Complexity of Judgment Aggregation: Safety of the Agenda

Ulle Endriss Umberto Grandi Daniele Porello

Institute for Logic, Language and Computation
University of Amsterdam

23 April 2010

Computational Complexity and Judgment Aggregation

Judgment aggregation is relevant for multi-agent systems

$$\textbf{JA} \quad \Longleftarrow \Longrightarrow \quad \textbf{MAS}$$

Societies of agents have to aggregate their opinions/judgments

> Is it feasible?
>
> Computational complexity provides a first measure.

Feedback for JA:

- precise formal framework;
- new problems arise: safety of the agenda.
- connection with work on complexity of manipulation in voting theory;

Outline

1. Basic framework of judgment aggregation for propositional logic:
   - Distinction between logical and syntactical properties of judgment sets;
   - Representation results in axiomatic approach.
2. Safety of the agenda:
   - Is consistency guaranteed for all aggregators?
   - Classical results in the literature are *possibility* results.
3. Characterisation results for agendas:
   - When is safety guaranteed for different classes of aggregators;
   - Introduce properties of sets of formulae: k-MP, SMP, SSMP.
4. Complexity results: checking safety is Π_p^2-complete.
   - Use characterisation results to reduce to SSMP-like problems;
   - Use *quantified boolean formulas* to prove membership;
   - Reduce co-QSAT_2 to SSMP to prove completeness.

Discursive Dilemma

JA was developed to generalise and study paradoxical situations that arises when a collective judgment has to be made on a set of correlated propositions.

Discursive Dilemma

	p	$p \rightarrow q$	q
Judge 1:	Yes	Yes	Yes
Judge 2:	No	Yes	No
Judge 3:	Yes	No	No
Majority:	Yes	Yes	No

Each individual is rational (i.e., has a consistent judgment) but the majority is contradictory!

Kornhauser and Sager, Unpacking the court. Yale Law Journal, 1986.

List and Pettit, Aggregating sets of judgments. Economics and Philosophy, 2002

Basic Definitions

A set N of individuals expressing judgments on a set of correlated propositions:

Definition

An agenda is a finite subset of propositional formulas $\Phi \subseteq \mathcal{L}_{PS}$ closed under complementation and not containing double negations.

A judgment set on an agenda Φ is a subset $J \subseteq \Phi$. Call J:

Complete: if for all $\alpha \in \Phi$ either α or its complement is in J. (easy)

Complement-free: α and its complement $\sim \alpha$ are never both in J. (easy)

Consistent: there is an assignment to make all formulas in J true. (not easy!)

Call $J(\Phi)$ the set of all consistent and complete judgment sets.

Definition

An aggregation procedure for agenda Φ and a set N of n individuals is a function $F : J(\Phi)^n \to 2^{\Phi}$.

Axioms for Aggregation Procedures I

Aggregation procedures have been studied using the axiomatic method, listing axioms as desirable properties of the functions:

Weak Rationality (WR): F is non-null, complete and complement-free.

These properties are all computationally easy.

Other standard axioms:

Unanimity (U): If $\varphi \in J_i$ for all i then $\varphi \in F(\mathbf{J})$.

Anonymity (A): F is symmetric with respect to individuals.

Neutrality (N): If φ and ψ share the same pattern of individuals' judgments then their outcome must be the same.

Independence (I): The outcome of F over φ depends only on the individuals' judgments over φ.

Systematicity (S)=(N)+(I).

Two forms of monotonicity for independent (M^I) and neutral procedures (M^N).

Axioms for Aggregation Procedures II

We can play with this formalism to get (small) interesting results:

Lemma

If an agenda Φ contains a tautology, then every aggregation procedure for Φ that satisfies (WR) *and* (S) *is unanimous* (U).

Proof. If $\varphi^\top$ is a tautology then it is in all individual judgment sets and $\varphi^\top \in F(\mathbf{J})$ for a certain profile $\mathbf{J}$ (by non-nullness). If a formula is unanimously accepted it behaves like $\varphi^\top$ and by systematicity (S) it must be accepted.

