Goal-Based Voting

Umberto Grandi

Institut de Recherche en Informatique de Toulouse (IRIT)
University of Toulouse 1 Capitole

11 April 2019

[Joint work with Arianna Novaro, Dominique Longin, Emiliano Lorini - IRIT]

Motivational Example

- A city trip for a group of friends: Ann, Barbara, and Camille
- Include a visit to a Church, to a Museum, to a Park?
- Goal of Ann: see all the points of interest
- Goal of Barbara: have a walk in the Park
- Goal of Camille: visit a single point of interest, does not care which one

Motivational Example

- A city trip for a group of friends: Ann, Barbara, and Camille
- Include a visit to a Church, to a Museum, to a Park?
- Goal of Ann: see all the points of interest
- Goal of Barbara: have a walk in the Park
- Goal of Camille: visit a single point of interest, does not care which one

Suppose we ask each question separately:

	Church	Museum	Park
Ann	☑	☑	☑
Barbara	☒	☒	☑
Camille	☒	☑	☒
Majority	☒	☑	☑

Camille is not happy! And she could have voted differently...

Motivational Example

- A city trip for a group of friends: Ann, Barbara, and Camille
- Include a visit to a Church, to a Museum, and to a Park?
- Goal Ann: see all the points of interest
- Goal Barbara: have a walk in the Park
- Goal Camille: visit a single point of interest, does not care about which one

Suppose we ask each question separately:

	Church	Museum	Park
Ann	☑	☑	☑
Barbara	☒	☒	☑
Camille	☒	☒	☒
Majority	☒	☒	☑

Camille is now happy! But she voted against her goal...

From aggregation games...

One possibility is to study aggregation games:

- Vote on multiple independent issues
- Collective decision by majority (or analogous issue-by-issue rule)
- Propositional goal for each the agents
- Dominant strategies? Equilibria?

Initial results in the first sections of recent work:

U. Grandi, D. Grossi, P. Turrini. Negotiable Votes. In *Journal of Artificial Intelligence Research (JAIR)*, 64:895–929. 2019

...to goal-based voting

We can otherwise collect the full goals from voters (using a compact propositional language) and directly aggregate using goals!

Three automated personal assistants arrange a business meal for their owners:

- Fancy restaurant (F), in the center (C), for lunch (L) or for dinner?
- $\gamma_1 = \neg C \to (\neg F \wedge L)$
- $\gamma_2 = \neg C \oplus \neg F$
- $\gamma_3 = F \wedge \neg L \wedge C$

The result should be a plan for the business meal: a decision on F, C and L

Starting questions:

1. Focus on (almost) resolute rules
2. Can manipulation be avoided at this level?
3. Relation with belief merging and judgment aggregation with abstentions?

Outline

1. Basic definitions of goal-based voting
2. Three versions of majoritarian rules
3. Axiomatic characterisations
4. Strategic voting
5. Computational complexity

Goal-based voting

Basic ingredients:

- $\mathcal{N}$ agents have to decide on m binary issues in $\mathcal{I}$
- Each agent expresses an individual goal as a propositional formula
- This forms a goal profile $\mathbf{\Gamma} = (\gamma_1, \ldots, \gamma_n)$

In our initial example:

- 3 voters, 3 issues {F,C,L}
- goal profile $\mathbf{\Gamma} = (\neg C \to (\neg F \land L), \neg C \oplus \neg F, F \land \neg L \land C)$

Each goal induces a set of models: Mod$(\gamma_1) = \{001, 010, 110, 011, 110\}$
(ie if the restaurant is far then it should not be fancy and it should be for lunch)

Goal-based voting rules

Definition

A goal-based voting rule is a collection of functions for all n and m:

$$F : (\mathcal{L}_\mathcal{I})^n \to \mathcal{P}(\{0,1\}^m) \setminus \emptyset$$

where $\mathcal{L}_\mathcal{I}$ is the propositional language over issues $\mathcal{I}$.

Two examples that are not resolute enough:

Conjunction rule: output $\bigwedge_i \gamma_i$ if consistent, otherwise a default model

Approval voting: output those models that satisfy the highest number of goals: $Approval(\mathbf{\Gamma}) = \arg\max_{v \in \mathsf{Mod}(\bigvee_{i \in \mathcal{N}} \gamma_i)} |\{i \in N \mid v \in \mathsf{Mod}(\gamma_i)\}|$.

