Strategic Majoritarian Voting with Propositional Goals

Extended Abstract

Arianna Novaro
IRIT, University of Toulouse
France
arianna.novaro@irit.fr

Umberto Grandi
IRIT, University of Toulouse
France
umberto.grandi@irit.fr

Dominique Longin
IRIT, CNRS, Toulouse
France
dominique.longin@irit.fr

Emiliano Lorini
IRIT, CNRS, Toulouse
France
emiliano.lorini@irit.fr

ABSTRACT

We study strategic behaviour in goal-based voting, where agents take a collective decision over multiple binary issues based on their individual goals (expressed as propositional formulas). We focus on three generalizations of the issue-wise majority rule, and study their resistance to manipulability in the general case, as well as for restricted languages for goals. We also study how computationally hard it is for an agent to know if they can profitably manipulate.

KEYWORDS

Computational Social Choice; Strategic Voting; Knowledge Representation; Preference Modeling

ACM Reference Format:
Arianna Novaro, Umberto Grandi, Dominique Longin, and Emiliano Lorini. 2019. Strategic Majoritarian Voting with Propositional Goals. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019,* IFAAMAS, 3 pages.

1 INTRODUCTION

A key aspect of agent-based architectures is endowing agents with goals [20], and propositional goals in particular are ubiquitous in models of strategic reasoning. When taking collective decisions in a multi-issue domain, agents share the control over the variables at stake while still holding individual goals. This happens, for instance, when they need to arrange a business meal and have to decide on its specifics: should the restaurant be in the center, should it be fancy, should the meal be dinner or lunch? First, agents need a procedure to decide over each issue. Second, strategic behavior needs to be taken into account. Two frameworks have been proposed in the AI literature to solve this and similar problems: belief merging (see, e.g., Konieczny and Pino Pérez [13]) and goal-based voting [18]. Given our primary concern of resoluteness of the voting outcome we choose the latter framework and we focus on majoritarian rules.

The appeal of majority lies not only in its intuitive definition and extensive application in real-world scenarios, but also on having been widely studied in the related fields of voting theory and judgment aggregation [2, 17]. When moving to goal-based voting many definitions of majority are possible. The three adaptations studied here strike a balance between different needs: that of providing a resolute result, and that of treating each issue independently while still considering the complex structure of propositional goals.

Each of these majoritarian goal-based voting rules is analyzed with respect to their resistance to several manipulation strategies. Negative results, i.e., finding that a rule *can* be manipulated in the general case, lead us to study the computational complexity of manipulation, as well as restricting the language of individual goals in the hope of discovering niches of strategy-proofness.

2 FORMAL FRAMEWORK

We recall the framework of goal-based voting by Novaro et al. [18].

A group of *agents*, represented by set $\mathcal{N} = \{1, \ldots, n\}$, has to take a collective decision over a number of *issues*, represented by set $\mathcal{I} = \{1, \ldots, m\}$ of propositional variables. We let $\mathcal{L}_\mathcal{I}$ be the propositional language over the atoms in $\mathcal{I}$, with the usual boolean connectives. Agent i expresses her *individual goal* by a consistent propositional formula γ_i of $\mathcal{L}_\mathcal{I}$. The languages $\mathcal{L}^\star$ for $\star \in \{\wedge, \vee, \oplus\}$, defined by the following BNF grammars $\varphi := p \mid \neg p \mid \varphi \star \varphi$, represent restrictions on the language of goals. A *goal-profile* $\mathbf{\Gamma} = (\gamma_1, \ldots, \gamma_n)$ collects the goals of all n agents.

An *interpretation* or *alternative* is a function $v : \mathcal{I} \to \{0, 1\}$ associating a binary value with each variable in $\mathcal{I}$, where 0 means the issue is rejected and 1 that is accepted. We assume that there is no integrity constraint: all interpretations over the issues are allowed. We write $v \models \varphi$ to indicate that interpretation v makes φ true (i.e., v is a *model* of φ) and the set $\text{Mod}(\varphi) = \{v \mid v \models \varphi\}$ contains all the models of formula φ. We denote the choices of agent i for issue j in the models of her goal γ_i as $v_i(j) = (m^1_{ij}, m^0_{ij})$, where $m^x_{ij} = |\{v \in \text{Mod}(\gamma_i) \mid v(j) = x\}|$ for $x \in \{0, 1\}$. Abusing notation, we let $v_i(j) = x$ in case $|\text{Mod}(\gamma_i)| = 1$ and $m^x_{ij} = 1$.