Lemma

If the number of individuals is even, then there exists no aggregation procedure that satisfies (WR), (A) *and* (N).

Proof. By (N) and (A) the acceptance of a formula φ depends only on the number of individuals supporting φ in profile $\mathbf{J}$. The profile where half of individuals accepts φ and half accepts $\neg\varphi$ is in contradiction with (WR).

Classes of Aggregation Procedures

Such axioms can be used to define different classes of aggregation procedures:

Set of axioms AX Agenda Φ	$\Longrightarrow$	Class of functions $\mathcal{F}_{\Phi}[\mathsf{AX}]$

Proposition (representation result)

$\mathcal{F}_{\Phi}[\mathsf{WR}, \mathsf{A}, \mathsf{S}, \mathsf{M}^{\mathsf{I}}] = \mathcal{F}_{\Phi}[\mathsf{WR}, \mathsf{A}, \mathsf{N}, \mathsf{M}^{\mathsf{N}}]$ *is the majority rule.*

We will concentrate on various weakening of the axiomatisation of majority:

$$\mathcal{F}_{\Phi}[\mathsf{WR}, \mathsf{A}, \mathsf{S}] \qquad \mathcal{F}_{\Phi}[\mathsf{WR}, \mathsf{A}, \mathsf{N}]$$

$$\mathcal{F}_{\Phi}[\mathsf{WR}, \mathsf{A}, \mathsf{I}] \qquad \mathcal{F}_{\Phi}[\mathsf{A}, \mathsf{S}, \mathsf{M}^{\mathsf{I}}]$$

Similar representation results can be proved for these classes.

E.g., $\mathcal{F}_{\Phi}[\mathsf{A}, \mathsf{S}, \mathsf{M}^{\mathsf{I}}]$ is the class of uniform quota rules, (Dietrich and List, 2007)

Safety of the Agenda I

Call a function consistent if $F(\mathbf{J})$ is consistent for all profiles $\mathbf{J}$ in the domain.

A discursive dilemma for F is a profile $\mathbf{J} = (J_1, \ldots, J_n)$ such that the outcome $F(\mathbf{J})$ is inconsistent.

Safety of the Agenda I

Call a function consistent if $F(\mathbf{J})$ is consistent for all profiles $\mathbf{J}$ in the domain.

Possibility (characterisation) results:

There exists a consistent F in $\mathcal{F}_\Phi$[AX] $\Leftrightarrow$ Φ has a certain property.

These results characterise agendas for which consistent aggregation is possible.

E.g., Nehring and Puppe (2006), Dietrich and List (2007), Dokow and Holzman (2005).

Safety (characterisation) results:

All F in $\mathcal{F}_\Phi$[AX] are consistent $\Leftrightarrow$ Φ has a certain property.

These results characterise agendas for which consistent aggregation is guaranteed.

Safety of the Agenda II

Definition

An agenda Φ is safe with respect to a class of aggregation procedures $\mathcal{F}_\Phi$ if every function in $\mathcal{F}_\Phi$ is consistent.

A discursive dilemma is a counterexample for the safety of an agenda.

Proposition

Φ is safe for $\mathcal{F}_\Phi[\mathrm{WR}, \mathrm{A}, \mathrm{S}, \mathrm{M^I}]$ iff every inconsistent subset of Φ contains an inconsistent subset of size 2 (median property MP).

Proof.

Majority rule is the only function in $\mathcal{F}_\Phi[\mathrm{WR}, \mathrm{A}, \mathrm{S}, \mathrm{M^I}]$.

Therefore proving safety = proving possibility.

