From Individual Goals to Collective Decisions

Extended Abstract

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

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

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

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

ABSTRACT

We introduce the problem of aggregating the individual goals of a group of agents to find a collective decision. Goals are represented by propositional formulas on a finite set of binary issues. We define some rules for carrying out the aggregation of goals and we show how to adapt axiomatic properties from the literature on Social Choice Theory to this setting. The type of problems we are interested in studying for our rules are axiomatic characterizations, as well as the computational complexity of computing the outcome.

KEYWORDS

Social Choice Theory; Judgment Aggregation; Preference Representation; Computational Complexity

ACM Reference Format:

Arianna Novaro, Umberto Grandi, Dominique Longin, and Emiliano Lorini. 2018. From Individual Goals to Collective Decisions. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), Stockholm, Sweden, July 10–15, 2018,* IFAAMAS, 3 pages.

1 INTRODUCTION

Social choice theory can be used in contexts where agents need to take collective decisions over a combinatorial space of alternatives [3, 23, 24]. Thus, many researchers have introduced compact languages for preference representation [5, 15, 22] to be used as input for procedures finding a collective choice (see, e.g., the survey by Lang and Xia [18]).

In particular, in judgment aggregation [9, 12] agents use an aggregation procedure to take a decision over a set of binary issues, based on their individual opinions. Numerous aggregators have been proposed in the literature to tackle the problem of obtaining a collective outcome that is consistent with a given integrity constraint (see, e.g., [10, 16, 20]). Applications of judgment aggregation have been found in multi-agent argumentation [1, 4] and in the collective annotation of linguistic corpora [21].

However, since agents are asked to give complete opinions over the issues, there are limitations on what they can express. For instance, if some agent wishes to express that they want to accept *exactly one* out of three issues, they would be forced to choose between submitting either ballot (100), (010) or (001), thus not fully representing their original more complex opinion.

Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, E. André, S. Koenig (eds.), July 10–15, 2018, Stockholm, Sweden.

In the model we propose here, agents can express their individual goals as propositional formulas and then an option is collectively chosen. Propositional goals are a compact way to represent dichotomous preferences over alternatives described by binary variables. Social choice with dichotomous preferences has been widely studied as a possible solution to the computational barriers affecting classical preference aggregation (see, e.g., the survey by Elkind et al. [8]). However, to the best of our knowledge, it has not been applied to combinatorial domains such as those of goal aggregation.

We define goal aggregation rules inspired from the literature in judgment aggregation [7] and logic-based belief merging [13, 14]. Since in the literature on social choice theory aggregation rules are typically studied with the axiomatic method, we define an initial set of axioms that will serve to study characterization results for our rules. From a computational perspective, we want to study the chosen rules by determining the complexity of computing the outcome, a problem known in the literature as *winner determination*.

We provide a characterization for one of the generalizations of issue-wise majority rule, and we establish complexity bounds for determining the outcome for some of our rules.

2 FRAMEWORK

A set $\mathcal{N} = \{1, \dots, n\}$ of *agents* has to take a collective decision over a set of propositional variables $\mathcal{I} = \{1, \dots, m\}$ representing the *issues* at stake. The *individual goal* γ_i of agent i is a consistent propositional formula of language $\mathcal{L}_\mathcal{I}$ over the atoms in $\mathcal{I}$ and closed under standard propositional connectives. An *interpretation* is a function $v : \mathcal{I} \to \{0, 1\}$ associating to each variable a binary value. We also interpret v as the binary vector $(v(1), \dots, v(m))$. The set $\text{Mod}(\varphi) = \{v \mid v \models \varphi\}$ consists of all the models of formula φ.

We let $v_i(j) = (m^1_{ij}, m^0_{ij})$ indicate all the choices of agent i for issue j in the different models of her goal, where $m^x_{ij} = |\{v \in \text{Mod}(\gamma_i) \mid v(j) = x\}|$ for $x \in \{0, 1\}$. The issues on which the agents have to take a decision are logically independent from one another (i.e., we assume no integrity constraints). By collecting for each agent i her individual goal γ_i we obtain a *goal-profile* $\boldsymbol{\Gamma} = (\gamma_1, \dots, \gamma_n)$. A *goal-based voting rule* $F : (\mathcal{L}_\mathcal{I})^n \to \mathcal{P}(\{0, 1\}^m) \setminus \emptyset$ is a function taking as input n consistent goal formulas submitted by the agents and returning a set of models as a collective outcome. In case the output set is not a singleton we call the rule *irresolute*. For resolute rules, we indicate with $F(\boldsymbol{\Gamma})_j = \{1\}$ (respectively, $\{0\}$) the acceptance (respectively, rejection) of issue $j \in \mathcal{I}$ by rule F.

