



HAL
open science

Simultaneous Elicitation of Committee and Voters' Preferences

Beatrice Napolitano, Olivier Cailloux, Paolo Viappiani

► **To cite this version:**

Beatrice Napolitano, Olivier Cailloux, Paolo Viappiani. Simultaneous Elicitation of Committee and Voters' Preferences. Rencontres des Jeunes Chercheurs en Intelligence Artificielle 2019, Jul 2019, Toulouse, France. pp.59-62. hal-02160455

HAL Id: hal-02160455

<https://hal.science/hal-02160455>

Submitted on 20 Jun 2019

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.

Simultaneous Elicitation of Committee and Voters' Preferences

B. Napolitano¹

O. Cailloux¹

P. Viappiani²

¹ LAMSADE, UMR 7243, CNRS and Université Paris Dauphine, PSL Research University, Paris, France

² LIP6, UMR 7606, CNRS and Sorbonne Université, Paris, France

beatrice.napolitano@dauphine.fr, olivier.cailloux@dauphine.fr, paolo.viappiani@lip6.fr

Abstract

Social choice deals with the problem of determining a consensus choice from the preferences of different voters. In the classical setting, the voting rule is fixed beforehand and full information concerning the preferences of the voters is provided. Recently, the assumption of full preference information has been questioned by a number of researchers and several methods for eliciting preferences have been proposed. In this paper we go one step further and we assume that both the voting rule and the voters' preferences are partially specified. In this setting, we present an interactive elicitation protocol based on minimax regret and develop several query strategies that interleave questions to the chair and questions to the voters in order to attempt to acquire the most relevant information in order to quickly converge to optimal or a near-optimal alternative.

Keywords

Computational Social Choice, uncertainty in AI, preference elicitation, minimax regret.

1 Introduction

In a traditional social choice setting, both the social choice function and the full preference orderings of the voters are expressed beforehand. Nevertheless, in real situations this is not always the case. When considering decisions with large sets of alternatives, requiring voters to express full preference orderings can be prohibitively costly and, perhaps, not necessary in order to determine a consensus choice. Furthermore, it is often difficult for non-expert users to formalize a voting rule on the basis of some generic preferences over a desired aggregation method. Consider as an example a situation in which the council of a city wants to build a community allotment and the decision of what to grow has to be taken. The voters are the residents of the neighborhood and the alternatives consist of all the vegetables that can grow in that particular area. An external observer, who is helping with the voting procedure, is responsible for collecting the preferences of both the residents (regarding the vegetable to grow) and the members of the council (regarding the method for aggregating the residents' preferences). Imagine that, after collecting some

preferences, she notices that some alternatives are always the least preferred. In this case, it is really unlikely that these alternatives would be selected as winners by any reasonable voting rule. Thus, she may consider to ask the voters only their preferences over few strategical alternatives instead of asking them the full preference orderings. Furthermore, if she notices that a particular vegetable is the most preferred of almost the totality of the residents there is no need to know the specific voting rule. Any reasonable aggregation procedure, in fact, would select this alternative as a winner. The reasoning behind this is that it is not unrealistic to assume that the committee, that has to decide how to aggregate these preferences, is not able to define a specific procedure.

These observations have motivated a number of recent works considering social choice with partial preference orders [14, 9, 4] and incremental elicitation [3, 8] of voters' preferences. Lu and Boutilier [7] proposed the use of minimax regret to drive incremental vote and preferences elicitation when the social choice function is fixed and known. Furthermore, several authors [12, 6, 13] worked on positional scoring rules with uncertain weights, assuming that the preferences of the voters are fully known. Some elicitation methods for a quite general class of rules based on weak orders have also been proposed by Cailloux and Endriss [2].

In this paper we focus on positional scoring rules, that are a particularly common method used to aggregate rankings, and we assume that both the voters' preferences and the social choice rule are partially specified. We develop methods for computing the minimax-optimal alternative using positional scoring rules and we provide incremental elicitation methods to acquire relevant preference information. We then discuss several heuristics that determine queries, either to a voter or to the committee, that quickly reduce minimax regret. While previous works have considered either partial information about the voters' preferences or a partially specified aggregation method, we do not know of any work considering both sources of uncertainty at the same time.