The proof follows from Nehring and Puppe (2006). □

Agenda Properties

An agenda Φ satisfies:

- the k-median property (kMP) for $k \geqslant 2$, if every inconsistent subset of Φ has itself an inconsistent subset of size at most k (2MP=MP).
- the simplified median property (SMP), if every nontrivially inconsistent subset of Φ (i.e., not containing any contradictory formula) has itself an inconsistent subset of the form $\{\varphi, \psi\}$ with $\models \varphi \leftrightarrow \neg\psi$.
- the syntactic simplified median property (SSMP), if every nontrivially inconsistent subset of Φ has an inconsistent subset of the form $\{\varphi, \neg\varphi\}$.

$$\text{SSMP} \Rightarrow \text{SMP} \Rightarrow \text{MP} \Rightarrow \text{kMP}$$

Agendas satisfying the SSMP are composed of uncorrelated formulas, i.e., they are essentially equivalent to agendas composed by atoms alone.
E.g. $\Phi = \{p \wedge q, \neg(p \wedge q), r \rightarrow (p \vee \neg p), \neg(r \rightarrow (p \vee \neg p))\}$

Characterisation Results I

Proposition

Φ is safe for $\mathcal{F}_\Phi$[WR, A, S] if and only Φ satisfies the SMP*.*

Proof.

$\Leftarrow$) If $F(\mathbf{J})$ is (non-trivially) inconsistent then it contains a subset $\{\varphi, \psi\}$ with $\varphi \leftrightarrow \neg\psi$. By individual rationality we know that $\varphi \in J_i \Leftrightarrow \neg\psi \in J_i$.
By systematicity this implies $\varphi \in F(\mathbf{J}) \Leftrightarrow \neg\psi \in F(\mathbf{J})$.
Since $\{\varphi, \psi\} \subseteq F(\mathbf{J})$ we should have also $\neg\psi \in F(\mathbf{J})$, contradicting (WR).

$\Rightarrow$) Suppose Φ violates the SMP: there exists two formulas s.t. $\varphi \vdash \neg\psi$ and $\neg\psi \not\vdash \varphi$ and we can build an inconsistent function in $\mathcal{F}_\Phi$[WR,A,S]:

	φ	ψ
Agent 1:	No	No
Agent 2:	No	Yes
Agent 3:	Yes	No
$F(\mathbf{J})$:	Yes	Yes

Safety Characterisation Results

We prove the following:

Proposition

Φ is safe for $\mathcal{F}_\Phi[\mathrm{WR}, \mathrm{A}, \mathrm{N}]$ if and only Φ satisfies the SMP *and does not contain a contradictory formula.*

Φ is safe for $\mathcal{F}_\Phi[\mathrm{WR}, \mathrm{A}, \mathrm{I}]$ if and only Φ satisfies the SSMP.

Reformulating work proved by Dietrich and List (2007) we obtain:

Proposition

*Let $k \geqslant 2$. An agenda Φ is safe for the class of uniform quota rules F_m for n individuals satisfying $m > n - \frac{n}{k}$ if and only if Φ satisfies the k*MP

.

Computational Complexity: Identifying the problem

Our problem:
Is an agenda Φ safe for the class $\mathcal{F}_\Phi[\mathrm{AX}]$?

We can define the following problem:

SAFETY[AX]	
Instance:	Agenda Φ of size n
Problem:	Is Φ safe for $\mathcal{F}_\Phi[\mathrm{AX}]$?

Using our characterisation, e.g. for $\mathcal{F}_\Phi[\mathrm{WR}, \mathrm{A}, \mathrm{I}]$, this boils down to checking the agenda property:

SSMP	
Instance:	Agenda Φ of size n
Problem:	Does Φ satisfy SSMP?

The same can be done for the remaining problems using SMP, MP and kMP.