A *goal-based voting rule* is a function for any n and m defined as $F : (\mathcal{L}_\mathcal{I})^n \to \mathcal{P}(\{0, 1\}^m) \setminus \{\emptyset\}$. If on every goal-profile F returns a singleton we call the rule *resolute*, and *irresolute* otherwise. The total number of acceptances and rejections of issue j in the outcome of $F(\mathbf{\Gamma})$ are defined as $F(\mathbf{\Gamma})_j = (F(\mathbf{\Gamma})^0_j, F(\mathbf{\Gamma})^1_j)$, where $F(\mathbf{\Gamma})^x_j = |\{v \in F(\mathbf{\Gamma}) \mid v_j = x\}|$ for $x \in \{0, 1\}$. In case we have $F(\mathbf{\Gamma})^x_j = 0$, we write $F(\mathbf{\Gamma})_j = 1 - x$ for simplicity.

The following are three variants of issue-wise majority defined for goal-based voting [18]. Firstly, *EMaj* interprets majority as the quota rule that accepts an issue if and only if more

Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), N. Agmon, M. E. Taylor, E. Elkind, M. Veloso (eds.), May 13–17, 2019, Montreal, Canada.

than half of the total number of votes are in its favor. Formally, $EMaj(\mathbf{\Gamma})_j = 1$ *iff* $\sum_{i \in \mathcal{N}} (\sum_{v \in \text{Mod}(\gamma_i)} \frac{v(j)}{|\text{Mod}(\gamma_i)|}) \geq \lceil \frac{n+1}{2} \rceil$. To guarantee *equality* among agents submitting formulas having a varying number of models, *EMaj* weights each model of an agent's goal inversely proportional to the total number of models of her goal.

Secondly, *TrueMaj* compares the total acceptances with the total rejections for an issue, setting the result to 1 (respectively, 0) if it is higher (respectively, lower) and to both 0 and 1 when tied. Formally, $TrueMaj(\mathbf{\Gamma}) = \Pi_{j \in \mathcal{I}} M(\mathbf{\Gamma})_j$ where for all $j \in \mathcal{I}$:

$$M(\mathbf{\Gamma})_j = \begin{cases} \{x\} & \text{if } \sum_{i \in \mathcal{N}} \frac{m_{ij}^x}{|\text{Mod}(\gamma_i)|} > \sum_{i \in \mathcal{N}} \frac{m_{ij}^{1-x}}{|\text{Mod}(\gamma_i)|} \\ \{0, 1\} & \text{otherwise} \end{cases}$$

As with *EMaj*, the models of an individual goal are weighted inversely to the total number of models for that goal.

Finally, *2sMaj* first applies majority to the models of the agents' individual goals, and then again to the result of the first step of aggregation, i.e., $2sMaj(\mathbf{\Gamma}) = Maj(Maj(\text{Mod}(\gamma_1)), \dots, Maj(\text{Mod}(\gamma_n)))$.

3 MANIPULATION OF MAJORITY RULES

The induced preference relation on the alternatives is dichotomous: agents equally prefer any model of their goal to any counter-model. As for irresolute rules the outcome may be a set of interpretations, different notions of satisfaction could be defined depending on how an agent compares two sets of interpretations.

Let $sat : \mathcal{L}_{\mathcal{I}} \times (\mathcal{P}(\{0,1\}^m) \setminus \emptyset) \to [0,1]$ be a function expressing the *satisfaction* of agent i towards the outcome of a rule F on profile $\mathbf{\Gamma}$. To simplify, we write $sat(i, F(\mathbf{\Gamma}))$ instead of $sat(\gamma_i, F(\mathbf{\Gamma}))$. The preference of agent i over outcomes is then defined as a complete and transitive relation $\preccurlyeq_i$, whose strict part is $\prec_i$, such that $F(\mathbf{\Gamma}) \preccurlyeq_i F(\mathbf{\Gamma}')$ *iff* $sat(i, F(\mathbf{\Gamma})) \geq sat(i, F(\mathbf{\Gamma}'))$.

