



Bayesian Vote Elicitation for Group Recommendations

Maeva Caillat, Nicolas N. Darcel, Cristina Manfredotti, Paolo Viappiani

► To cite this version:

Maeva Caillat, Nicolas N. Darcel, Cristina Manfredotti, Paolo Viappiani. Bayesian Vote Elicitation for Group Recommendations. From Multiple Criteria Decision Aid to Preference Learning (DA2PL 2020), Andrea Passerini (University of Trento); Vincent Mousseau (Centrale Supélec), Nov 2020, Trento, Italy. hal-03006639

HAL Id: hal-03006639

<https://hal.sorbonne-universite.fr/hal-03006639>

Submitted on 16 Nov 2020

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.

Bayesian Vote Elicitation for Group Recommendations

Maeva Caillat¹ and Nicolas Darcel² and Cristina Manfredotti³ and Paolo Viappiani⁴

Abstract. Elicitation of preferences is a critical task in modern application of voting protocols such as group recommender systems.

This paper introduces a Bayesian elicitation paradigm for social choice. The system maintains a probability distribution over the preferences (rankings) of the voters. At each step the system asks the question to one of the voters, and the distribution is conditioned on the response. We consider strategies to pick the next question based on value of information, conditional entropy, and a mix of these two notions.

We develop this idea focusing on scoring rules and compare different elicitation strategies in the case of Borda rule.

1 Introduction

Aggregation of preferences is an important task studied in social choice [5] and as well in the more recent research field of group recommender systems [1].

In modern applications of social choice and as well in recommender systems, it cannot be given for granted that preferences are readily available. Realizing this fact, a number of researchers considered voting procedures with incomplete preferences and as well elicitation procedures for voting; to our knowledge the first of these work is the one by Konczak and Lang [3] that introduced the notion of possible winners and necessary winners. Previous approaches considered robust winner approximation and elicitation with the minimax regret criterion [4]. A key insight is that, often, not all user preferences are needed in order to reach a “winning item”.

In our work, we consider a Bayesian approach, as Dery et al. [6], but our solution differs in a number of key points. We provide a principled quantification of uncertainty using the notion of expected loss allowing a principled termination condition, and we develop more effective elicitation strategies, as shown by our experiments.

We focus on Borda count, but our method can be extended to other social choice functions.

Notation and preliminaries We assume that $V = \{1, \dots, n\}$ is the set voters and C the set of candidates; with $|C| = m$. A profile (r^1, \dots, r^n) is a vector of linear orders on C , one for each voter $i \in V$. The set of all $m!$ linear orders on C is denoted as L ; hence the set of all possible profiles is L^n . We use $r(x)$ to denote the position of alternative x in ranking r ; and r_i to denote the alternative in the i -th position in r .

A social choice function (or voting rule) $f : L^n \rightarrow 2^C \setminus \emptyset$ maps a profile to a non-empty subset of candidates (the “winners”). Scoring rules are social choice functions that rank alternatives accord-

ing to their score, computed by summing up the number of points they receive in each ranking; the score of x with respect to rankings r^1, \dots, r^n is:

$$s(x; r^1, \dots, r^n) = \sum_{i=1}^n w(r^i(x)), \quad (1)$$

where w is a function that assigns each position $1, \dots, m$ to a number of points.

A special type of scoring rule is Borda, that assumes the weights to be: $w(i) = m - i$; the first ranked item is assigned $m - 1$ points, the second obtains $m - 2$, etc.

2 Bayesian vote elicitation

We adopt a Bayesian approach to preference elicitation and approximate winner determination. The system maintains a probability distribution $\mathbb{P}(r^1, \dots, r^n)$ over the preferences (rankings) r^1, \dots, r^n of the voters. The distributions give zero probability to all rankings that are not completion of the known preferences of the voters. We assume that the voters preferences are independent, thus $\mathbb{P}(r^1, \dots, r^n) = \mathbb{P}(r^1) \cdot \dots \cdot \mathbb{P}(r^n)$.

