Multi-issue Voting with Propositional Goals

Doctoral Consortium

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

ACM Reference Format:
Arianna Novaro. 2018. Multi-issue Voting with Propositional Goals. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), Stockholm, Sweden, July 10–15, 2018,* IFAAMAS, 2 pages.

1 BACKGROUND

I present the motivating example of my thesis, as well as the relevant background literatures on judgment aggregation, belief merging, preference representation and social networks that constitute the starting point of my study.

Example. *Let Ann, Barbara and Camille be three friends who want to spend a day visiting a seaside city. They narrowed down the options to* (i) *visit the lighthouse,* (ii) *have a picnic on the beach, and* (iii) *go to the acquarium. Ann wants to do everything, Barbara prefers to just relax on the beach, and Camille wants to do a single activity but she does not care about which one. How should the situation be modeled?*

Judgment Aggregation. The most natural choice for modelling a scenario such as the one presented in the motivating example is to use judgment aggregation, an area of computational social choice where agents have to deliberate over a set of interconnected issues in order to take a collective decision [2, 12]. The domain of each issue is binary, meaning that agents can either accept it or reject it. In particular, one of the frameworks of judgment aggregation is *binary aggregation with integrity constraints*, where the opinions of the agents are modelled via a vector of binary values, and the dependencies among the issues are expressed by a propositional formula [5]. Inspired by voting theory, many aggregation procedures have been defined and studied thoroughly in the literature from an axiomatic, computational and strategic viewpoint [3, 11]. There are some subtleties according to which judgment aggregation falls short as a model for examples like the one proposed above. First of all, judgment aggregation originated as a model to aggregate individual opinions of agents in order to get a collective decision, while here we want to aggregate individual goals to obtain a collective plan. Secondly, there is no need to introduce integrity constraints since the issues are independent from one another. Last, and most importantly, by being forced to submit a unique vector of binary values, goals such as the one expressed by Camille cannot be captured.

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.

Belief Merging. Another framework that we could think of using for modeling our example is that of belief merging, inspired from the work in belief revision. In this case each agent is assumed to have a set of propositional formulas capturing her beliefs (i.e., a *belief base*), and different aggregation procedures are used to compute a collective belief [9]. As it can be observed, the frameworks of belief merging and judgment aggregation are indeed quite close [4]. The central idea of using propositional formulas to model mental attitudes could be applied to the goals of the agents in our example, so that even more elaborated goals as Camille's one could be expressed. Nevertheless some aspects of belief merging may not make it the best candidate to solve our problem. In the first place, the merging operators used in the belief merging literature are essentially non-independent (since it is often the case that aggregation takes place in the presence of integrity constraints). Moreover, since we are looking for a definite plan for the agents telling them what to do, we want to avoid irresoluteness (i.e., the presence of tied outcomes) as much as possible, something that is in general not guaranteed by belief merging operators.

Preference Representation. In the proposed example, agents have to express their goals regarding issues that have a binary domain. More generally, the problem of voting on combinatorial domains has generated some study on how to represent the preferences of the agents [10]. By choosing to represent the goals of the agents as a propositional formula, the models of that fomula would naturally become the preferred outcomes as the final choice for a plan, while each countermodel would be equally disliked. Therefore, we are drawn to model this situation with *dicothomous preferences*. Depending on the nature of the issues at stake and how agents value them, another option would be to endow agents with preferences based on the Hamming distance.

Social Networks. Observe that one assumption in the motivating example is that we are modelling a group of friends having to take a decision together. This detail leads us to think that we might want to capture the existing relationship among the agents via a social network. In recent years, some work has focussed on modelling the formation of opinions in groups of agents that are connected in a social network (modelled by a graph whose nodes are the agents) and who update their currently held opinion according to their neighbours' opinions [7].

2 CURRENT RESEARCH

In this section I present the current state of the investigations we performed in the research areas introduced above.

Goal-based voting [13]. Let $\mathcal{N} = \{1, \ldots, n\}$ be a set of *agents*, who have to decide over a set $\mathcal{I} = \{1, \ldots, m\}$ of propositional variables representing *issues*. The domain of each issue $j \in \mathcal{I}$ is the set $\{0, 1\}$, where 0 denotes rejection and 1 denotes acceptance. We assume the issues to be independent from one another, meaning there is no integrity constraint.

Let $\mathcal{L}_\mathcal{I}$ be a propositional language whose atoms are the variables in $\mathcal{I}$ and that is closed under standard propositional connectives. The individual goal of each agent $i \in \mathcal{N}$ is a propositional formula γ_i of $\mathcal{L}_\mathcal{I}$. For instance, we have that in our motivating example the goals of Ann, Barbara and Camille are $\gamma_1 = 1 \wedge 2 \wedge 3$, $\gamma_2 = \neg 1 \wedge 2 \wedge \neg 3$ and $\gamma_3 = ((1 \oplus 2) \oplus 3) \wedge \neg(1 \wedge 2 \wedge 3)$ respectively, where $\oplus$ represents the exclusive disjunction.

