Computational Complexity of Judgment Aggregation

Ronald de Haan

Computational Social Choice: Spring 2019

Institute for Logic, Language and Computation
University of Amsterdam

Plan for today

- We will look at computational complexity considerations in Judgment Aggregation
- Various computational problems arise:
  - Outcome determination
  - Problems related to strategic behavior
  - Agenda safety
  - *(and more..)*
- We will use the Kemeny procedure as illustrating example

Computational Complexity

- Remember: P, NP, polynomial-time reductions
- $\Theta_2^{\mathsf{p}} = \mathsf{P}^{\mathsf{NP}}[\log]$:
  - Solvable in polynomial time with $O(\log n)$ NP oracle queries
- $\Sigma_2^{\mathsf{p}} = \mathsf{NP}^{\mathsf{NP}}$:
  - Solvable in nondeterministic polynomial time with an NP oracle
- $\Pi_2^{\mathsf{p}} = \mathsf{coNP}^{\mathsf{NP}}$:
  - Complement of the problem in $\mathsf{NP}^{\mathsf{NP}}$

$$\mathsf{P} \subseteq \mathsf{NP} \subseteq \Theta_2^{\mathsf{p}} \subseteq \Sigma_2^{\mathsf{p}}, \Pi_2^{\mathsf{p}}$$

Judgment Aggregation with an Integrity Constraint

- Agenda: a set $\Phi = \{x_1, \neg x_1, \ldots, x_m, \neg x_m\}$ of propositional variables and their negations
- Integrity constraint: a propositional formula Γ
- Judgment set: $J \subseteq \Phi$
  - consistent if $J \cup \{\Gamma\}$ is satisfiable
  - complete if $\{x_i, \neg x_i\} \cap J \neq \emptyset$ for each $1 \leq i \leq m$
  - admissible if consistent and complete
  - $\mathcal{J}(\Phi, \Gamma)$ denotes the set of all admissible judgment sets
- Profile: a sequence $\boldsymbol{J} = (J_1, \ldots, J_n)$ of admissible judgment sets
- Judgment aggregation procedure: a function F that assigns to each profile $\boldsymbol{J}$ a set $F(\boldsymbol{J})$ of judgment sets (the outcomes)

The Kemeny Rule in JA

- The Kemeny rule selects those admissible judgment sets J that minimize the cumulative distance to the profile $\boldsymbol{J}$:

$$F_{\text{Kemeny}}(\boldsymbol{J}) = \underset{J \in \mathcal{J}(\Phi,\Gamma)}{\text{argmin}} \sum_{i \in N} H(J, J_i), \quad \text{where } H(J, J_i) = |J \setminus J_i|$$

- Example:

$$\Gamma = \begin{array}{l}(\neg x_1 \vee \neg x_2 \vee \neg x_3) \wedge \\ (\neg x_2 \vee \neg x_3 \vee \neg x_4)\end{array}$$

$$F_{\text{Kemeny}}(\boldsymbol{J}) = \{\{x_1, x_2, \neg x_3, x_4\}, \{x_1, \neg x_2, x_3, x_4\}\}$$

$\boldsymbol{J}$	x_1	x_2	x_3	x_4
J_1	1	0	1	1
J_2	1	1	0	1
J_3	0	1	1	0
J_4	1	0	1	1
J_5	1	1	0	1
maj	1	1	1	1

Outcome Determination

- Ultimately, we want to find outcomes: this is a search problem
- There are several ways to cast this as a decision problem
- (Note: *"Does there exist some* $J \in F(\boldsymbol{J})$*?"* is trivial)
- We will use the following variant:

Outcome-Determination(F)

Input: An agenda Φ, an integrity constraint Γ, a profile $\boldsymbol{J} \in \mathcal{J}(\Phi, \Gamma)^+$, and a formula $\varphi^* \in \Phi$ from the agenda.

Question: Is there a judgment set $J^* \in F(\boldsymbol{J})$ such that $\varphi^* \in J^*$?