Our incremental elicitation approach is based on looping through the following steps:

- It computes the current best alternative x^* achieving the highest score *in expectation*
- If x^* meets the stopping criterion, the procedure stops and outputs x^*
- Otherwise, it selects an elicitation question to ask the user, and asks it.
- It conditions the probability distribution on the response.

We now discuss the different steps.

Computing expected scores We identify the “approximate winner” as the candidate that yields the highest expected score under the current probability distribution of preference rankings. Each alternative x is associated to its expected score $\mathbb{E}[s(x)]$, that we denote as $\bar{s}(x)$. We observe that for scoring rules $\bar{s}(x)$ can be efficiently computed as:

$$\bar{s}(x) = \sum_{i=1}^n \sum_{r^i \in L} \mathbb{P}(r^i) w(r^i(x))$$

and for Borda, the expression further simplifies to:

$$\bar{s}(x) = mn - \sum_{i=1}^n \sum_{r^i \in L} \mathbb{P}(r^i) r^i(x).$$

Let $s^* = \max_{x \in C} \mathbb{E}[s(x)]$ be the maximum value of expected score given the current uncertainty, and x^* the associated alternative (the “winner in expectation”); i.e. $s^* = \bar{s}(x^*)$.

¹ AgroParisTech and Ecole Centrale Nantes, email: caillat.maeva@gmail.com

² AgroParisTech, email: nicolas.darcel@agroparistech.fr

³ AgroParisTech, email: cristina.manfredotti@agroparistech.fr

⁴ LIP6, CNRS and Sorbonne université, email: paolo.viappiani@lip6.fr

Expected loss and stopping criterion At some point of the interaction, we have a posterior distribution over the rankings, and an associated alternative maximizing the expected score. We estimate the regret or loss of stopping the elicitation and recommending x^* . The user's loss is the difference between her expected utility, under true preferences, of the optimal alternative x^* , and her expected utility under the recommended alternative x .

The loss $\ell(r^1, \dots, r^n)$ is the regret of choosing x^* occurred when the true voters preferences are (r^1, \dots, r^n) :

$$\ell(r^1, \dots, r^n) = \max_{y \in C} s(y; r^1, \dots, r^n) - s(x^*; r^1, \dots, r^n).$$

Since we do not know the true preferences, but we know their distributions, we consider the *expected loss* $\mathbb{E}[\ell]$ (in a way analogous to [2] that considered Bayesian elicitation in influence diagrams) that quantifies how far we are from the true optimum in expectation:

$$\begin{aligned} \mathbb{E}_{r^1 \sim \mathbb{P}(r^1), \dots, r^n \sim \mathbb{P}(r^n)}[\ell(r^1, \dots, r^n)] &= \\ &= \left[\sum_{r^1 \in L} \dots \sum_{r^n \in L} \mathbb{P}(r^1) \dots \mathbb{P}(r^n) \max_{y \in C} s(y; r^1, \dots, r^n) \right] - s^*. \end{aligned}$$

In order to compute the above expression we approximate it using a Monte Carlo method. We sample the voters preference rankings from $\mathbb{P}(r^1), \dots, \mathbb{P}(r^n)$, compute the scores of alternatives, and compute the loss for these preferences. We repeat the procedure N times and take the average. To set N (the number of samples) we use Chebyshev inequality. N should be at least $\frac{b^2}{4\delta\epsilon^2}$ where ϵ is the required precision, δ is the confidence and b is an upper bound of the variance; in our case we set $b = n(m-1) - s^*$ (the highest possible Borda score less the current best expected score).

The elicitation procedure continues until the expected loss is lower than a given threshold. If the goal is to find a necessary winner with certainty, the threshold can be set to zero.

Elicitation strategies We consider different strategies to decide which query to ask at any stage of the elicitation process. The strategies aim at reducing uncertainty over the voters preferences in such a way to improve the quality of the approximated winner.

We focus on pairwise comparisons. In the following we denote with $q_{a,b}^v$ the query asking voter v to compare alternatives a and b .

