Judgment Aggregation

Markus Brill
brill@cs.duke.edu

Based on *Tutorial on Judgment Aggregation* by **Ulle Endriss**
https://staff.fnwi.uva.nl/u.endriss/teaching/aamas-2013/

Duke UNIVERSITY

Alexander von Humboldt Stiftung/Foundation

Doctrinal Paradox

- Three judges have to decide on a case of an alleged breach of contract
- The need to decide whether (a) the contract is **valid** and whether (b) the contract has been **breached**.
- Legal doctrine stipulates that the defendant is **liable** if and only if (a) and (b) hold.

	Valid?	*Breach?*	*Liable?*
Judge 1	**Yes**	**Yes**	**Yes**
Judge 2	**Yes**	**No**	**No**
Judge 3	**No**	**Yes**	**No**

Discursive Dilemma

- Instead of expressing **judgments on atoms** (valid, breached, liable) and imposing **constraints** (liable iff valid and breached), we could allow **judgments on compound formulas**:

	p	q	$p \wedge q$
Judge 1	**Yes**	**Yes**	**Yes**
Judge 2	**Yes**	**No**	**No**
Judge 3	**No**	**Yes**	**No**
Majority	**Yes**	**Yes**	**No**

	p	q	$r \leftrightarrow p \wedge q$	r
Judge 1	**Yes**	**Yes**	**Yes**	**Yes**
Judge 2	**Yes**	**No**	**Yes**	**No**
Judge 3	**No**	**Yes**	**Yes**	**No**
Majority	**Yes**	**Yes**	**Yes**	**No**

Why Paradox?

	p	q	$p \wedge q$
Judge 1	**Yes**	**Yes**	**Yes**
Judge 2	**Yes**	**No**	**No**
Judge 3	**No**	**Yes**	**No**
Majority	**Yes**	**Yes**	**No**

- **Reason 1:** **Premise-based procedure** and **conclusion-based procedure** produce different outcomes.
- **Reason 2:** Even though each individual judgment is logically consistent, the **majority outcome** is not.

Outline

- Formal framework
- Axioms and impossibilities
- Aggregation Rules
  - quota-based rules
  - distance-based rules
  - premise-based rules
- Computational aspects
- Strategic aspects

Formal Framework

- An **agenda** A is a finite nonempty subset of propositional formulas closed under complementation (i.e., $\phi \in A \Rightarrow \neg\phi \in A$)
- A **judgment set** J on an agenda A is a subset of A
  - J is **complete** if $\phi \in J$ or $\neg\phi \in J$ for all $\phi \in A$
  - J is **consistent** if there exists an assignment satisfying all formulas in A
- A finite set of **individuals** $N=\{1,\dots,n\}$ express judgements on the formulas in A, producing a **profile** $J=(J_1,\dots,J_n)$
  - we assume that $n \geq 3$ and n is odd
- An **aggregation rule** maps a profile of complete and consistent individual judgement sets to a single collective judgment set

Example: Majority Rule

- $N = \{1,2,3\}$
- $A = \{p, \neg p, q, \neg q, (p \vee q), \neg(p \vee q)\}$

	p	q	$p \vee q$
Agent 1	**Yes**	**Yes**	**Yes**
Agent 2	**Yes**	**No**	**Yes**
Agent 3	**No**	**No**	**No**

- **Majority rule** returns all propositions accepted by $> n/2$ agents
  - in the example: $F_{maj}(J) = \{p, \neg q, p \vee q\}$ (complete and consistent!)

