting opinions? One way, typically employed by the social choice community, suggests the use of *axioms*, i.e., properties of aggregation rules that are normatively appealing [14]. Such axioms are well-studied in the

[1] Indeed, most of the real preference data collected in the website `preflib.org` is in fact incomplete [6].

Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), B. An, N. Yorke-Smith, A. El Fallah Seghrouchni, G. Sukthankar (eds.), May 9–13, 2020, Auckland, New Zealand.

standard frameworks where completeness is supposed, but present numerous open questions when incompleteness sets in.

Recall Example 1: the agents can be expected to hold incomplete preferences, which may also display different degrees of incompleteness (for instance, a frequent traveller has probably visited and can compare more hotels than the average person). It is then natural to focus on a family of aggregation rules under which the *weight* assigned to an agent's ranking of two alternatives depends on how many other pairs of alternatives she ranks as well.[2] We characterise axiomatically rules in this family by two axioms [9]:[3]

- *Majoritarianism:* Respect the will of the majority on the pairwise ranking of alternatives, when possible.
- *Splitting:* If several agents have mutually compatible preferences, it should be possible for them to form a pact and all report the union of their individual preference sets—without this change affecting the outcome.

We find that majoritarianism characterises the rule with constant weights and that the splitting axiom (in certain forms) characterises the rule where an agent's weight is inversely proportional to the size of her reported preference set.

But weight rules are not the only interesting ones in such a setting. Indeed, by selecting a different set of axioms, viz., *anonymity*, *neutrality*, *reinforcement*, and *continuity* (counterparts to the classical axioms for the complete case), we obtain the distinct family of *positional scoring rules* [5], which assign scores to alternatives with respect to their position in a preference ranking.

Various other axioms can be considered, leading to different notions of compromise. Depending on which we find appealing, we can argue for the use of the corresponding rules that satisfy them. This idea is obviously applicable in frameworks of collective decision making beyond voting, such as judgment aggregation.[4]

3 DISCOVERING A GROUND TRUTH

Collective decision making often takes place for the purposes of discovering the true state of a given situation. For example, doctors form committees to identify the exact illness of a patient, and reviewers are collectively asked to express their judgments on the soundness of the proofs in a submitted paper. It is well-known that the accuracy of a group exceeds the accuracy of a single agent, provided that all members of the group independently make better than random judgments. But the (in)completeness of individual opinions constitutes a vital assumption in such contexts as well.

Suppose, more specifically, that you need to determine the correct answer to a question (e.g., whether a colleague should get promoted) that depends on two independent premises (e.g., whether the colleague is excellent at her work and whether she is a good team player). You can ask several agents to each evaluate either just one of those premises (which they can do with relatively high accuracy) or both (in which case their need to multitask will lower their accuracy). When these accuracies are known, we find that the optimal rule to aggregate the judgments reported by the agents is a weight rule as presented in Section 2; we also determine how many agents it is optimal to ask for how many judgments [10].

This work introduces the concept of multitasking in formal models of collective decision making and exemplifies its significance within a relatively simple scenario in judgment aggregation—this approach could be further developed in related settings, like in the classical framework of voting.

4 AVOIDING STRATEGIC MANIPULATION

The goal of collective decision making is to find a compromise between the individuals involved, or to discover a ground truth. But this may not be satisfactory for the members of a group that simply want a better outcome for themselves. In particular, agents that have the freedom to report incomplete opinions may lie in three ways (only the last of which is possible in the complete case):

- by *hiding* their truthful opinion on some issues;
- by *inventing* a new untruthful opinion on some issues;
- by *reversing* their truthful opinion on some issues.

Are there positional scoring rules for incomplete preferences that are immune to these types of manipulation? We answer positively for manipulation moves that are restricted to just a single one of the above types, as well as for the combination of the two last types, but negatively for every other possible manipulation [5].