The first approach, adopted from [6], is called *Information Gain for Borda (IGB)* and is based on the notion of entropy. Define $P_{\text{win}}(a)$ as the probability that a wins, that can be found by summing up the probability of all preference combinations that make a a winner. Let $H(W)$ be the entropy associated to the distribution P_{win} . The query $q_{a,b}^v$ is associated with its *information gain* (i.e. the conditional entropy):

$$\text{IG}(q_{a,b}^v) = p_{a \succ_v b}^v H(W|a \succ_v b) + p_{b \succ_v a}^v H(W|b \succ_v a).$$

where $p_{a \succ_v b}^v$, the probability that voter v prefers a to b , can be computed by marginalization.


The second approach, *ESB*, also from [6], computes the *a posteriori* improvement of the maximum of P_{win} (details omitted).

The third strategy adopts myopic *Expected Value of Information (EVOI)*, that has been shown to be very effective [7] in single-user preference elicitation. $\text{EVOI}(q_{a,b}^v)$ is

$$p_{a \succ_b}^v \max_{x \in C} \mathbb{E}[s(x)|a \succ b] + p_{b \succ_a}^v \max_{x \in C} \mathbb{E}[s(x)|b \succ a] - s^*.$$

The selected query is the one with highest EVOI.

While EVOI can often identify very informative queries, in preliminary tests we realized that it can sometimes happen that *myopic*



EVOI of all candidate queries is null. Motivated by this observation, we designed *EVOI+IGB* that asks the query with highest EVOI if its value is positive, and otherwise asks a question using IGB.

Updating the distributions Whenever a query is answered, the distributions are updated using Bayes theorem. In fact, since there is no noise in user feedback, this means assigning zero probability to rankings that are inconsistent with the user's input and to renormalize.

3 Experiments

We provide some preliminary experimental results evaluating the performance of the Bayesian elicitation approach comparing the different elicitation strategies. In the Figure above we plot the expected loss as a function of the number of questions asked using the sushi dataset⁵.

The experiments show that a necessary winner can be found with relatively few questions. Somewhat surprisingly, we found ESB to perform worse than IGB, contrary to what reported in [6]. EVOI+IGB is the most efficient query strategy in our tests.

Future Works We are currently testing the approach in a realistic setting of food recommendations.

Since the current *flat* representation of distributions is not scalable, we are planning to adopt probabilistic ranking models, such as Plackett-Luce. We are also interested in dealing with other aggregation methods and handling more query types.

REFERENCES

- [1] M. Stettinger A. Felfernig, L. Boratto and M. Tkalcić, *Group Recommender Systems. An Introduction*, Springer, 2018.
- [2] Urszula Chajewska, Daphne Koller, and Ronald Parr, 'Making rational decisions using adaptive utility elicitation', in *Proceedings of AAAI 2000*.
- [3] Kathrin Konczak and Jerome Lang, 'Voting procedures with incomplete preferences', *Proceedings of the Multidisciplinary IJCAI-05 Workshop on Advances in Preference Handling*, (01 2005).
- [4] Tyler Lu and Craig Boutilier, 'Preference elicitation and robust winner determination for single- and multi-winner social choice', *Artif. Intell.*, **279**, (2020).
- [5] Nicolas Maudet, Ulle Endriss, and Yann Chevaleyre, 'A short introduction to computational social choice', in *Proceedings of the 33rd conference on Current Trends in Theory and Practice of Computer Science*, (01 2005).
- [6] Lihi Naamani-Dery, Inon Golan, Meir Kalech, and Lior Rokach, 'Preference elicitation for group decisions using the Borda voting rule', *Group Decision and Negotiation*, **24**(6), 1015–1033, (2015).
- [7] Paolo Viappiani and Craig Boutilier, 'On the equivalence of optimal recommendation sets and myopically optimal query sets', *Artificial Intelligence*, **286**, 103328, (2020).

⁵ <http://www.kamishima.net/sushi/>