

Voting with Partial Information: What Questions to Ask?

(Extended Abstract)

Ning Ding
Dept. of Comp. Sci. & Eng.
HKUST
dning@cse.ust.hk

Fangzhen Lin
Dept. of Comp. Sci. & Eng.
HKUST
flin@cse.ust.hk

ABSTRACT

Voting is a way to aggregate individual voters' preferences. Traditionally a voter's preference is represented by a total order on the set of candidates. However, sometimes one may not have complete information about a voter's preference, and in this case, can only model a voter's preference as a partial order. Given this framework, there has been work on computing the possible and necessary winners of a (partial) profile. In this paper, we take a step further, look at sets of questions to ask in order to determine the outcome of such a partial profile. Specifically, we call a set of questions a deciding set for a candidate if the outcome of the vote for the candidate is determined no matter how the questions are answered by the voters, and a possible winning (losing) set if there is a way to answer these questions to make the candidate a winner (loser) of the vote. We discuss some interesting properties about these sets of queries, prove some complexity results about them under some well-known voting rules such as plurality and Borda, and consider their application in vote elicitation.

Categories and Subject Descriptors

I.2 [Computing Methodologies]: ARTIFICIAL INTELLIGENCE

Keywords

Minimal Deciding Set, Comparison Query

1. INTRODUCTION

Traditionally, a voter's preference is assumed to be a complete linear order over possible candidates (outcomes, or alternatives). One can easily imagine situations where this assumption is too strong, either because the voter herself cannot rank all of the possibilities linearly or because as an observer, we do not have a complete knowledge about her preferences. In the context of voting, there has been work in this direction as well. Given a partial ordering for each voter, Konczak and Lang [5] considered the problem of deciding whether a candidate is a *necessary winner* and *possible winner*. A necessary winner is a candidate who is always a winner in every possible completion of the given partial preference profile, while a possible winner is one who is a winner in some of the completions. The complexities of these two problems under a variety of voting rules, especially the so-called positional scoring rules,

have been extensively studied [7, 8, 1]. More recently, Conitzer *et al.* [4] considered a notion of manipulations in voting with partial information.

In this paper, we continue this line of work. Given a partial preference profile, we consider in general how much additional information is still needed to make a particular candidate a winner or loser under a voting rule. If the candidate is already a necessary winner or a necessary loser, then no additional information is needed. Otherwise, one may want to know which voter is crucial in deciding the outcome, and for that voter what would be the important questions to ask. These are obviously important issues to consider when doing voter preference elicitation, and should have some interesting applications. For instance, in an election, a candidate's team may want to know that given what they already know about a group of people, whether more knowledge about their voting preferences would make any differences to the outcome.

2. PRELIMINARIES

We assume a finite set $N = \{1, \dots, n\}$ for voters, and a finite set O for candidates. A *preference ordering* p_i of a voter i is a total (linear) order on O , and a *preference profile* p is a tuple of preference orderings, one for each voter.

A voting rule f is a function from preferences profiles to non-empty sets of outcomes. For a preference profile p , $f(p)$ is the set of winners. When a single winner is desired, a tie-breaking rule can be used to select the one from $f(p)$. Or f is required to be single-valued.

We consider the situation when the preference ordering of a voter may not be total, either because the onlooker who is studying the voting does not have a complete knowledge of the voter's preference or that the voter herself is not certain of her own preferences.

Formally a *partial preference ordering* p_i of voter i is a partial order on the set O of candidates: for each $o \in O$, $(o, o) \in p_i$ (reflexivity), if both (o_1, o_2) and (o_2, o_1) are in p_i , then $o_1 = o_2$ (anti-symmetry), and if (o_1, o_2) and (o_2, o_3) are in p_i , then $(o_1, o_3) \in p_i$ (transitivity). A *partial preference profile* is then a tuple of partial preference orderings, one for each voter.

Given a partial preference ordering p_i , an *extension* of p_i is a partial preference ordering p'_i such that $p_i \subseteq p'_i$. An extension of p_i that is a total order is called a *completion* of p_i . A completion of a partial preference profile can be similarly defined.

Under a voting rule f , a candidate o is said to be a *necessary winner* of a partial preference profile p , if for all completion p' of p , $o \in f(p')$. If there exists such a completion, then o is said to be a *possible winner* [5].

Appears in: *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013)*, Ito, Jonker, Gini, and Shehory (eds.), May 6–10, 2013, Saint Paul, Minnesota, USA. Copyright © 2013, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

3. DEFINITIONS

Our interest in this paper is on getting additional information to decide the outcome of a vote. This additional information will be in the form of comparison queries [2] to voters.

DEFINITION 1. A (comparison) query to voter i is one of the form $i:\{a, b\}$ that asks i to rank candidates a and b .

When presented with the query $i:\{a, b\}$, the voter i has to answer either “a” (she prefers a over b) or “b” (she prefers b over a). Formally, an answer to a set Q of questions is a function σ from Q to O such that for any $i:\{a, b\} \in Q$, $\sigma(i:\{a, b\}) \in \{a, b\}$.

