Collective Decision-Making with Goals

Arianna Novaro

PhD Thesis Defense

12th of November 2019

Supervised by

Umberto Grandi   Dominique Longin   Emiliano Lorini

The Research (Fields) Behind the Title

Collective Decision-Making with Goals

Multi-Agent Systems
Interactions of multiple agents acting towards a goal.

Computational Social Choice
Aggregation of preferences or opinions of a group of agents.

Game Theory
Strategic agents trying to maximize their utilities.

Logical Languages
To represent goals, agents and their interactions.

Challenges in Collective Decision-Making

Compact Input

Please input your preferences
over the 50 options as a linear order.

Strategic Behavior

The new vote of agent 5
changes the winner.

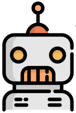

Easy Computation

Please wait 80 hours
while I calculate the result.

I found 9 equally good plans
satisfying your query.

Decisive Result

A Tale of Two Research Questions

1. How can we design aggregation procedures to help a group of agents having compactly expressed goals and preferences make a collective choice?

2. How can we model agents with conflicting goals who try to get a better outcome for themselves by acting strategically?

Presentation Roadmap

① Aggregation

1. Goal-based Voting
2. Aggregation of gCP-nets

② Strategic Behavior

3. Strategic Goal-based Voting
4. Strategic Disclosure of Opinions on a Social Network
5. Relaxing Exclusive Control in Boolean Games

Part I: Aggregation

Compact Languages | Goals and Preferences

Propositional Logic Goals

$$\varphi ::= p \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid \varphi_1 \vee \varphi_2$$

“$\texttt{fish} \wedge \texttt{white\_w}$”

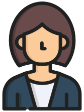

gCP-nets

$$\varphi := \psi : p_1 \triangleright p_2$$

“$\texttt{fish} \vee \texttt{chocolate} : \texttt{white\_w} \triangleright \texttt{red\_w}$”

Goal-based Voting | Framework

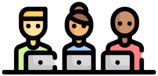

- n agents in $\mathcal{A}$ have to decide over m binary issues in $\mathcal{I}$
  - $\mathcal{A} = \{A, B, C\}$ and $\mathcal{I} = \{\texttt{morning}, \texttt{guest\_talks}, \texttt{lunch}\}$
- agent i's goal is prop. formula γ_i with models $\text{Mod}(\gamma_i)$
  - $\gamma_C = \texttt{guest\_talks} \land (\texttt{morning} \rightarrow \texttt{lunch})$
  - $\text{Mod}(\gamma_C) = \{(111), (011), (010)\}$
- a goal-profile $\mathbf{\Gamma} = (\gamma_1, \dots, \gamma_n)$ contains all agents' goals
- no integrity constraints

Novaro, Grandi, Longin, Lorini. *Goal-Based Collective Decisions: Axiomatics and Computational Complexity*. IJCAI-18.

Goal-based Voting | Rules

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

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

Approval: Return all interpretations satisfying the most goals.
Majority: . . . how to generalize to propositional goals?

agent i	Mod(γ_i)
A	(000)
B	(010) (100)
C	(111) (011) (010)

EMaj Majority with equal weights to models.

TrueMaj Majority with equal weights to models and fair treatment of ties.

2sMaj Majority done in two steps: on goals, and then on result of step one.

Goal-based Voting | Rules

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

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

Approval: Return all interpretations satisfying the most goals.
Majority: ... how to generalize to propositional goals?

agent i	Mod(γ_i)
A	(000)
B	(010) (100)
C	(111) (011) (010)

EMaj Majority with equal weights to models.

TrueMaj Majority with equal weights to models and fair treatment of ties.

2sMaj Majority done in two steps: on goals, and then on result of step one.

Goal-based Voting | Axioms

The axiomatic method in Social Choice Theory is an established approach studying which properties are satisfied by voting rules.

- Challenge: How to generalize axioms to goal-based voting?

Two interpretations for unanimity (and others)

	issue-wise		model-wise
A	$(\mathbf{0}10)$	A	$(\mathbf{010})$
B	$(\mathbf{0}10)$	B	$(\mathbf{010})$
C	$(\mathbf{0}10)$	C	$(\mathbf{010})$
	$(\mathbf{0}11)$		(011)

Goal-based Voting | Axiomatic Results

- Negative results: Axioms often incompatible.

Theorem. No resolute F can satisfy both anonymity and duality.

- Positive results: Characterization of the rule *TrueMaj*.

Theorem. A rule is egalitarian, independent, neutral, anonymous, monotonic, unanimous and dual if and only if it is *TrueMaj*.

Goal-based Voting | Complexity Results

How hard is it to compute the outcome of a rule F?

$\textsc{WinDet}(F)$
Given profile $\mathbf{\Gamma}$ and issue $j \in \mathcal{I}$, is it the case that $F(\mathbf{\Gamma})_j = 1$?