2 Partial information

We consider a set A of m alternatives (products, restaurants, public projects, etc.) and a set $\{1, \dots, n\}$ of voters. Each voter j comes from an infinite set \mathbb{N} of potential voters and is associated to her “real” preference order \succ_j which is a linear order (connex, transitive, asymmetric relation) over the alternatives. A *profile* is the association of a preference to each voter and it is equivalently represented by $(\succ_1, \dots, \succ_n)$ or by $\mathbf{v} = (v_1, \dots, v_n) \in V$ where $v_j(i) \in \{1, \dots, m\}$ denotes the rank (position) of alternative i in the preference order \succ_j and V is the set of all possible preference profiles. A social choice function associates a profile with a set of winners; we consider *positional scoring rules (PSR)* which attach weights to positions according to the scoring vector (w_1, \dots, w_m) . Without loss of generality, we assume $w_1 = 1$ and $w_m = 0$. An alternative x obtains a score that depends on the rank obtained in each of the preference orders:

$$s^{\mathbf{v}, \mathbf{w}}(x) = \sum_{j=1}^n w_{v_j(x)} = \sum_{r=1}^m \alpha_r^x w_r$$

where α_r^x is the number of times that alternative x was ranked in the r -th position. The winners are the alternatives with highest scores.

In this work we assume fixed, but unknown, a profile $\mathbf{v}^* = (\succ_1^*, \dots, \succ_n^*)$ and a weight vector \mathbf{w}^* . Our knowledge at a given time of the preference of voter j is encoded by a partial order over the alternatives, thus a transitive and asymmetric binary relation, denoted by \succ_j^p ; we assume that preference information is truthful, i.e. $a \succ_j^p b \implies a \succ_j^* b$. An incomplete profile $\mathbf{p} = (\succ_1^p, \dots, \succ_n^p)$ maps each voter to a partial preference. A completion of \succ_j^p is any linear order \succ that extends \succ_j^p and we indicate with $C(\succ_j^p) = \{\succ \in \mathcal{L}(A) \mid \succ_j^p \subseteq \succ\}$ the set of possible completions of \succ_j^p , where $\mathcal{L}(A)$ is the set of linear orders on A .

Our knowledge regarding the weights of the positional scoring rule is represented by a set of constraints restraining the possible values that they can take. We consider a convex sequence of weights and we assume that the chair is able to specify additional preferences about how the social choice function should behave. Such requirements are encoded with linear constraints about the vector \mathbf{w} and the set of these constraints is denoted by \mathcal{C}_W . Given the set of convex weight vectors \mathcal{W} , we use $W \subseteq \mathcal{W}$ to denote the set of weight vectors compatible with the chair’s preferences. We will provide an example of such constraints, used to specify information about weights, in Section 4.

3 Robust winner determination

As a decision criterion to determine a winner, we propose to use minimax regret [11]. This concept has been used for robust optimization under data uncertainty [5] as well as in decision-making with uncertain utility values [10, 1].

Lu and Boutilier [7] have adopted this criterion for winner determination in social choice with the preferences of the voters that are only partially known, while the social choice function is predetermined and known. We use *maximum regret* to quantify the worst-case error, which measures, intuitively, how far an alternative is from the optimal one given current knowledge; the alternatives that minimize this error are selected as tied winners. The maximum regret of an alternative is the highest possible difference between its score and the one achievable by any alternative in any state compatible with our current knowledge. It is considered by assuming that an adversary can both 1) extend the partial profile \mathbf{p} into a complete one, and 2) instantiate the weights choosing among any weight vector in W . We formalize the notion of minimax regret in multiple steps. First of all, $\text{Regret}^{\mathbf{v}, \mathbf{w}}(x)$ is the “regret” of selecting x as a winner instead of choosing the optimal alternative under \mathbf{v} and \mathbf{w} :

$$\text{Regret}^{\mathbf{v}, \mathbf{w}}(x) = \max_{y \in A} s^{\mathbf{v}, \mathbf{w}}(y) - s^{\mathbf{v}, \mathbf{w}}(x).$$

