Complexity of Election Evaluation and Probabilistic Robustness

Extended Abstract

Dorothea Baumeister
Institut für Informatik
Heinrich-Heine-Universität Düsseldorf
d.baumeister@uni-duesseldorf.de

Tobias Hogrebe
Institut für Informatik
Heinrich-Heine-Universität Düsseldorf
tobias.hogrebe@uni-duesseldorf.de

ABSTRACT

When dealing with election data it is reasonable to assume that the votes are incomplete or noisy. The reasons are manifold and range from cost-intensive elicitation to manipulation. We study the important questions of evaluating elections with incomplete data and the robustness of elections with noisy data from a computational point of view. To capture different motivations, we consider three models for the distribution of preferences: the uniform distribution over the completions of incomplete preferences inspired by the possible winner problem, the dispersion around complete preferences, also called Mallows noise model, and a model in which the distribution over the votes of each voter is explicitly given. We consider both approval vector preferences and linear order preferences and show that the complexity of the problem can vary greatly depending on the voting rule, the distribution model, and the parameterization.

KEYWORDS

probabilistic social choice; computational complexity; voting

ACM Reference Format:
Dorothea Baumeister and Tobias Hogrebe. 2020. Complexity of Election Evaluation and Probabilistic Robustness. 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

Elections are an integral part of any democracy, be it for the collective decision-making of a whole country or just for any group of people, a sports club or employees of a company. In addition to these classic applications, elections are also considered in connection with software agents and automation. Here, the applications of elections range from multi-agent planning (see, e.g., Ephrati and Rosenschein [6]) and meta-search engines (see, e.g., Dwork et al. [5]) to recommender systems (see, e.g., Ghosh et al. [8]) and email classification (see, e.g., Cohen et al. [3]). In the classic case, we assume that we have perfect knowledge about the preferences of the voters or agents regarding the candidates or alternatives and are able to use a voting rule to determine the, with respect to the specific rule, rightful winners. However, in many realistic scenarios, we can not assume that we have perfect information about voter preferences. Nevertheless, a decision must often be made or at least in some way a result of the election must be presented. The reasons for imperfect election data are manifold. First, we often can not assume that the election data we receive is complete. In the case of actual elections, the collection of complete election data is often cost-intensive, complicated, or simply not possible under the given circumstances. The same holds for the creation of election forecasts based on partial data aggregated from social networks or polls, where a complete collection of election data is not appropriate. On the other hand, even if we receive complete election data, in many situations we can not assume that it has not been corrupted in transmission, by manipulation, or through the elicitation itself. In these situations, the question arises how robust and thereby justified a candidate's victory is if we assume that the election data has been corrupted to a certain degree.

Therefore, we study the problem of determining the probability that a particular candidate wins an election for a given distribution over the preferences of the voters. Conitzer and Sandholm [4] were the first to study this problem and called it the evaluation problem. The relevance of the problem is immense, as it captures many different, and in particular the previously presented, scenarios, such as the winner determination on incomplete data, the creation of election forecasts, and the examination of the justification or robustness of a candidate's victory if corruption of the data is possible. To cover those different motivations, we consider three models for the distribution of preferences: the uniform distribution over the completions of incomplete preferences inspired by the possible winner problem (see Bachrach et al. [1]), the dispersion around complete preferences, also called Mallows noise model (see Mallows [10]), and a model in which the distribution over the votes of each voter is explicitly given (see Conitzer and Sandholm [4]).