2.1 Aggregation Rules

The *conjunction rules* return all the models in common among the individual goals, and a chosen default model if there are none.

$$Conj_v(\Gamma) = \begin{cases} \text{Mod}(\gamma_1 \wedge \cdots \wedge \gamma_n) & \text{if non-empty} \\ \{v\} \text{ for } v \in \{0,1\}^m & \text{otherwise} \end{cases}$$

The *threshold rules* are comparable to quota rules in judgment aggregation [7]. Let $\mu_\varphi : \text{Mod}(\varphi) \to \mathbb{R}$ be a function associating to each model v of φ some weight $\mu_\varphi(v)$. Then:

$$TrSh^\mu(\Gamma)_j = \{1\} \text{ iff } \Big(\sum_{i \in \mathcal{N}} (w_i \cdot \sum_{v \in \text{Mod}(\gamma_i)} v(j) \cdot \mu_{\gamma_i}(v))\Big) \geq q_j$$

such that $0 \leq q_j \leq n+1$ for all $j \in \mathcal{I}$ is the quota of issue j, where for each $v \in \text{Mod}(\gamma_i)$ we have $\mu_{\gamma_i}(v) \neq 0$ and $w_i \in [0,1]$ is the individual weight of agent i. For readability, we omit vector $(q_1, \ldots, q_m)$, with the quotas for the issues, from the name of $TrSh^\mu$.

We call *EQuota* rule the $TrSh^\mu$ rules having $\mu_{\gamma_i}(v) = \frac{1}{|\text{Mod}(\gamma_i)|}$ and $w_i = 1$ for all $v \in \text{Mod}(\gamma_i)$ and $i \in \mathcal{N}$. Since $\sum_{v \in \text{Mod}(\gamma_i)} \mu_{\gamma_i}(v) = 1$ for all $i \in \mathcal{N}$, the number of models of a goal formula is irrelevant and agents have an equal impact on the outcome.

Finally, we introduce three alternative definitions of the majority rule. The first one is the *EMaj* (resolute) rule, which is an *EQuota* rule having $q_j = \lceil \frac{n}{2} \rceil$ for all $j \in \mathcal{I}$. The second version of majority is irresolute and it compares for each issue the number of acceptances with the number of rejections, weighting each model of an individual goal in the same way as in *EQuota* rules:

$$TrueMaj(\Gamma) = \Pi_{j \in \mathcal{I}} M(\Gamma)_j$$

where for all $j \in \mathcal{I}$

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

Thus, for each issue $j \in \mathcal{I}$ we compute the value of $M(\Gamma)_j$ by setting it to $\{1\}$ (respectively, $\{0\}$) if strictly more than half of the agents accept (respectively, reject) issue j, and to $\{0,1\}$ if exactly half of the agents accept/reject. The third version of majority is *2sMaj*, defined as *Maj*(*Maj*(γ_1), ..., *Maj*(γ_n)), where the strict majority rule is applied first to the models of each individual goal and then again to the result obtained in the first step.

3 AXIOMATIC CHARACTERIZATION

The *axiomatic method* in social choice theory evaluates aggregation rules by first defining some general properties (axioms) and then proving whether or not aggregation rules satisfy them. We define adaptations of known axioms for rules aggregating goals.

(A) An *anonymous* aggregation rule F is such that for any profile Γ and any permutation $\sigma : \mathcal{N} \to \mathcal{N}$, we have that $F(\gamma_1, \ldots, \gamma_n) = F(\gamma_{\sigma(1)}, \ldots, \gamma_{\sigma(n)})$.

(I) An *independent* aggregation rule F is such that for any two profiles Γ and Γ', for all $j \in \mathcal{I}$ and for all $i \in \mathcal{N}$, we have that $v_i(j) = v'_i(j)$ implies $F(\Gamma)_j = F(\Gamma')_j$.