For $\mathbf{\Gamma} = (\gamma_i)_{i \in \mathcal{N}}$, let $(\mathbf{\Gamma}_{-i}, \gamma_i') = \mathbf{\Gamma}' = (\gamma_1, \dots, \gamma_i', \dots, \gamma_n)$ be the profile where only agent i changed her goal from γ_i to γ_i'. Agent i has an *incentive* to manipulate by submitting goal γ_i' in place of goal γ_i if and only if $F(\mathbf{\Gamma}_{-i}, \gamma_i') \prec_i F(\mathbf{\Gamma})$. A rule F is *strategy-proof* if and only if for all profiles $\mathbf{\Gamma}$ there is no agent i who has an incentive to manipulate.

Everaere et al. [9] propose three manipulation strategies an agent i may perform depending on how much they are allowed to deviate from their truthful goal: *unrestricted* when i can send any γ_i' instead of her truthful γ_i, *erosion* when i can only send a γ_i' such that $\text{Mod}(\gamma_i') \subseteq \text{Mod}(\gamma_i)$ and *dilatation* when i can send only a γ_i' such that $\text{Mod}(\gamma_i) \subseteq \text{Mod}(\gamma_i')$. If a rule can be manipulated by erosion or dilatation it is manipulable in the general case, while if it is strategy-proof for unrestricted manipulation it is also strategy-proof for erosion and dilatation.

The issue-wise majority rule is known to be strategy-proof in the context of judgment aggregation [3], while this is not true anymore when moving to propositional goals:

Theorem 3.1. *EMaj, TrueMaj and 2sMaj can be manipulated by both erosion and dilatation.*

Since strategy-proofness cannot be guaranteed in general, we study the manipulability of the proposed rules when the agents' goals are restricted to conjunctions (corresponding to the framework of judgment aggregation with abstentions [4, 5]), disjunctions or exclusive disjunctions. Our results are summarized in Table 1.

	$\mathcal{L}^{\wedge}$		$\mathcal{L}^{\vee}$		$\mathcal{L}^{\oplus}$	
	E	D	E	D	E	D
EMaj	SP	SP	M	SP	M	M
TrueMaj	SP	SP	M	SP	M	M
2sMaj	SP	SP	SP	SP	M	M

Table 1: *E stands for erosion, D for dilatation, SP for strategy-proof and M for manipulable.*

4 COMPLEXITY OF MANIPULATION

We also study how computationally difficult would it be for an agent to find a goal allowing them to get a better outcome for the rules *EMaj* and *2sMaj*. The formal definition of the Manip(F) problem is in line with analogous work in judgment aggregation [8]: the input is a profile $\mathbf{\Gamma}$ and an agent i, and the question is whether there is a γ_i' such that $F(\mathbf{\Gamma}_{-i}, \gamma_i') \prec_i F(\mathbf{\Gamma})$. Let PP, for Probabilistic Polynomial Time, be the class of problems that can be solved in nondeterministic polynomial time with acceptance condition that more than half of the computations accept. We show that:

Theorem 4.1. *Manip(2sMaj) and Manip(EMaj) are PP-hard.*

5 RELATED WORK AND CONCLUSIONS

The literature on combinatorial voting (see, e.g., the chapter by Lang and Xia [16]) provides solutions to tackle the combinatorial explosion entailed by the structure of the alternatives, such as voting sequentially over issues using tractable voting rules. The work of Lang [15] on voting in multi-issue domains with compactly represented preferences is the starting point of our considerations. Propositional goals are perhaps the simplest compact language for preferences, linked to the literature on social choice with dichotomous preferences [6, 7]. The framework of belief merging [13, 14], also studies the aggregation of propositional formulas, focusing on aggregators satisfying a set of desirable properties inspired from belief revision. Closely related work is the study of strategy-proofness in judgment aggregation [3, 8], where the input is a complete binary choice over all issues rather than a propositional goal, as well as in belief merging [9]. Manipulation of voting rules has been amply studied in voting theory, starting from the seminal result of Gibbard and Satterthwaite [10, 19] to more recent studies aimed at finding barriers to manipulation (see, e.g., the survey by Conitzer and Walsh [1]). Propositional goals in a strategic setting have been extensively studied in the literature on boolean games [12, 21]. Here, however, issues are not exclusively controlled by agents, a closer model being that of aggregation games [11].