PP: Probabilistic Polynomial Time

$\textsc{WinDet}(F)$	membership	hardness	
Approval	Θ_p^2-complete		
EMaj	PSPACE	PP	
2sMaj	P^{PP}	PP	
TrueMaj	PSPACE	PP	☹
$\gamma_i \in \mathcal{L}^{\wedge}, \mathcal{L}^{\vee}$ *EMaj*, *2sMaj*, *TrueMaj*	P		☺

gCP-nets | Framework

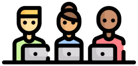

- A variable X has values $x_1, x_2, \dots$ on which agents express *ceteris paribus* preferences via CP statements
  - price $= \{\texttt{cheap}, \texttt{high}\}$, area $= \{\texttt{Capitole}, \texttt{Blagnac}, \dots\}$
  - $\texttt{high} : \texttt{Capitole} \rhd \texttt{Blagnac}$
- A CP-net N induces an order $>_N$ on possible outcomes

$$
\begin{array}{lrl}
(\varphi_1) & \top : & b_2 \rhd b_1 \\
(\varphi_2) & c \vee b_2 : & \overline{a} \rhd a
\end{array}
$$

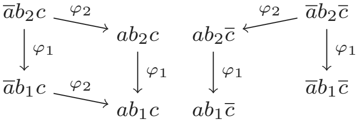

Haret, **Novaro**, Grandi. *Preference Aggregation with Incomplete CP-nets*. KR-18.

gCP-nets | Semantics

Aggregate dominance relations in
the individual CP-nets by using four semantics.

Pareto	Dominance stays if all agents have it
maj	Dominance stays if a majority of agents have it
max	Dominance stays if a majority of non-indifferent agents have it
rank	Sum of length of longest path to a non-dominated dominance class

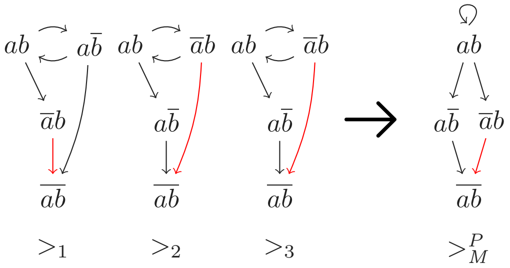

gCP-nets | Computational Problems

Dominance

DOMINANCE: $o_1 >_N o_2$

Consistency

CONSISTENCY: there is no o such that $o >_N o$

Dominance for o

wNON-DOM'ED: $o' >_N o$ implies $o >_N o'$ for all o'

NON-DOM'ED: there is no o' so that $o' >_N o$ (including $o' = o$)

DOM'ING: $o >_N o'$ for all o'

STR-DOM'ING: o is dominating and non-dominated in $>_N$

Existence

∃NON-DOM'ED: there is a non-dominated outcome in $>_N$

∃DOM'ING: there is a dominating outcome in $>_N$

∃STR-DOM'ING: there is a strongly dominating outcome in $>_N$

gCP-nets | Complexity Results

	one gCP-net	*Pareto*	*maj*	*max*	*rank*
Dominance	PSPACE-c	PSPACE-c	PSPACE-c	PSPACE-c	PSPACE-h
Consistency	PSPACE-c	PSPACE-c	PSPACE-h	PSPACE-h	—
wNon-Dom'ed	PSPACE-c	PSPACE-c	PSPACE-c	PSPACE-h	PSPACE-h
Non-Dom'ed	P	PSPACE-c	PSPACE-c	in PSPACE	—
Dom'ing	PSPACE-c	PSPACE-c	PSPACE-c	PSPACE-c	PSPACE-h
Str-Dom'ing	PSPACE-c	PSPACE-c	PSPACE-c	PSPACE-c	—
∃Non-Dom'ed	NP-c	PSPACE-c	NP-h	NP-h	—
∃Dom'ing	PSPACE-c	PSPACE-c	PSPACE-c	PSPACE-c	—
∃Str-Dom'ing	PSPACE-c	PSPACE-c	PSPACE-c	PSPACE-c	—

Most results do **not** become **harder**
when moving from one to multiple gCP-nets.

Part II: Strategic Behavior

Strategic Goal-based Voting | Example

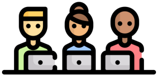

A: "Morning, guest talks, lunch."

B: "Afternoon, guest talks, no lunch."

C: "Either *afternoon, guest talks and lunch*, or *no guest talks and no lunch*."

A	(111)	(111)
B	(010)	(010)
C	(011) (100) (000)	(001)
TrueMaj	(010)	(011)

Novaro, Grandi, Longin, Lorini. *Strategic Majoritarian Voting with Propositional Goals (EA)*. AAMAS-19.

Strategic Goal-based Voting | Framework

F is resolute if it always returns a singleton output.

- An agent i is satisfied with $F(\mathbf{\Gamma})$ *iff* $F(\mathbf{\Gamma}) \subset \text{Mod}(\gamma_i)$.

F is weakly resolute $F(\mathbf{\Gamma}) = \text{Mod}(\varphi)$ for φ a conjunction on all $\mathbf{\Gamma}$.