2 SCENARIO

Formally an election is a pair (C, V), with $C = \{c_1, \dots, c_m\}$ with $m \geq 2$ being the set of *candidates* and $V = (v_1, \dots, v_n)$ with $n \geq 1$ a *profile* consisting of n votes over C. We consider the two most prominent types of votes: approval vectors and linear orders. In the case of approval vectors, we use *approval voting* (AV) to determine the winners. By k-AV we denote the variant of AV in which each voter must distribute exactly k approvals with fixed $k \geq 1$ for $m > k$. In the case of linear orders we consider positional scoring rules, namely k-*approval* and k-*veto* with fixed $k \geq 1$ for $m > k$, *Borda* and the scoring rule characterized by the vector $(2, 1, \dots, 1, 0)$. Note that k-AV and k-approval essentially describe the same voting rule and differ only in the amount of information we are given about the preferences of the voters, since for k-approval we additionally have given rankings over all candidates. Interestingly, this very distinction leads to differing complexity results in some cases, as we will see later. In the course of this work we also encounter elections with partial information, in the form of partial profiles, i.e. profiles consisting of partial votes. In the case of approval vectors,

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.

Table 1: Complexity results for $\mathcal{E}$-Evaluation in the non-unique winner case.

	PPIC		Mallows		EDM	
	general	const. n	general	const. n	general	const. n
AV	FP		FP		#P-hard	FP
k-AV ($k \geq 1$)	#P-hard	FP	?	FP	#P-hard	FP
k-approval ($k \geq 1$), k-veto ($k \geq 1$), $(2, 1, \ldots, 1, 0)$	#P-hard	#P-hard ($n = 1$)	#P-hard	FP	#P-hard	FP
Borda			#P-hard	FP ($n = 1$)		

a partial vote is represented by an approval vector in which entries may be undetermined. In the case of linear orders, a partial vote consists of a partial order over the set of candidates that is, an irreflexive and transitive, but on the contrary to linear orders, not necessarily connex relation. The set of all possible completions of a profile $\tilde{V}$ consisting of partial votes is denoted by $\Lambda(\tilde{V})$.

As mentioned earlier, there may be some uncertainty about the votes in elections due to several reasons or an election should be evaluated according to some predictions. Thus we assume some distribution over possible profiles, and then investigate the problem of determining the winning probability of a certain candidate. This is formalized as follows for a given voting rule $\mathcal{E}$.

$\mathcal{E}$-Evaluation	
Given:	A set of candidates C, a distribution $\mathcal{P}$ of profiles over C, and a candidate $p \in C$.
Question:	What is the probability Φ that p is a winner of the election with respect to $\mathcal{E}$ assuming $\mathcal{P}$?

Here we focus mainly on the *non-unique winner* case. The distribution as part of the input means that the respective distribution is specified by the necessary parameters as part of the input. In the following we will present the three distributions for profiles considered in this paper.

PPIC. The first distribution model we consider is the normalized variant of the possible winner motivated model of Bachrach et al. [1]. We refer to this model as *partial profile impartial culture model* (PPIC). Given a set of candidates C, a partial profile $\tilde{V} = (\tilde{v}_1, \ldots, \tilde{v}_n)$ over C. The probability of a profile $V = (v_1, \ldots, v_n)$ over C according to PPIC is given by $\Pr_{\text{PPIC}}(V \mid \tilde{V}) = 1/|\Lambda(\tilde{V})|$. Thereby, each completion of the partial profile is equally likely, hence 'impartial culture'. Referring back to the motivations stated at the beginning, PPIC covers the scenarios in which we want to make an election forecast based on partial data aggregated from social networks or polls, or perform a winner determination on partial data. Note, that for linear orders, the computation of the probability of a given profile is #P-hard, since the calculation of the normalization $|\Lambda(\tilde{V})|$ itself is already a #P-hard problem as shown by Brightwell and Winkler [2].

Mallows. The second model we are considering is the *Mallows noise model* (Mallows [10]). The basic idea is that some reference profile is given, and the probability of another profile is measured according to its distance to the reference profile. Given a set of candidates C, a profile $\hat{V} = (\hat{v}_1, \ldots, \hat{v}_n)$ over C and dispersion $\varphi \in (0, 1)$, the probability of a profile $V = (v_1, \ldots, v_n)$ over C according to the Mallows model is given by $\Pr_{\text{Mallows}}(V \mid \hat{V}, \varphi) = \varphi^{d(V,\hat{V})}/Z^n$ with distance d and normalization constant Z. In the original case of linear orders, the summed up vote-wise swap distance is used. In the case of approval vectors, we propose to use the summed up vote-wise Hamming distance. Mallows noise model captures the scenarios in which the data was corrupted to a certain degree or we expect the preferences to have changed over time and we are interested in how likely and thereby robust the victory of a candidate is.

