Restricted Manipulation in Iterative Voting: Condorcet Efficiency and Borda Score

Umberto Grandi[1], Andrea Loreggia[1], Francesca Rossi[1], Kristen Brent Venable[2], and Toby Walsh[3]

[1] University of Padova
{umberto.uni,andrea.loreggia}@gmail.com, frossi@math.unipd.it
[2] Tulane University and IHMC
kvenabl@tulane.edu
[3] NICTA and UNSW
toby.walsh@nicta.com.au

Abstract. In collective decision making, where a voting rule is used to take a collective decision among a group of agents, manipulation by one or more agents is usually considered negative behavior to be avoided, or at least to be made computationally difficult for the agents to perform. However, there are scenarios in which a restricted form of manipulation can instead be beneficial. In this paper we consider the iterative version of several voting rules, where at each step one agent is allowed to manipulate by modifying his ballot according to a set of restricted manipulation moves which are computationally easy and require little information to be performed. We prove convergence of iterative voting rules when restricted manipulation is allowed, and we present experiments showing that iterative restricted manipulation yields a positive increase in the Condorcet efficiency and Borda score for a number of standard voting rules.

1 Introduction

In multi-agent systems often agents need to take a collective decision. A voting rule can be used to decide which decision to take, mapping the agents' preferences over the possible candidate decisions into a winning decision for the collection of agents. In these kind of scenarios, it may be desirable that agents do not have any incentive to manipulate, that is, to misreport their preferences in order to influence the result of the voting rule in their favor. Indeed, manipulation is usually seen as bad behavior from an agent, to be avoided or at least to be made computationally difficult to accomplish. While we know that every voting rule is manipulable when no domain restriction is imposed on the agents' preferences (Gibbard, 1973; Satterthwaite, 1975), we can at least choose a voting rule that is computationally difficult to manipulate for single agents or coalitions.

In this paper we consider a different setting, in which instead manipulation is allowed in a fair way. More precisely, agents express their preferences over the set of possible decisions and the voting rule selects the current winner as in the usual case. However, this is just a temporary winner, since at this point a single agent may decide to manipulate, i.e., to change her preference if by doing so the result changes in her favor. The process repeats with a new agent manipulating until we eventually reach a convergence state, i.e., a profile where no single agent can get a better result by manipulating. We call

P. Perny, M. Pirlot, and A. Tsoukiàs (Eds.): ADT 2013, LNAI 8176, pp. 181–192, 2013.

such a process *iterative voting*. In this scenario, manipulation can be seen as a way to achieve consensus, to give every agent a chance to vote strategically (a sort of fairness), and to account for inter-agent influence over time.

There are two prototypical situations in which iterative manipulation takes place. The first example is represented by the response of an electorate to a series of information polls about the result of a political election. At each step individuals may realize that their favorite candidate does not have chances to win and report a different preference in the subsequent poll. The second example is Doodle,[1] a very popular on-line system to select a time slot for a meeting. In Doodle, each participant can approve as many time slots as she wants, and the winning time slot is the one with the largest number of approvals. At any point, each participant can modify her vote in order to get a better result, and this can go on for several steps.

Iterative voting has been the subject of numerous publications in recent years. Previous work has focused on iterating the plurality rule (Meir *et al.*, 2010), on the problem of convergence for several voting rules (Lev and Rosenschein, 2012), and on the convergence of plurality decisions between multiple agents (Airiau and Endriss, 2009). Lev and Rosenschein (2012) showed that, if we allow agents to manipulate in any way they want (i.e., to provide their best response to the current profile), then the iterative version of most voting rules do not converge. Therefore, an interesting problem is to seek restrictions on the manipulation moves to guarantee convergence of the associated iterative rule. Restricted manipulation moves are good not only for convergence, but also because they can be easier to accomplish for the manipulating agent. In fact, contrarily to what we aim for in classical voting scenarios, here we do not want manipulation to be computationally difficult to achieve. It is actually desirable that the manipulation move be easy to compute while not requiring too much information for its computation.

An example of a restricted manipulation move is called k-pragmatists in Reijngoud and Endriss (2012): a k-pragmatist just needs to know the top k candidates in the collective candidate order, and will move the most preferred of those candidates to the top position of her preference. To compute this move, a k-pragmatist needs very little information (just the top k current candidates), and with this information it is computationally easy to perform the move. This move assures convergence of all positional scoring rules, Copeland, and Maximin, with linear tie-breaking. Note that each agent can apply this manipulation rule only once (since the top k candidates are always the same), and this is the main reason for convergence.

In this paper we introduce two restricted manipulation moves within the scenario of iterative voting and we analyze some of their theoretical and practical properties. Both manipulation moves we consider are polynomial to compute and require little information to be used. We show that convergence is guaranteed under both moves for those rules we consider, except for STV for which we only have experimental evidence of convergence. Moreover, we show that if a voting rule satisfies some axiomatic properties, such as Condorcet consistency or unanimity, then its iterative version will also satisfy the same properties as well. We then perform an experimental analysis of four restrictions on the set of manipulation strategies. For voting rules that are not Condorcet consistent, we test whether their Condorcet efficiency (that is, the probability to elect

[1] `http://doodle.com/`