Membership in $\Theta_2^p = P^{NP}[\log]$

- To show that Outcome-Determination(Kemeny) is in Θ_2^p, we describe a polynomial-time algorithm that queries an NP oracle $\log(n \cdot m)$ times:

1. Find the minimum cumulative Hamming distance k^* of any $J \in \mathcal{J}(\Phi, \Gamma)$ to $\boldsymbol{J}$:
   - Use binary search to find k^* by querying the NP oracle to answer questions *"Is there some $J \in \mathcal{J}(\Phi, \Gamma)$ whose cumulative Hamming distance to $\boldsymbol{J}$ is $\leq k$?"*
2. Then ask the NP oracle: *"Is there some $J \in \mathcal{J}(\Phi, \Gamma)$ whose cumulative Hamming distance to $\boldsymbol{J}$ is k^* with $\varphi^* \in J$?"* and return the same answer

- All oracle queries are problems in NP, so we can do this with a single NP-complete oracle (with polynomial overhead)

Θ_2^p-hardness

- To show that Outcome-Determination(Kemeny) is Θ_2^p-hard, we will give a polynomial-time reduction from the following Θ_2^p-complete problem:

Max-Model	
Input:	A satisfiable propositional logic formula ψ, and some $x^* \in \text{var}(\psi)$.
Question:	Is there a maximal model of ψ that sets x^* to true?

- A maximal model of ψ is a truth assignment to $\text{var}(\psi)$ that satisfies ψ and that sets a maximum number of variables in $\text{var}(\psi)$ to true (among those that satisfy ψ)

Θ_2^{p}-hardness (the reduction)

- Let (ψ, x^*) be an instance of Max-Model, with $\text{var}(\psi) = \{x_1, \dots, x_m\}$. We construct $\Phi, \Gamma, \boldsymbol{J}, \varphi^*$ as follows:
  - $\Phi = \text{lit}(\psi) \cup \{z_{i,j}, \neg z_{i,j} : 1 \leq i \leq 3, 1 \leq j \leq 2m\}$
  - $\Gamma = \psi \vee \bigvee_{1 \leq i \leq 3} \bigwedge_{1 \leq j \leq 2m} z_{i,j}$
  - $\varphi^* = x^*$
  - $\boldsymbol{J} = (J_1, J_2, J_3)$:

$\boldsymbol{J}$	x_1	x_2	$\cdots$	x_m	$z_{1,1}$	$z_{2,1}$	$z_{3,1}$	$\cdots$	$z_{1,2m}$	$z_{2,2m}$	$z_{3,2m}$
J_1	1	1	$\cdots$	1	1	0	0	$\cdots$	1	0	0
J_2	1	1	$\cdots$	1	0	1	0	$\cdots$	0	1	0
J_3	1	1	$\cdots$	1	0	0	1	$\cdots$	0	0	1

Θ_2^p-hardness (correctness of the reduction)

For any judgment set J to be Γ-consistent, either (i) $J \cup \{\psi\}$ needs to be consistent, or (ii) $J \cup \{\bigvee_{1 \leq i \leq 3} \bigwedge_{1 \leq j \leq 2m} z_{i,j}\}$.
In case (i), $\sum_{1 \leq i \leq n} H(J, J_i) \leq 3m$. In case (ii), $\sum_{1 \leq i \leq n} H(J, J_i) \geq 4m$.

($\Rightarrow$) Suppose x^* is made true by some maximal model α of ψ.

Take $J_\alpha = \{x_i : 1 \leq i \leq m, \alpha(x_i) = 1\} \cup \{\neg x_i : 1 \leq i \leq m, \alpha(x_i) = 0\} \cup \{\neg z_{i,j} : 1 \leq i \leq 3, 1 \leq j \leq 2m\}$. J_α is Γ-consistent, contains x^* and has cumulative Hamming distance $\leq 3m$ to the profile $\boldsymbol{J}$.

There is no $J' \in \mathcal{J}(\Phi, \Gamma)$ with smaller cumulative Hamming distance to $\boldsymbol{J}$—if such a J' would exist, there would be some α' satisfying ψ that sets more variables to true than α. Thus, $J \in F_{\text{Kemeny}}(\boldsymbol{J})$.

($\Leftarrow$) Suppose there is some $J \in F_{\text{Kemeny}}(\boldsymbol{J})$ with $x^* = \varphi^* \in J$. We know that $J \cup \{\psi\}$ is satisfiable.
Let α be the truth assignment such that $\alpha(x_i) = 1$ if and only if $x_i \in J$, for each $1 \leq i \leq n$. Then α satisfies ψ and sets x^* to true.