EDM. Finally, we consider a model introduced by Conitzer and Sandholm [4] and later studied by Hazon et al. [9]. Due to its nature, we refer to it as the *explicit distribution model* (EDM). Given a set of candidates C, for each voter $i \in \{1, \ldots, n\}$ we are given a probability distribution π_i over the votes over C through a list of votes paired with their non-zero probabilities. Each unspecified vote has probability 0. The probability of a profile $V = (v_1, \ldots, v_n)$ over C according to EDM for $\pi = (\pi_1, \ldots, \pi_n)$ is given by $\Pr_{\text{EDM}}(V \mid \pi) = \prod_{i=1}^{n} \pi_i(v_i)$. In its generality, EDM can be used to cover various distributions and thus scenarios, but may require lists with exponential length.

We summarize our main results for the non-unique winner case in Table 1. The results presented here additionally hold for the *unique* winner case and *random* and *lexicographic tie-breaking*, except for k-veto assuming PPIC with a constant number of voters in the unique winner case for which trivial FP results follow. For a constant number of voters we receive FP results for the distributions and voting rules considered here through applying the dynamic programming approach by Hazon et al. [9]. Our results regarding scoring rules assuming PPIC and EDM in Table 1 actually hold for all scoring rules for which the scoring vector can be determined efficiently.

We showed that the complexity of the problem varies greatly depending on the voting rule, the distribution model, and the parameterization. As future work, there are two open cases, namely the complexity regarding Mallows and k-AV in general and Borda for a constant number of voters. Besides that, it is interesting to consider further distribution models and the fine-grained parameterized counting complexity, see Flum and Grohe [7].

ACKNOWLEDGMENTS

This work is supported by the DFG-grant BA6270/1-1.

REFERENCES

[1] Y. Bachrach, N. Betzler, and P. Faliszewski. 2010. Probabilistic possible winner determination. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence*. AAAI Press, 697–702.

[2] G. Brightwell and P. Winkler. 1991. Counting linear extensions. *Order* 8, 3 (1991), 225–242.

[3] W. Cohen, R. Schapire, and Y. Singer. 1998. Learning to Order Things. In *Advances in Neural Information Processing Systems*. 451–457.

[4] V. Conitzer and T. Sandholm. 2002. Complexity of Manipulating Elections with Few Candidates. In *Proceedings of the 18th National Conference on Artificial Intelligence*. AAAI Press, 314–319.

[5] C. Dwork, R. Kumar, M. Naor, and D. Sivakumar. 2001. Rank Aggregation Methods for the Web. In *Proceedings of the 10th International Conference on World Wide Web*. ACM, 613–622.

[6] E. Ephrati and J. Rosenschein. 1993. Multi-Agent Planning as a Dynamic Search for Social Consensus. In *Proceedings of the 13th International Joint Conference on Artificial Intelligence*. 423–429.

[7] J. Flum and M. Grohe. 2004. The parameterized complexity of counting problems. *SIAM J. Comput.* 33, 4 (2004), 892–922.

[8] S. Ghosh, M. Mundhe, K. Hernandez, and S. Sen. 1999. Voting for movies: the anatomy of a recommender system. In *Proceedings of the 3rd Annual Conference on Autonomous Agents*. ACM, 434–435.

[9] N. Hazon, Y. Aumann, S. Kraus, and M. Wooldridge. 2012. On the evaluation of election outcomes under uncertainty. *Artificial Intelligence* 189 (2012), 1–18.

[10] C. Mallows. 1957. Non-null ranking models. *Biometrika* 44 (1957), 114–130.