In this paper we studied the strategic component of the recent framework of goal-based voting for three variants of issue-wise majority [18]. We find that *EMaj*, *TrueMaj* and *2sMaj* are manipulable, even for erosion and dilatation strategies. As positive results, an agent with a goal in the language of conjunctions cannot manipulate, and for disjunctions *TrueMaj* and *EMaj* are only manipulable by erosion. While not strategy-proof in general, these resolute majority rules are PP-hard for an agent to manipulate.

REFERENCES

[1] Vincent Conitzer and Toby Walsh. 2016. Barriers to Manipulation in Voting. In *Handbook of Computational Social Choice*, F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia (Eds.). Cambridge University Press, Chapter 6.

[2] Franz Dietrich and Christian List. 2007. Judgment aggregation by quota rules: Majority voting generalized. *Journal of Theoretical Politics* 19, 4 (2007), 391–424.

[3] Franz Dietrich and Christian List. 2007. Strategy-proof judgment aggregation. *Economics & Philosophy* 23, 3 (2007), 269–300.

[4] Franz Dietrich and Christian List. 2008. Judgment Aggregation Without Full Rationality. *Social Choice and Welfare* 31, 1 (2008), 15–39.

[5] Elad Dokow and Ron Holzman. 2010. Aggregation of Binary Evaluations with Abstentions. *Journal of Economic Theory* 145, 2 (2010), 544–561.

[6] Edith Elkind and Martin Lackner. 2015. Structure in Dichotomous Preferences. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*.

[7] Edith Elkind, Martin Lackner, and Dominik Peters. 2017. Structured Preferences. In *Trends in Computational Social Choice*, Ulle Endriss (Ed.). AI Access, Chapter 10, 187–207.

[8] Ulle Endriss, Umberto Grandi, and Daniele Porello. 2012. Complexity of Judgment Aggregation. *J. of Artificial Intelligence Research* 45 (2012), 481–514.

[9] Patricia Everaere, Sébastien Konieczny, and Pierre Marquis. 2007. The Strategy-Proofness Landscape of Merging. *Journal of Artificial Intelligence Research* 28 (2007), 49–105.

[10] Allan Gibbard. 1973. Manipulation of voting schemes: a general result. *Econometrica* (1973), 587–601.

[11] Umberto Grandi, Davide Grossi, and Paolo Turrini. 2015. Equilibrium Refinement through Negotiation in Binary Voting.. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*. 540–546.

[12] Paul Harrenstein, Wiebe van der Hoek, John-Jules Ch. Meyer, and Cees Witteveen. 2001. Boolean games. In *Proceeding of the 8th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*.

[13] Sébastien Konieczny and Ramón Pino Pérez. 2002. Merging Information Under Constraints: A Logical Framework. *Journal of Logic and Computation* 12, 5 (2002), 773–808.

[14] Sébastien Konieczny and Ramón Pino Pérez. 2011. Logic Based Merging. *Journal of Philosophical Logic* 40, 2 (2011), 239–270.

[15] Jérôme Lang. 2004. Logical Preference Representation and Combinatorial Vote. *Annals of Mathematics and Artificial Intelligence* 42, 1-3 (2004), 37—71.

[16] Jérôme Lang and Lirong Xia. 2016. Voting in Combinatorial Domains. In *Handbook of Computational Social Choice*, F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia (Eds.). Cambridge University Press, Chapter 9.

[17] Kenneth O May. 1952. A set of independent necessary and sufficient conditions for simple majority decision. *Econometrica: Journal of the Econometric Society* (1952), 680–684.

[18] Arianna Novaro, Umberto Grandi, Dominique Longin, and Emiliano Lorini. 2018. Goal-based Voting with Propositional Goals. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*.

[19] Mark Allen Satterthwaite. 1975. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory* 10, 2 (1975), 187–217.

[20] Michael Wooldridge. 2009. *An introduction to multiagent systems*. John Wiley and Sons.

[21] Michael Wooldridge, Ulle Endriss, Sarit Kraus, and Jérôme Lang. 2013. Incentive Engineering in Boolean Games. *Artificial Intelligence* 195 (2013), 418–439.