- An agent i is satisfied with $F(\mathbf{\Gamma}) \ldots$ depends on if she is an optimist, a pessimist or an expected utility maximizer.

F is strategy-proof if for all $\mathbf{\Gamma}$ there is no agent i who would get a preferred outcome by submitting goal γ_i'.

Strategic Goal-based Voting | Results

Agents may know each other and have some ideas about their goals . . .

Unrestricted: i can send any γ_i' instead of her truthful γ_i
Erosion: i can only send a γ_i' s.t. $\text{Mod}(\gamma_i') \subseteq \text{Mod}(\gamma_i)$
Dilatation: i can send only a γ_i' s.t. $\text{Mod}(\gamma_i) \subseteq \text{Mod}(\gamma_i')$

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

Theorem. $\textsc{Manip}$(*2sMaj*) and $\textsc{Manip}$(*EMaj*) are PP-hard.

Strategic Disclosure of Opinions | Framework

"Is Toulouse the best city?"

- Agents have binary opinions on issues and they can decide to use their influence power on others

- States consist of all opinions and use of influence of agents
- An influence network is a directed irreflexive graph $E \subseteq N \times N$ s.t.

$$(i, j) \in E \text{ iff agent } i \text{ influences agent } j$$

Strategic Disclosure of Opinions | Games

The opinions update process:

1. Agents activate (or not) their influence power on (some) issues
2. Agents update opinions via unanimous aggregation

Influence Games: agents, issues, influence network, aggregation functions, initial state and individual goals (Linear Temporal Logic)

$$\text{influence}(i, C, J) = \Diamond\Box \bigwedge_{p \in J} \big((\text{op}(i,p) \to \bigcirc \text{pcon}(C,p)) \wedge (\neg\text{op}(i,p) \to \bigcirc \text{ncon}(C,p)) \big)$$

Strategic Disclosure of Opinions | A Result

Prop. Using influence is not a dominant strategy for Influence goal.

	→	→	→	
s_0	0	1	0	1
s_1	0	0	1	0
s_2	0	0	0	1

- Agent has the goal Influence(, , p)
- : always use influence power over p
- : use influence power over p unless agree on p

⇒ does not use her influence power over p in s_0

Shared and Exclusive Control | Framework

In different situations, control over issues is exclusive or shared.

A Potluck

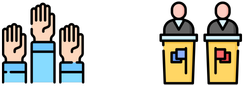

Group Decisions

- Iterated games where agents have goals in LTL
- Logics ATL and ATL* to reason about the games, interpreted over Concurrent Game Structures

Belardinelli, Grandi, Herzig, Longin, Lorini, **Novaro**, Perrussel. *Relaxing Exclusive Control in Boolean Games*. TARK-17.

Shared and Exclusive Control | Result

Theorem. Verification of ATL* formulas on CGS with shared control (SPC) reducible to CGS with exclusive control (EPC).

- Define a corresponding CGS-EPC from a given CGS-SPC
- Define a translation function tr within ATL*
- Show that the CGS-SPC satisfies φ if and only if the corresponding CGS-EPC satisfies $tr(\varphi)$

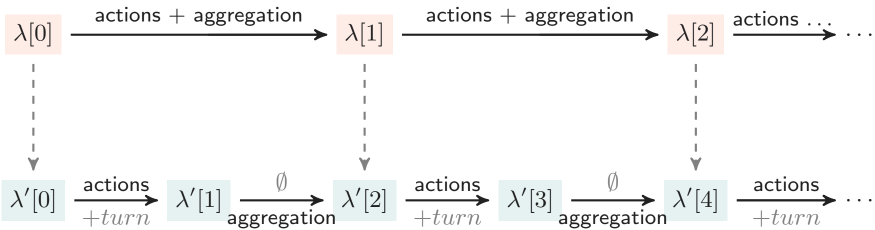

Conclusion and Perspective

Conclusion | →

1. How can we design aggregation procedures to help a group of agents having compactly expressed goals and preferences make a collective choice?

Goal-based Voting

Framework where agents can express complex goals compactly

Many interesting rules, and characterization result for *TrueMaj*

WinDet hard in general, but restrictions make it tractable

Aggregation of gCP-nets

Agents can state incomplete preferences, then aggregated

Most results do not become harder with respect to a single agent

Conclusion | 

2. How can we model agents with conflicting goals who try to get a better outcome for themselves by acting strategically?

Majoritarian Goal-based Voting

Strategy-proofness for restrictions on language and strategies

Disclosure of Opinions on Networks

Intuitive idea, complex dynamic: results for specific graphs and goals

Shared and Exclusive Control in Concurrent Game Structures

Natural model for shared control, still reducible to exclusive control

Perspectives

- **Explain** axioms to users (choose when incompatible)
- **Characterize** language restrictions giving tractability
- Opinion **delegation** rather than diffusion

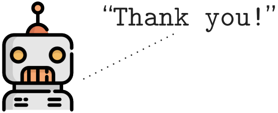

Credits to Freepik, Lyolya, Nikita Golubev, smalllikeart at `flaticon.com` for the icons.