(N) A *neutral* aggregation rule F is such that for all profiles Γ, for all two issues $j, k \in \mathcal{I}$, and for all agents $i \in \mathcal{N}$ we have that $v_i(j) = v_i(k)$ implies $F(\Gamma)_j = F(\Gamma)_k$.

(U) A *unanimous* aggregation rule F is such that for all profiles Γ and for all $j \in \mathcal{I}$, if $m_{ij}^x = 0$ for all $i \in \mathcal{N}$ then $F(\Gamma)_j = \{1-x\}$.

(PR) Profiles Γ and Γ' are *comparable* if and only if for all $i \in \mathcal{N}$ we have that $|\text{Mod}(\gamma_i)| = |\text{Mod}(\gamma'_i)|$. An aggregation rule satisfies *positive responsiveness* if for all comparable profiles Γ and $\Gamma' = (\gamma_1, \ldots, \gamma'_i, \ldots, \gamma_n)$, for all issues $j \in \mathcal{I}$ and for all $i \in \mathcal{N}$, if $m'^x_{ij} \geq m^x_{ij}$ for $x \in \{0,1\}$, then $[F(\Gamma)_j = \{x\}$ or $F(\Gamma)_j = \{0,1\}]$ implies $F(\Gamma')_j = \{x\}$.

(E) An aggregation rule F is *egalitarian* if and only if for all profiles Γ, if we construct a profile Γ' with $|\mathcal{N}'| = \text{lcm}(|\text{Mod}(\gamma_1)|, \ldots, |\text{Mod}(\gamma_n)|)$, and for all $i \in \mathcal{N}$ and each $v \in \text{Mod}(\gamma_i)$ we have $\frac{|\mathcal{N}'|}{|\mathcal{N}| \cdot |\text{Mod}(\gamma_i)|}$ agents in $\mathcal{N}'$ voting v in Γ', then $F(\Gamma) = F(\Gamma')$.

(D) An aggregation rule satisfies *duality* when for all profiles Γ and for all issues $j \in \mathcal{I}$, if $F(\Gamma)_j = \{x\}$ then $F(\overline{\Gamma})_j = \{1-x\}$, where $\overline{\Gamma}$ is such that $\overline{v_i(j)} = (\overline{m^1_{ij}}, \overline{m^0_{ij}}) = (m^0_{ij}, m^1_{ij})$ for all $j \in \mathcal{I}$ and $i \in \mathcal{N}$.

Following the seminal result of May [19], where an axiomatization of the majority rule in the context of voting over two alternatives is provided, we also axiomatically characterize *TrueMaj*, the most intuitive among our proposed generalizations of majority.

THEOREM 3.1. *For arbitrary $\mathcal{N}$ and $\mathcal{I}$, a goal-aggregation rule satisfies* (E), (I), (A), (N), (PR), (U) *and* (D) *if and only if it is TrueMaj.*

4 COMPUTATIONAL COMPLEXITY

The *winner determination* problem asks how hard it is to compute the outcome of aggregation rules [2, 6, 11, 17]. For resolute rules we define it as follows (and analogously for irresolute rules):

WINDET(F)
Input: profile Γ, issue j
Question: Is it the case that $F(\Gamma)_j = \{1\}$?

For a special case of $TrSh^\mu$ we get completeness for the class NP.

THEOREM 4.1. *WINDET($TrSh^\mu$) is NP-complete, for $\mu_{\gamma_i}(v) = 1$ constant and $w_i = 1$ for all $i \in \mathcal{N}$.*

Let PP be the complexity class Probabilistic Polynomial Time. Let *TrueMaj** be a resolute version of *TrueMaj* that in case of equal support for issue j returns 0 in the outcome.

THEOREM 4.2. *Problems WINDET(EMaj), WINDET(TrueMaj*) and WINDET(2sMaj) are PP-hard.*

Since Γ contains formulas, if a rule has to manipulate the models to compute the outcome some form of satisfiability is needed — thus starting from the complexity class NP. Asking agents to provide the models of their goals would lower the complexity of some results, but the input would become demanding for the agents and also of exponential size in the number of issues in the worst case.