Strongly resolute rules

A rule is strongly resolute if $|F(\mathbf{\Gamma})| = 1$ in all profiles $\mathbf{\Gamma}$.

There is a general tension between resoluteness and fairness notions:

Anonymity: permuting the goals in the input should not change the output

Duality: if $\overline{\mathbf{\Gamma}}$ is obtained from profile $\mathbf{\Gamma}$ by replacing every issue with its negation, then $F(\overline{\mathbf{\Gamma}}) = (1 - F(\mathbf{\Gamma})_1, \ldots, 1 - F(\mathbf{\Gamma})_m)$

Theorem

There is no strongly resolute rule F satisfying both anonymity and duality.

Proof idea. Consider a symmetric profile such as $\gamma_1 = 1 \land \neg 2$ and $\gamma_2 = \neg 1 \land 2$.

Strongly resolute rules

Two adaptations of the majority rule (biased towards 0) can be proposed:

EMaj: a quota rule accepting an issue if and only if more than half of the discounted total votes are in its favor

$$\mathit{EMaj}(\mathbf{\Gamma})_j = 1 \;\; \mathit{iff} \;\; \sum_{i \in \mathcal{N}} \Big(\sum_{v \in \mathrm{Mod}(\gamma_i)} \frac{v(j)}{|\mathrm{Mod}(\gamma_i)|} \Big) \geq \left\lceil \frac{n+1}{2} \right\rceil$$

2-step Majority: first apply *Maj* to the models of the agents' goals, and then to the result of the first step

$$\mathit{2sMaj}(\mathbf{\Gamma}) = \mathit{Maj}(\mathit{Maj}(\mathrm{Mod}(\gamma_1)), \ldots, \mathit{Maj}(\mathrm{Mod}(\gamma_n)))$$

Weakly Resolute Rules

A rule is weakly resolute if on every profile $\mathbf{\Gamma}$,
$F(\mathbf{\Gamma}) = \text{Mod}(\varphi)$ for some conjunction $\varphi \in \mathcal{L}^{\wedge}$

A weakly resolute rule either accepts, rejects, or abstains on any of the issues.

Informally, an independent goal-based voting rule decides each issue independently from one another (similar to classic JA definition):

Theorem

Every independent goal-aggregation rule is weakly resolute.

Weakly Resolute Rules

A rule is weakly resolute if on every profile $\mathbf{\Gamma}$,
$F(\mathbf{\Gamma}) = \text{Mod}(\varphi)$ for some conjunction $\varphi \in \mathcal{L}^{\wedge}$

A weakly resolute rule either accepts, rejects, or abstains on any of the issues.

Informally, an independent goal-based voting rule decides each issue independently from one another (similar to classic JA definition):

Theorem

Every independent goal-aggregation rule is weakly resolute.

And here is a weak-resolute version of the majority rule:

TrueMaj: compare issue-by-issue the total discounted acceptances with the total discounted rejections, setting the result to 1 (respectively, 0) if higher (respectively, lower), and to both 0 and 1 if tied.

The three majoritarian rules differ

	$\mathbf{\Gamma}^1$	$\mathbf{\Gamma}^2$	$\mathbf{\Gamma}^3$
Mod(γ_1)	(111)	(111)	(000)
Mod(γ_2)	(001)	(111) (011) (000)	(111) (110) (011)
Mod(γ_3)	(100) (010) (101)	(111) (011) (000)	(111) (110) (011)
$EMaj$	(001)	—	(010)
$TrueMaj$	(101)	(111)	—
$2sMaj$	—	(011)	(111)

Axiomatic characterisation

Egalitarianism (E): for all profiles $\mathbf{\Gamma}$, if we construct a profile $\mathbf{\Gamma}'$ with $|\mathcal{N}'| = \text{lcm}(|\text{Mod}(\gamma_1)|, \dots, |\text{Mod}(\gamma_n)|)$, 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 $\mathbf{\Gamma}'$, then $F(\mathbf{\Gamma}) = F(\mathbf{\Gamma}')$.

Unanimity (U): ...

Positive Responsiveness (PR): ...

Neutrality (N): ...

Theorem

A rule satisfies (E), (I), (N), (A), (PR), (U) and (D) iff it is TrueMaj.

Proof idea for $\Rightarrow$. Using (E) we can reduce to a judgment aggregation profile in which goals have a single model, compute the result using the axioms, and show that it corresponds to *TrueMaj*.