The pairwise max regret $\text{PMR}^{\mathbf{p}, W}(x, y)$ of x relative to y given partial profile \mathbf{p} and the set of weights W is the worst-case loss of choosing x instead of y under all possible realizations of the full profile *and* weights. Max regret $\text{MR}^{\mathbf{p}, W}(x)$ is the worst-case loss of x and $\text{MMR}^{\mathbf{p}, W}$ is the value of minimax regret obtained when recommending a minimax optimal alternative.

$$\text{PMR}^{\mathbf{p}, W}(x, y) = \max_{\mathbf{w} \in W} \max_{\mathbf{v} \in C(\mathbf{p})} s^{\mathbf{v}, \mathbf{w}}(y) - s^{\mathbf{v}, \mathbf{w}}(x) \quad (1)$$

$$\text{MR}^{\mathbf{p}, W}(x) = \max_{y \in A} \text{PMR}^{\mathbf{p}, W}(x, y) \quad (2)$$

$$\text{MMR}^{\mathbf{p}, W} = \min_{x \in A} \text{MR}^{\mathbf{p}, W}(x) \quad (3)$$

$$x_{\mathbf{p}, W}^* = \operatorname{argmin}_{x \in A} \text{MR}^{\mathbf{p}, W}(x) \in A_{\mathbf{p}, W}^* \quad (4)$$

By picking as consensus choice an alternative associated with minimax regret $x_{\mathbf{p}, W}^* \in A_{\mathbf{p}, W}^*$, we can provide a recommendation that gives worst-case guarantees, giving some robustness in face of uncertainty. In cases of ties in minimax regret, we can either decide to return all minimax alternatives $A_{\mathbf{p}, W}^*$ as winners or to pick just one of them using a tie-breaking strategy. Because the constraints in \mathcal{C}_W are not necessarily linear PMR is not straightforwardly optimized, but it can be computed by adapting the reasoning of Lu and Boutilier [7] to the case of uncertain weights using linear programming. This optimization strategy uses the assumption of convexity.

4 Interactive elicitation

We propose an incremental elicitation method based on minimax regret. At each step, the system may ask a question either to one of the voters or to the chair. The goal is to acquire relevant information to reduce minimax regret as quickly as possible. As termination condition of elicitation, we can check whether minimax regret is lower than a

threshold or, if we wish optimality, we can perform elicitation until minimax regret drops to zero.

Question types We distinguish between questions asked to the voters and questions asked to the chair. For the first we consider comparison queries that ask a particular voter to compare two alternatives. Then, the partial profile \mathcal{p} is augmented on the basis of the voter’s answer. Thus, if voter j answers a comparison query stating that alternative a is preferred to b , then the partial order $\succ_j^{\mathcal{p}}$ is augmented with $a \succ_j^{\mathcal{p}} b$ and by transitive closure. The query for the chair, instead, aims at refining our knowledge about the scoring rule; in particular, we assume we can acquire constraints of the type:

$$w_r - w_{r+1} \geq \lambda(w_{r+1} - w_{r+2})$$

for $r \in \{1, \dots, m-2\}$, relating the difference between the importance of consecutive ranks.

Elicitation strategies A strategy is a function (or a random process) that, given our partial knowledge so far, gives us a question that should be asked.

The *Random* strategy first decides with a probability of $1/2$ each whether it will ask a question about weights or about a preference ordering. Then it equiprobably draws a question among the set of the possible ones. This strategy is clearly non efficient but it gives us a baseline for comparisons.

The *Extreme completions* strategy considers the solution of the minimax regret game, given the current knowledge, and estimates the contribution to the regret of our uncertainty about the weights and the profile. Then, it selects a question for the chair or for the voters depending on which of these two values is the highest.