5 CONCLUSIONS

We presented a framework to handle the aggregation of individual goals in a multi-agent setting. We defined a number of procedures taking as input the consistent goal formulas of the agents and returning a collective decision in the form of a set of valuations for the issues at stake. We introduced three alternative definitions of the majority rule and characterized one of them axiomatically. We also studied computationally the problem of determining the outcome for some of our proposed rules.

REFERENCES

[1] Edmond Awad, Richard Booth, Fernando Tohmé, and Iyad Rahwan. 2017. Judgement Aggregation in Multi-Agent Argumentation. *Journal of Logic and Computation* 27, 1 (2017), 227—259.

[2] Dorothea Baumeister, Jörg Rothe, and Ann-Kathrin Selker. 2017. Strategic Behavior in Judgment Aggregation. In *Trends in Computational Social Choice*, Ulle Endriss (Ed.). AI Access, Chapter 8, 145—168.

[3] Felix Brandt, Vincent Conitzer, Ulle Endriss, Ariel D Procaccia, and Jérôme Lang. 2016. *Handbook of Computational Social Choice*. Cambridge University Press.

[4] Martin Caminada and Gabriella Pigozzi. 2011. On Judgment Aggregation in Abstract Argumentation. *Autonomous Agents and Multi-Agent Systems* 22, 1 (2011), 64—102.

[5] Yann Chevaleyre, Ulle Endriss, Jérôme Lang, and Nicolas Maudet. 2008. Preference Handling in Combinatorial Domains: From AI to Social Choice. *AI Magazine* 29, 4 (2008), 37—46.

[6] Ronald de Haan and Marija Slavkovik. 2017. Complexity Results for Aggregating Judgments using Scoring or Distance-Based Procedures. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*.

[7] Franz Dietrich and Christian List. 2007. Judgment Aggregation by Quota Rules: Majority Voting Generalized. *Journal of Theoretical Politics* 19, 4 (2007), 391—424.

[8] 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.

[9] Ulle Endriss. 2016. Judgment Aggregation. In *Handbook of Computational Social Choice*, F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia (Eds.). Cambridge University Press, Chapter 17, 399—426.

[10] Ulle Endriss and Umberto Grandi. 2014. Binary Aggregation by Selection of the Most Representative Voters. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*.

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

[12] Davide Grossi and Gabriella Pigozzi. 2014. *Judgment Aggregation: A Primer*. Morgan and Claypool.

[13] Sébastien Konieczny, Jérôme Lang, and Pierre Marquis. 2004. DA2 Merging Operators. *Artificial Intelligence* 157, 1-2 (2004), 49—79.

[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, Gabriella Pigozzi, Marija Slavkovik, Leon van der Torre, and Srdjan Vesic. 2017. A partial taxonomy of judgment aggregation rules and their properties. *Social Choice and Welfare* 48, 2 (2017), 327—356.

[17] Jérôme Lang and Marija Slavkovik. 2014. How Hard is it to Compute Majority-Preserving Judgment Aggregation Rules?. In *Proceedings of the 21st European Conference on Artificial Intelligence (ECAI)*.

[18] Jérôme Lang and Lirong Xia. 2016. Voting in Combinatorial Domains. In *Handbook of Computational Social Choice*, Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel Procaccia (Eds.). Cambridge University Press, Chapter 9, 197—222.

[19] 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.

[20] Gabriella Pigozzi. 2005. Two Aggregation Paradoxes in Social Decision Making: the Ostrogorski Paradox and the Discursive Dilemma. *Episteme: A Journal of Social Epistemology* 2, 2 (2005), 33—42.

[21] Ciyang Qing, Ulle Endriss, Raquel Fernández, and Justin Kruger. 2014. Empirical Analysis of Aggregation Methods for Collective Annotation. In *Proceedings of the 25th International Conference on Computational Linguistics (COLING)*.

[22] Francesca Rossi, Kristen Brent Venable, and Toby Walsh. 2011. *A Short Introduction to Preferences: Between Artificial Intelligence and Social Choice*. Morgan & Claypool Publishers.

[23] Yoav Shoham and Kevin Leyton-Brown. 2009. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press.

[24] Michael Wooldridge. 2009. *An Introduction to Multiagent Systems* (second ed.). John Wiley & Sons.