Belief merging and judgment aggregation

Belief merging is a widely studied setting to combine propositional formulas into sets of models: exactly the same as goal-based voting!

- Some correspondences: Approval voting is called $\Delta_{\mu}^{\Sigma,d}$
- All belief merging rules are neither strongly nor weakly resolute:

Proposition

The belief merging axiom (IC2), i.e., $F(\mathbf{\Gamma}) = \bigwedge_i \gamma_i$ *if consistent, is incompatible with both strong and weak resoluteness.*

Proof idea. Take a profile where $\bigwedge_i \gamma_i = 1 \vee 2$.

Belief merging and judgment aggregation

Belief merging is a widely studied setting to combine propositional formulas into sets of models: exactly the same as goal-based voting!

- Some correspondences: Approval voting is called $\Delta_{\mu}^{\Sigma,d}$
- All belief merging rules are neither strongly nor weakly resolute:

Proposition

The belief merging axiom (IC2), i.e., $F(\mathbf{\Gamma}) = \bigwedge_i \gamma_i$ *if consistent, is incompatible with both strong and weak resoluteness.*

Proof idea. Take a profile where $\bigwedge_i \gamma_i = 1 \vee 2$.

Judgment aggregation with abstentions: special case of individual goals restricted to conjunctions. See strategic manipulation part.

P. Everaere, S. Konieczny and P. Marquis. An Introduction to Belief Merging and its Links with Judgment Aggregation. In *Trends in COMSOC*, 2017.

Z. Terzopoulou, U. Endriss, R. de Haan. Aggregating Incomplete Judgments: Axiomatisations for Scoring Rules. In *COMSOC-2018*.

Strategic voting

Agents can report a different goal to change the outcome in their favour:

Unrestricted: agents try any γ_i' instead of their γ

Dilatation: agents try any $\gamma' \supseteq \gamma$
Example: vote for unwanted combinations to block specific issues

Erosion: agents try any $\gamma' \subseteq \gamma$
Example: be excessively strict on Doodle, or the example below for agent 3

Strategic voting

Agents can report a different goal to change the outcome in their favour:

Unrestricted: agents try any γ_i' instead of their γ

Dilatation: agents try any $\gamma' \supseteq \gamma$
Example: vote for unwanted combinations to block specific issues

Erosion: agents try any $\gamma' \subseteq \gamma$
Example: be excessively strict on Doodle, or the example below for agent 3

	$\mathbf{\Gamma}$	$\mathbf{\Gamma}'$
Mod(γ_1)	(111)	(111)
Mod(γ_2)	(001)	(001)
Mod(γ_3)	(101) (010) (000)	(101)
EMaj	(001)	(101)
[True/2s]Maj	(001)	(101)

Strategic voting in goal-based voting

To have a more refined picture, we considered restricted goal languages: $\mathcal{L}^\star$ for $\star \in \{\wedge, \vee, \oplus\}$

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

Instructions:

- SP=strategy-proof, M=manipulable
- E=erosion, D=dilatation
- For weak resolute rules (*TrueMaj*) we extend preferences to sets of models

Computational Complexity

We study two classical decision problems (strongly resolute versions):

Winner determination: is $F(\mathbf{\Gamma})_j = 1$?

Manipulation: Does agent i have any incentive to manipulate?

Both problems are PP-hard for our rules. PP for probabilistic polynomial time: problems solvable by a deterministic Turing-machine that accepts in strictly more than half of the non-deterministic computations iff the answer is yes.

Theorem

Winner determination and manipulation are PP-hard for 2sMaj and EMaj.

Not a good news: PP problems are very hard (PP contains NP and co-NP)

Is PP pervasive in majoritarian reasoning with propositional goals? The problem of determining whether a majority of the models of a goal makes a proposition true or false is PP-complete...

Conclusions

- We aimed at defining resolute goal aggregation rules: from individual propositional goals to a single model or to models of a conjunction
- Three different versions of the majority rule: *TrueMaj* can be characterised axiomatically, but is weakly resolute
- Manipulation cannot be avoided, but when goals are conjunctive (JA with abstentions) or disjunctive the picture is nicer
- The computational complexity is high (is it avoidable?). Future work: study restrictions on the goal language to tame the complexity

If you want to know more:

A. Novaro, U. Grandi, D. Longin, E. Lorini. Goal-Based Collective Decisions: Axiomatics and Computational Complexity. In *Proceedings of IJCAI-2018*.