When subjective opinions (preferences or judgments) are concerned, it is not difficult to describe an agent's desired outcome in relation to her truthful opinion. But if a ground truth is involved, various sources of strategic behaviour need to be investigated, depending on whether the agents are motivated by (i) the group tracking the truth, by (ii) maximising their own reputation, or by (iii) maximising the agreement of the group's findings with their own personal opinion. Our work analyses how these parameters affect the manipulability of the optimal truth-tracking rule [10].

5 CONCLUSION AND FUTURE WORK

Many questions remain to be answered regarding incomplete opinions in collective decision making. Some of those concern the complexity of the aggregation rules developed so far; some involve the definition of new rules, generalising known ones from the standard frameworks with completeness but tailoring them to our specific contexts of interest. In addition, after having established our theoretical background, it would be worth reviewing popular applications where incompleteness of the agents' opinions is a central feature (e.g., recommender systems, crowdsourcing sites), and make use of experimental evidence to test and improve our models. This is what I plan to work on for the remainder of my thesis.

[2] Weight rules reduce to the known Kemeny rule in the complete framework.

[3] We have also examined similar axioms and obtained analogous characterisation results in the framework of judgment aggregation [12].

[4] In recent work, I analyse the family of quota rules with incomplete judgments, also from an axiomatic perspective [8].

REFERENCES

[1] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (Eds.). 2016. *Handbook of Computational Social Choice.* CUP.

[2] Franz Dietrich and Christian List. 2008. Judgment aggregation without full rationality. *Social Choice and Welfare* 31, 1 (2008), 15–39.

[3] Ulle Endriss, Maria Silvia Pini, Francesca Rossi, and K. Brent Venable. 2009. Preference Aggregation over Restricted Ballot Languages: Sincerity and Strategy-Proofness. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI).* 122–127.

[4] Kathrin Konczak and Jérôme Lang. 2005. Voting procedures with incomplete preferences. In *Proceedings of the IJCAI-05 Multidisciplinary Workshop on Advances in Preference Handling.*

[5] Justin Kruger and Zoi Terzopoulou. 2020. Strategic Manipulation with Incomplete Preferences: Possibilities and Impossibilities for Positional Scoring Rules. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.
[6] Nicholas Mattei and Toby Walsh. 2013. Preflib: A library of preference data.. In *Proc. of the 3rd International Conference on Algorithmic Decision Theory (ADT)*.
[7] Maria Silvia Pini, Francesca Rossi, Kristen Brent Venable, and Toby Walsh. 2006. Strategic voting when aggregating partially ordered preferences. In *Proceedings of the 5th International Joint Conference on Autonomous Agents and Multiagent Systems (AAAMAS)*. 685–687.
[8] Zoi Terzopoulou. ND. Quota Rules for Incomplete Judgments. (ND). Unpublished.
[9] Zoi Terzopoulou and Ulle Endriss. 2019. Aggregating Incomplete Pairwise Preferences by Weight. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*.
[10] Zoi Terzopoulou and Ulle Endriss. 2019. Optimal Truth-Tracking Rules for the Aggregation of Incomplete Judgments. In *Proceedings of the 12th International Symposium on Algorithmic Game Theory (SAGT)*.
[11] Zoi Terzopoulou and Ulle Endriss. 2019. Strategyproof Judgment Aggregation under Partial Information. *Social Choice and Welfare* 53, 3 (2019), 415–442.
[12] Zoi Terzopoulou, Ulle Endriss, and Ronald de Haan. 2018. Aggregating Incomplete Judgments: Axiomatisations for Scoring Rules. In *Proceedings of the 7th International Workshop on Computational Social Choice (COMSOC)*.
[13] Zoi Terzopoulou, Alexander Karpov, and Svetlana Obraztsova. 2020. Restricted Domains of Dichotomous Preferences with Possibly Incomplete Information (Extended Abstract). In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.
[14] William Thomson. 2001. On the axiomatic method and its recent applications to game theory and resource allocation. *Social Choice and Welfare* 18, 2 (2001), 327–386.