A Complexity Class: Π_2^p

If M is Turing maching running in polynomial time, the following holds:

$$\mathbf{NP} \sim \exists \underline{x} M(z, \underline{x}) = 1 \quad \textbf{co-NP} \sim \forall \underline{x} M(z, \underline{x}) = 1$$

Above **P** and **NP** sits the polynomial hierarchy **PH**. Its second layer is:

$$\Sigma_2^p \sim \exists \underline{x} \forall \underline{y} M(z, \underline{x}, \underline{y}) = 1 \quad \Pi_2^p \sim \forall \underline{x} \exists \underline{y} M(z, \underline{x}, \underline{y}) = 1$$

A complete problem for Π_2^p is the following satisfiability problem:

coQSAT$_2$

Instance: Quantified boolean formula (QBF)
$\forall x_1, \ldots, x_r . \exists y_1 \ldots, y_s . \varphi(x_1, \ldots, x_r, y_1, \ldots, y_s)$

Problem: Is the QBF satisfiable (i.e., true)?

SSMP and QBF

Let us reformulate the simplified syntactic median property as follows:

SSMP

$\Phi = \Phi^+ \cup \neg\Phi^+$ *satisfies the* SSMP *if and only if for any set of formulas in* Φ^+ *the conjunction of those formulas and of the negation of all other formulas in* Φ^+ *is consistent. (Valid for agendas without tautologies)*

This suggest that our SoA problem can be embedded into a QBF. Indeed:

$$[\forall y_1 \cdots \forall y_{n^+} \exists x_1 \cdots \exists x_m : (y_1 \leftrightarrow \varphi_1^+) \wedge \cdots \wedge (y_{n^+} \leftrightarrow \varphi_{n^+}^+)]$$

is a QBF expressing that the agenda Φ satisfies the reformulation of the SSMP.

This proves membership to Π_2^p.

Problems in Π_2^p are very hard ones, but there exist theorem provers for QBFs: http://www.qbflib.org

Computational Complexity of SSMP

Checking safety turns out to be quite hard:

Theorem

SAFETY[WR, A, I] *is* Π_2^p*-complete.*

Proof.

- By characterisation this problem is equivalent to SSMP;
- SSMP is in Π_2^p (membership): reduction to co-QSAT$_2$.
- SSMP is Π_2^p-complete (hardness): by reduction from co-QSAT$_2$. For every QBF Δ we can build an agenda Φ_Δ such that if Φ_Δ satisfies the SSMP iff the initial QBF Δ is satisfiable.

□

Complexity results

The following theorem is an overview of our complexity results:

Theorem

Checking the safety of an agenda is Π_2^p-complete for any of these classes of aggregation procedures:

1. *the majority rule, corresponding to* $\mathcal{F}_\Phi[\mathsf{WR}, \mathsf{A}, \mathsf{S}, \mathsf{M}^\mathsf{I}]$ *and* $\mathcal{F}_\Phi[\mathsf{WR}, \mathsf{A}, \mathsf{N}, \mathsf{M}^\mathsf{N}]$*;*
2. *systematic rules:* $\mathcal{F}_\Phi[\mathsf{WR}, \mathsf{A}, \mathsf{S}]$*;*
3. *neutral rules:* $\mathcal{F}_\Phi[\mathsf{WR}, \mathsf{A}, \mathsf{N}]$*;*
4. *independent rules:* $\mathcal{F}_\Phi[\mathsf{WR}, \mathsf{A}, \mathsf{I}]$*;*
5. *for any $k \geqslant 2$, the class of uniform quota rules F_m with $m > n - \frac{n}{k}$, where n is the number of individuals.* $\mathcal{F}_\Phi^{k,m,n}[A,S,M^I]$.

Proof.

SSMP, SMP, kMP for $k \geqslant 2$, are all Π_2^p-complete. □

Conclusions and Future Work

In conclusion:

- To the best of our knowledge the safety of the agenda is a new problem relevant to an application-driven study of JA.
- We have proved several characterisation results to recognise safe agendas for different classes of functions.
- We have proved that this problem is Π_2^p-complete for all classes of functions we considered.

Guaranteeing safety for various classes of functions is a very restrictive requirement, and checking these simple properties is computationally hard.

Future work includes:

- Study the complexity of the safety problem for other classes of functions.
- Introduce and study the safety of a profile.
- Study the complexity of strategic manipulation in JA.