There is no α' satisfying ψ that sets more variables to true than α—if such an α' would exists, there would be some J' with smaller cumulative Hamming distance to $\boldsymbol{J}$. Thus, α is a maximal model of ψ.

Strategic Behavior: Manipulation

- Strategic manipulation: an individual submitting an insincere judgment set to get a preferred outcome
- There are several ways to cast this as a decision problem. We will use the following variant:

Manipulation(F)

Input: An agenda Φ, an integrity constraint Γ, a profile $\boldsymbol{J} = (J_1, \ldots, J_n)$, and a set $L \subseteq \Phi$.

Question: Is there an admissible judgment set $J' \in \mathcal{J}(\Phi, \Gamma)$ such that for all $J^* \in F_{\text{Kemeny}}(J', J_2, \ldots, J_n)$ it holds that $L \subseteq J^*$?

Strategic Behavior: Manipulation

- **Theorem:** Manipulation(Kemeny) is Σ_2^{p}-complete

  - Intuition why the problem is in $\Sigma_2^{\mathrm{p}} = \mathrm{NP}^{\mathrm{NP}}$:

  1. Guess a (strategizing) judgment set J'
     (nondeterministic/NP guess)
  2. Solve the problem of outcome determination for $(J', J_2, \ldots, J_n)$
     (using NP oracle queries)

  - Σ_2^{p}-hardness by reduction from $\exists\forall$-TQBF

- One can see this hardness as a barrier against manipulation

R. de Haan. *Complexity results for manipulation, bribery and control of the Kemeny judgment aggregation procedure.* In: Proceedings of AAMAS 2017, pp. 1151–1159.

Agenda Safety

- An agenda Φ and an integrity constraint Γ are safe for the majority rule if and only if there is no minimally Γ-inconsistent subset $L \subseteq \Phi$ of size > 2
  - Safety: for every possible profile $\boldsymbol{J}$, the outcome is Γ-consistent
  - Minimally Γ-inconsistent set L: $L \cup \{\Gamma\}$ is unsatisfiable, and for each $L' \subsetneq L$, $L' \cup \{\Gamma\}$ is satisfiable
- Idea: if there is some minimally Γ-inconsistent L of size ≥ 3, you can construct a "doctrinal paradox" situation

Agenda-Safety

Input: An agenda Φ, and an integrity constraint Γ.

Question: Is there no minimally Γ-inconsistent $L \subseteq \Phi$ of size > 2?

Agenda Safety

- **Theorem:** Agenda-Safety is Π_2^p-complete

  - Intuition why the problem is in $\Pi_2^p = \text{coNP}^{\text{NP}}$:

  1. Quantify over all possible $L \subseteq \Phi$ of size ≥ 3
     (nondeterministic/coNP guess)
  2. Quantify over all truth assignments for $L \cup \{\Gamma\}$,
     and check that none is satisfying
     (nondeterministic/coNP guess)
  3. Check that all $L' \subsetneq L$ are Γ-consistent (using NP oracle queries)

  - Π_2^p-hardness by reduction from $\forall\exists$-TQBF

U. Endriss, U. Grandi, and D. Porello. Complexity of Judgment Aggregation.
Journal of Artificial Intelligence Research (JAIR), 45, 481–514, 2012.

All bad news?

- Computational complexity results for the Kemeny rule in JA are generally negative
- Similar results for other rules (at least those that work for any agenda and that guarantee consistent outcomes)
- Does this mean that we cannot use Judgment Aggregation to model social choice scenarios in practice?

- **No!** Research: find particular cases where, say, Outcome-Determination(Kemeny) is efficiently solvable
  - Simple example: if Γ is in DNF, we can solve Outcome-Determination(Kemeny) in polynomial time
  - Idea: iterate over all disjuncts of the DNF and find which one allows for minimum cumulative Hamming distance to the profile

Conclusion

- We looked at several computational problems that arise in the setting of Judgment Aggregation, and their computational complexity (using the Kemeny rule as example)
- Most results are worst-case intractability results
  - Some are obstacles (e.g., for outcome determination)
  - Some can be seen as helpful (e.g., for strategic manipulation)
- To use Judgment Aggregation as an applied general system to model social choice applications, computational complexity considerations are important