Let p be a partial preference profile and Q a set of queries. An answer σ to Q is legal under p if for each voter i , the transitive closure of the following set

$$p_i \cup \{(a, b) \mid i:\{a, b\} \in Q \wedge \sigma(i:\{a, b\}) = a\}$$

which we denote by $p_i(\sigma, Q)$, is a partial order on O , the set of candidates. Given a legal answer σ to Q under p , the resulting partial preference profile is then

$$p(\sigma, Q) = (p_1(\sigma, Q), \dots, p_n(\sigma, Q)),$$

In the following, unless stated otherwise, we always assume that answers to sets of questions are legal under the given partial preference profile.

We can now define the sets of questions that we are interested in this paper. A deciding set of queries for a candidate o determines the outcome of the vote for o no matter how the queries in the set are answered.

DEFINITION 2. Let p be a partial preference profile, o a candidate, and f a voting rule. A set Q of queries is a deciding set for o (in p under f) if for every answer σ , o is either a necessary winner or a necessary loser in the new partial profile $\sigma(p, Q)$ under f . Q is a minimal deciding set for o if it is a deciding set and there is no other deciding set Q' such that $Q' \subset Q$.

Sometimes one may also be interested in knowing the ways to make a candidate a winner or a loser in a vote. In this case, one may want to find sets of queries that when answered properly will lead to the candidate being a winner (or loser).

DEFINITION 3. Let p be a partial preference profile, o a candidate, and f a voting rule. A set Q of queries is a possible winning (losing) set for o (in p under f) if there is an answer σ such that o is a necessary winner (loser) in the new partial profile $\sigma(p, Q)$ under f . Q is a minimal possible winning (losing) set for o if it is a possible winning (losing) set for o , and there is no other possible winning (losing) set Q' for o such that $Q' \subset Q$.

4. MAIN RESULTS

THEOREM 1. For any voting rule f , partial preference profile p , and candidate o , there is a unique minimal deciding set for o in p under f .

THEOREM 2. Let S be the set of all comparison queries. For any candidate o , and any partial preference profile p , a query $q = i:\{a, b\}$ is in the minimal deciding set if and only if there is an answer σ to $S \setminus \{q\}$ such that it can be extended to two answers σ_1 and σ_2 to S such that $\sigma_1(q) = a$, $\sigma_2(q) = b$, and the outcome of o is different in $p(\sigma_1, S)$ and $p(\sigma_2, S)$.

This theorem is the basis for our algorithm for computing minimal deciding sets.

THEOREM 3. The problem of checking if a query q is in the minimal deciding set for a candidate o in a partial profile p is in P under plurality or veto rule, and NP -complete under Borda rule.

The same complexity results hold for deciding if a query set is a winning set.

5. VOTE ELICITATION

Our notion of deciding sets generalizes the notions of necessary and possible winners [5]. It is also closely related to work on vote elicitation (e.g. [3, 6]). Generally speaking, vote elicitation is about getting information about voters' preferences. Assuming that we are after voters' ranking of pairs of candidates, i.e. their answers to pair-wise queries, one important consideration is whether these queries are asked one at a time. Our notion of minimal deciding sets is obviously suitable for vote elicitation when we can ask a set of comparison questions. One can easily imagine situations when this is necessary. For instance, given the time and resource constraints, there may be only one chance to communicate with voters. In other cases, we may have the luxury of asking questions one at a time. This means that we can ask one question first, and depending on the answer given to that question, decide which question to ask next. To distinguish these two different ways of doing vote elicitation, we call the latter dynamic vote elicitation and former one-step vote elicitation.

For one-step vote elicitation, if we are interested in a particular candidate, then the obvious way is to compute the minimal deciding set for this candidate, and send out the questions in the minimal deciding set to respective voters. If we want to find out the outcomes for a set of candidates, we will need to compute the minimal deciding set for each of the candidates, and send out the union of all the minimal deciding sets to the voters.

A straightforward way of using minimal deciding set for dynamic vote elicitation is to pick a query from the unique minimal deciding set in the current partial profile. This simple strategy can be effective in many situations, but may not be optimal.

6. FUTURE WORK

There are many possible directions for future work, including more general notions of partial information and query sets, as well as experiments on the above strategy for dynamic vote elicitation.

7. REFERENCES

- [1] N. Betzler and D. Britta. Towards a dichotomy for the possible winner problem in elections based on scoring rule. *Journal of Computer and System Sciences*, 76:812–836, 2010.
- [2] V. Conitzer. Eliciting single-peaked preferences using comparison queries. *JAIR*, 35:161–191, 2009.
- [3] V. Conitzer and T. Sandholm. Vote elicitation: Complexity and strategy-proofness. In *AAAI-02*, pages 392–397. AAAI, 2002.
- [4] V. Conitzer, T. Walsh, and L. Xia. Dominating manipulations in voting with partial information. In *IJCAI-2011 Workshop on Social Choice and Artificial Intelligence*, 2011.
- [5] K. Konczak and J. Lang. Voting procedures with incomplete preferences. In *Proc. IJCAI- 05 Multidisciplinary Workshop on Advances in Preference Handling*, 2005.
- [6] Ariel D. P. *Computational Voting Theory: Of the Agents, By the Agents, For the Agents*. PhD thesis, The Hebrew University of Jerusalem, Sep. 2008.
- [7] M.S. Pini, F. Rossi, K. B. Venable, and T. Walsh. Incompleteness and incomparability in preference aggregation. In *IJCAI-07*, pages 1464–1469, 2007.
- [8] L. Xia and V. Conitzer. Determining possible and necessary winners under common voting rules given partial orders. *JAIR*, 41:25–67, May 2011.