Relation to Voting

- Preference aggregation can be **embedded into JA** as follows:
  - Suppose the set of alternatives is given by $\{a,b,c\}$
  - Take atomic propositions to be $p_{a>b}$, $p_{a>c}$, etc.
  - All individuals accept the following propositions:
    - $p_{a>b} \vee p_{b>a}$ etc. (completeness)
    - $p_{a>b} \wedge p_{b>c} \longrightarrow p_{a>c}$ etc. (transitivity)
- The **Condorcet paradox** in JA language:

	$p_{a>b}$	$p_{a>c}$	$p_{b>c}$	corresponding order
Agent 1	**Yes**	**Yes**	**Yes**	a>b>c
Agent 2	**No**	**No**	**Yes**	b>c>a
Agent 3	**Yes**	**No**	**No**	c>a>b
Majority	**Yes**	**No**	**Yes**	**not a linear order**

Axioms

- **Anonymity**: treat all individuals symmetrically
- **Neutrality**: treat all propositions symmetrically
  - if two formulas are accepted by the same judges, then the collective must accept either both or neither of them
- **Independence**: only the “pattern of acceptance” should matter
  - if a formula is accepted in some profile, then it is accepted in all profiles that have the same judgements on that formula
- **Monotonicity**: if an collectively accepted formula is accepted by an additional agent, it should still get collectively accepted
- **Completeness**: the collective judgement is always complete
- **Consistency**: the collective judgement is always consistent

Impossibility

- Majority rule violates consistency
- Question: Are there "reasonable" **consistent aggregation rules**?
- **Theorem** (List and Pettit, 2002): There is no aggregation rule that simultaneously satisfies **anonymity**, **neutrality**, **independence**, **completeness**, and **consistency**.
    - assuming the agenda is sufficiently rich (e.g., $\{p,q,p \wedge q\} \subseteq A$) and n>1
    - proof sketch: the first three properties implies that acceptance of a formula only depends on the number of judges that accept it
- What now? Weaken axioms? Which axioms?

Quota-based Rules

- A **quota rule** F_q is defined by a function $q\colon A \rightarrow \{0,1,\ldots,n+1\}$ and accepts a formula ϕ iff ϕ is accepted by at least $q(\phi)$ individuals
  - a quota rule is called **uniform** if there is a k s.t. $q(\phi)=k$ for all $\phi \in A$
- Examples of uniform quota rules:
  - the **constant rule** F_0 (F_{n+1}) accepts all (no) formulas
  - the **intersection rule** F_n accepts ϕ iff everybody does
  - the **(strict) majority rule** F_{maj} has quota $q = \lceil (n+1)/2 \rceil$
- What are the advantages of high/low quotas?

Axiomatic Characterizations

- **Proposition** (Dietrich and List, 2007): An aggregation rule is anonymous, independent and monotonic iff it is a **quota rule**.

Clearly, a quota rule is neutral iff it is uniform. Therefore:

- **Corollary:** An aggregation rule is anonymous, neutral, independent and monotonic (= **ANIM**) iff it is a **uniform quota rule**.
- **Corollary:** An aggregation rule is ANIM, complete and complement-free iff it is the **(strict) majority rule**.
  - a rule is **complement-free** if it never accepts both ϕ and $\neg\phi$
  - this holds for odd n; for even n, no rule satisfies these properties

Distance-based Rules

- **Idea:** Find a **consistent** judgment set that minimizes the "distance" to the profile
- **Hamming distance** between two judgment sets is given by the number of disagreements
  - $H(J_i, J_i') = 1/2 * |J_i \Delta J_i'|$
  - distance to a profile given by sum of distances to individual judgment sets in the profile

p	*q*	*r*
No	No	Yes
Yes	Yes	Yes

- Two ways to define aggregation rule based on Hamming distance:
  - minimize Hamming distance to profile **← generalized Kemeny rule**
  - minimize Hamming distance to majority outcome **← generalized Slater rule**

Example

ϕ_1 and ϕ_2 are both equivalent to
$p \vee (q_1 \wedge q_2) \vee (r_1 \wedge r_2 \wedge r_3)$

	p	q_1	q_2	r_1	r_2	r_3	ϕ_1	ϕ_2
1 agent	Yes	No	No	No	No	No	Yes	Yes
10 agents	No	Yes	Yes	No	No	No	Yes	Yes
10 agents	No	No	No	Yes	Yes	Yes	Yes	Yes