The *Pessimistic* strategy selects among all the possible questions the one that leads to minimal regret in the worst case. Because all our questions are binary questions, for each of them the pessimistic strategy first assumes the answer is “No”, it updates the knowledge according to this negative answer and computes the minimax regret. Then, it does the same assuming the answer is positive. Finally, it uses some aggregation of these two MMR values (for example *max* or *avg*) as a score of the question. The question with the highest score is selected.

The *Two phase* strategy asks a predefined, non adaptive sequence of $m-2$ questions to the chair in order to gather informations about the weights w of the scoring rule, using one question per rank except the extreme ones that we are assuming known ($w_1 = 1, w_m = 0$). Then it only asks questions about the voters, using questions as defined in the *Pessimistic* strategy.

5 Conclusions

In this paper we have considered a social choice setting with partial information and we have proposed the use of minimax regret both as a means of robust winner determination as well as a guide to the process of simultaneous elicitation of preferences and voting rule.

The next step is to test our strategies in order to compare the performance of our interleaved elicitation approach that mixes questions to the chair and to the voters, to a more classical approach that elicits the rule first and then the voters’ preferences (or vice-versa).

Further development of elicitation strategies, considering alternative heuristics, is also an important direction for future work together with the extension of the approach to voting rules beyond scoring rules.

Acknowledgements: We thank the reviewers for comments helping to improve the paper.

References

- [1] C. Boutilier, R. Patrascu, P. Poupart, and D. Schuurmans. Constraint-based Optimization and Utility Elicitation using the Minimax Decision Criterion. *Artificial Intelligence*, 170(8–9):686–713, 2006.
- [2] O. Cailloux and U. Endriss. Eliciting a suitable voting rule via examples. In *ECAI 2014 - 21st European Conference on Artificial Intelligence, 18-22 August 2014, Prague, Czech Republic - Including Prestigious Applications of Intelligent Systems (PAIS 2014)*, pages 183–188, 2014.
- [3] M. Kalech, S. Kraus, G. A. Kaminka, and C. V. Goldman. Practical voting rules with partial information. *Autonomous Agents and Multi-Agent Systems*, 22(1):151–182, 2011.
- [4] K. Konczak and J. Lang. Voting procedures with incomplete preferences. In *IJCAI Workshop on Advances in Preference Handling*, pages 124–129, 2005.
- [5] P. Kouvelis and G. Yu. *Robust Discrete Optimization and Its Applications*. Kluwer, Dordrecht, 1997.
- [6] B. Llamazares and T. Peña. Aggregating preferences rankings with variable weights. *European Journal of Operational Research*, 230(2):348–355, 2013.
- [7] T. Lu and C. Boutilier. Robust approximation and incremental elicitation in voting protocols. In *Proceedings of IJCAI 2011*, pages 287–293, 2011.
- [8] L. Naamani-Dery, I. Golan, M. Kalech, and L. Rokach. Preference elicitation for group decisions using the Borda voting rule. *Group Decision and Negotiation*, 24(6):1015–1033, Nov 2015.
- [9] M. S. Pini, F. Rossi, K. B. Venable, and T. Walsh. Aggregating partially ordered preferences. *J. Log. Comput.*, 19(3):475–502, 2009.
- [10] A. Salo and R. P. Hämäläinen. Preference ratios in multiattribute evaluation (PRIME)–elicitation and decision procedures under incomplete information. *IEEE Trans. on Systems, Man and Cybernetics*, 31(6):533–545, 2001.

- [11] L. J. Savage. *The Foundations of Statistics*. Wiley, New York, 1954.
- [12] W. E. Stein, P. J. Mizzi, and R. C. Pfaffenberger. A stochastic dominance analysis of ranked voting systems with scoring. *European Journal of Operational Research*, 74(1):78 – 85, 1994.
- [13] P. Viappiani. Positional scoring rules with uncertain weights. In *Scalable Uncertainty Management - 12th International Conference, SUM 2018, Milan, Italy, October 3-5, 2018, Proceedings*, pages 306–320, 2018.
- [14] L. Xia and V. Conitzer. Determining possible and necessary winners under common voting rules given partial orders. In *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence, AAAI 2008, Chicago, Illinois, USA, July 13-17, 2008*, pages 196–201, 2008.