A *goal-profile* is a vector $\mathbf{\Gamma} = (\gamma_1, \ldots, \gamma_n)$ collecting the individual goals of the agents. In order to obtain a collective plan for the agents, we use a *goal-based voting rule* defined as $F : (\mathcal{L}_\mathcal{I})^n \rightarrow \mathcal{P}(\{0, 1\}^m) \setminus \emptyset$. A rule f thus takes as input a vector of propositional formulas (capturing the agents' goals) and it outputs a set of models (capturing the chosen plans). If for every profile the output set is always a singleton, we call the rule *resolute*, and we call it *irresolute* otherwise.

Once the framework of goal-based voting has been established, we investigated multiple aspects of this research agenda that I only briefly list here due to space constraints.

- We **defined multiple goal-based voting rules**, trying to favour resoluteness and an independent treatment of the issues (given the absence of integrity constraints). In particular, we proposed multiple generalisation for the well-known majority rule from its definition in judgment aggregation.
- We studied the defined rules from an **axiomatic perspective**, i.e., investigating which mathematical properties they satisfy and providing a characterisation for one of them.
- We studied the **computational complexity** of computing the outcome for the proposed rules, a problem that is known in the judgment aggregation literature as the *winner determination* problem.

Strategic Aspects in Agents Networks [1, 6]. We are interested in situations where agents may know one another and be influenced when forming their opinions. Therefore, we explored the strategic aspect of a simple model of social influence among agents. The connections among the agents are modelled via a directed graph, where an edge pointing to an agent means that the first influences the second. Agents express binary views on a set of issues, and they iteratively update them by taking into account the expressed opinions of their influencers. They can also strategically choose to actively influence others or not.

From a more theoretical perspective, we studied in general strategic situations where agents are endowed with (temporal) goals. Starting from the literature on *iterated boolean games* [8] we introduced structures where propositional atoms can be controlled by multiple agents. We showed that these structures can account for a variety of strategic scenarios, such as the ones modelling social influence described above, and that (quite surprisingly) they can be reduced to structures of exclusive control.

3 FUTURE WORK

In the rest of my thesis, I want to delve deeper into the themes that I have shortly presented in this abstract. More precisely, I want to continue developing the framework of goal-based voting by exploring its rules and axiomatics, by studying the strategical aspects of agents who may want to lie to get a preferable outcome, and by bridging the two lines of research we worked on (i.e., agents may have their goals shaped by those of their influencers on a social network).

REFERENCES

[1] Francesco Belardinelli, Umberto Grandi, Andreas Herzig, Dominique Longin, Emiliano Lorini, Arianna Novaro, and Laurent Perrussel. 2017. Relaxing Exclusive Control in Boolean Games. In *Proceedings of the 16th conference on Theoretical Aspects of Rationality and Knowledge (TARK)*.

[2] 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.

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

[4] Patricia Everaere, Sébastien Konieczny, and Pierre Marquis. 2015. Belief merging versus judgment aggregation. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

[5] Umberto Grandi and Ulle Endriss. 2011. Binary aggregation with Integrity Constraints. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*.

[6] Umberto Grandi, Emiliano Lorini, Arianna Novaro, and Laurent Perrussel. 2017. Strategic disclosure of opinions on a social network. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. 1196–1204.

[7] Umberto Grandi, Emiliano Lorini, and Laurent Perrussel. 2015. Propositional Opinion Diffusion. In *Proceedings of the 14th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

[8] Julian Gutierrez, Paul Harrenstein, and Michael Wooldridge. 2015. Iterated boolean games. *Information and Computation* 242 (2015), 53–79.

[9] Sébastien Konieczny and Ramón Pino Pérez. 1999. Merging with integrity constraints. In *European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty*. 233—244.

[10] Jérôme Lang. 2004. Logical preference representation and combinatorial vote. *Annals of Mathematics and Artificial Intelligence* 42, 1 (2004), 37—71.

[11] 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)*.

[12] Christian List and Philip Pettit. 2002. Aggregating sets of judgments: An impossibility result. *Economics & Philosophy* 18, 1 (2002), 89–110.

[13] Arianna Novaro, Umberto Grandi, Dominique Longin, and Emiliano Lorini. 2018. Goal-based Collective Decisions: Axiomatics and Computational Complexity. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence and the 23rd European Conference on Artificial Intelligence (IJCAI-ECAI 2018)*.