Kemeny	No	Yes	Yes	No	No	No	Yes	Yes
Slater	Yes	No	No	No	No	No	Yes	Yes

Premise-based Rules (1)

- Idea of premise-based rules (PBRs): Divide formulas into **premises** and **conclusions**, apply **majority rule on premises**, and accept conclusions that **logically follow** from accepted premises
- A PBR yields complete and **consistent outcomes** if
  - the set of premises is the set of literals
  - the agenda is closed under propositional letters

	p	q	$p \wedge q$
Judge 1	**Yes**	**Yes**	**Yes**
Judge 2	**Yes**	**No**	**No**
Judge 3	**No**	**Yes**	**No**
PBR	**Yes**	**Yes**	**Yes**

Premise-based Rules (2)

- A PBR might violate **unanimity**
  - … and “composition-consistency”

	$p \vee q$	r	$(p \vee q) \vee r$
Judge 1	**Yes**	**No**	**Yes**
Judge 2	**Yes**	**No**	**Yes**
Judge 3	**No**	**Yes**	**Yes**
PBR	**Yes**	**No**	**Yes**

	p	q	r	$p \vee q \vee r$
Judge 1	**Yes**	**No**	**No**	**Yes**
Judge 2	**No**	**Yes**	**No**	**Yes**
Judge 3	**No**	**No**	**Yes**	**Yes**
PBR	**No**	**No**	**No**	**No**

Winner Determination

- How hard is it to **compute the outcome** of an aggregation rule?
- **Fact:** Winner determination for quota-based rules is in P.
- **Fact:** Winner determination for premise-based rule is in P.
    - Proof: counting (for premises) + model checking (for conclusions)
- **Theorem** (Endriss et al., 2012): The winner determination problems of the **generalized Kemeny rule** and the **generalized Slater rule** are both NP-hard.
    - Proof: reduction from corresponding voting problem
- Efficiently computable variants of distance-based rules: **representative-voter rules**
    - only search through judgment sets proposed by individuals
    - find "most representative voter"

Strategic Manipulation

- Do aggregation rules **incentivise** judges to report their judgments **truthfully**?
  - what are the **preferences** of a judge?
- **Example:** Hamming preferences and the PBR
  - **Hamming preferences:** prefer outcomes closer to own judgment set

	p	q	r	$p \vee q$	$p \vee r$
Judge 1	**No**	**No**	**No**	**No**	**No**
Judge 2	**Yes**	**No**	**No**	**Yes**	**Yes**
Judge 3	**No**	**Yes**	**Yes**	**Yes**	**Yes**

Preventing Manipulation

Sometimes manipulation is impossible.

- **Proposition** (Dietrich and List, 2007): An aggregation rule is **immune to manipulation** if and only if it is both **independent** and **monotonic**.
  - but: independent and monotonic rules are not very attractive…

Sometimes it is hard.

- **Proposition** (Endriss et al., 2012): **Manipulating** the premise-based rule with Hamming preferences is **NP-complete.**
  - but: this is only a worst-case result…

Summary

- JA provides a framework to aggregate binary opinions
- Motivated by **paradoxes** and **impossibilities**
- Concrete **aggregation rules** violate at least some **axioms**
- **Computational** and **strategic** considerations are important
- Further reading:
  - U. Endriss. Judgment aggregation. In *Handbook of Computational Social Choice*, chapter 17. Cambridge University Press, 2015.
  - C. List. The Theory of Judgment Aggregation: An Introductory Review. *Synthese*, 187(1):179–207, 2012.
  - U. Endriss, U. Grandi, and D. Porello. Complexity of Judgment Aggregation. *Journal of Artificial Intelligence Research*, 45:481–514, 2012.
- **Discussion:** Which rules/axioms are appealing? How is this material relevant to our goal of deriving societal tradeoffs? …