Tutorial on Judgment Aggregation

Ulle Endriss
Institute for Logic, Language and Computation
University of Amsterdam

[http://www.illc.uva.nl/~ulle/teaching/wine-2012/]

Example

Three judges have to decide whether the defendant is *guilty* (q). Relevant premises are whether those *fingerprints* are his (p) and whether that would be *sufficient evidence* for a conviction ($p \rightarrow q$).

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

What should be their *collective decision*?

Purpose of this Tutorial

Judgment aggregation (JA) is a rich modelling tool for reasoning about collective decision making. The basic ideas originate in *Legal Theory* and have been developed in *Philosophy*, *Economics*, and *Logic*.

Recently, people in *Computer Science* and *AI* also got interested, but so far work of an algorithmic nature has been very limited.

My goal today is

- to provide a "classical" introduction to JA in some detail and
- to provide pointers to the little algorithmic work there is.

Outline

- Doctrinal Paradox (Kornhauser and Sager, 1993)
- Formal framework for judgment aggregation
- Examples for concrete aggregation procedures
- Examples for axioms (desirable properties of procedures)
- Axiomatic characterisation of aggregation procedures
- Basic Impossibility Theorem (List and Pettit, 2002)
- Ways of circumventing the impossibility
- Agenda characterisation results: possibility and safety theorems
- Complexity of judgment aggregation
- Pointers to the literature

The Doctrinal Paradox

Consider a court with three judges. Suppose legal doctrine stipulates that the defendant is *liable* *iff* there has been a *valid* contract (p) and that contract has been *breached* (q). So we need to worry about $p \wedge q$.

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

Paradox: Taking majority decisions on the *premises* (p and q) and then inferring the conclusion ($p \wedge q$) gives a different result from taking a majority decision on the *conclusion* ($p \wedge q$) directly. Also: individual judgement sets are *consistent*, but the collective judgment set obtained by *majority* is not.

L.A. Kornhauser and L.G. Sager. The One and the Many: Adjudication in Collegial Courts. *California Law Review*, 81(1):1–59, 1993.

Formal Framework

<u>Notation:</u> Let $\sim\varphi := \varphi'$ if $\varphi = \neg\varphi'$ and let $\sim\varphi := \neg\varphi$ otherwise.

An *agenda* Φ is a finite nonempty set of propositional formulas (w/o double negation) closed under complementation: $\varphi \in \Phi \;\Rightarrow\; \sim\varphi \in \Phi$.

A *judgment set* J on an agenda Φ is a subset of Φ. We call J:

- *complete* if $\varphi \in J$ or $\sim\varphi \in J$ for all $\varphi \in \Phi$
- *complement-free* if $\varphi \notin J$ or $\sim\varphi \notin J$ for all $\varphi \in \Phi$
- *consistent* if there exists an assignment satisfying all $\varphi \in J$

Let $\mathcal{J}(\Phi)$ be the set of all complete and consistent subsets of Φ.

Now a finite set of *individuals* $\mathcal{N} = \{1, \ldots, n\}$, with $n \geqslant 2$, express judgments on the formulas in Φ, producing a *profile* $\boldsymbol{J} = (J_1, \ldots, J_n)$.

An *aggregation procedure* for an agenda Φ and a set of n individuals is a function mapping a profile of complete and consistent individual judgment sets to a single collective judgment set: $F : \mathcal{J}(\Phi)^n \to 2^\Phi$.

Outcome-Related Properties of Aggregators

We extend the concepts of completeness, complement-freeness, and consistency of *judgment sets* to properties of *aggregators* F:

- F is *complete* if $F(\boldsymbol{J})$ is complete for any $\boldsymbol{J} \in \mathcal{J}(\Phi)^n$
- F is *complement-free* if $F(\boldsymbol{J})$ is c.-f. for any $\boldsymbol{J} \in \mathcal{J}(\Phi)^n$
- F is *consistent* if $F(\boldsymbol{J})$ is consistent for any $\boldsymbol{J} \in \mathcal{J}(\Phi)^n$

Only consistency involves logic *proper*. Complement-freeness and completeness are purely syntactic concepts, not involving any model-theoretic ideas (they are also computationally easy to check).

F is called *collectively rational* if it is both complete and consistent (and thus also complement-free).

Aggregation Procedures

Ideas that come to mind for how to design an aggregation procedure:

- *Majority rule:* accept φ if a strict majority does (natural choice, but we have already seen that this does not preserve consistency)
- *Quota rules:* accept φ if at least, say, $\geqslant \frac{2}{3}$ of the individuals do
- *Premise-based rule:* decide on "premises" (maybe literals?) by majority; then logically *infer* truth values for "conclusions"
- *Distance-based approach:* define a notion of *distance* between judgment sets and choose an outcome that minimises, say, the *sum* of distances to the individual judgment sets
- *Average-voter rule:* identify the "most representative" individual and copy her judgment set

How to choose? The *axiomatic method* can help to make various *normative desiderata* precise . . .

Axioms

What makes for a "good" aggregation procedure F? The following axioms all express intuitively appealing properties:

- *Unanimity:* if $\varphi \in J_i$ for all i, then $\varphi \in F(\boldsymbol{J})$.
- *Anonymity:* for any profile $\boldsymbol{J}$ and any permutation $\pi : \mathcal{N} \to \mathcal{N}$ we have $F(J_1, \ldots, J_n) = F(J_{\pi(1)}, \ldots, J_{\pi(n)})$.
- *Neutrality:* for any φ, ψ in the agenda Φ and any profile $\boldsymbol{J}$, if for all i we have $\varphi \in J_i \Leftrightarrow \psi \in J_i$, then $\varphi \in F(\boldsymbol{J}) \Leftrightarrow \psi \in F(\boldsymbol{J})$.
- *Independence:* for any φ in the agenda Φ and any profiles $\boldsymbol{J}$ and $\boldsymbol{J'}$, if $\varphi \in J_i \Leftrightarrow \varphi \in J'_i$ for all i, then $\varphi \in F(\boldsymbol{J}) \Leftrightarrow \varphi \in F(\boldsymbol{J'})$.
- *Systematicity* = neutrality + independence
- *Monotonicity:* for any $\varphi \in \Phi$ and profiles $\boldsymbol{J}, \boldsymbol{J'}$, if $\varphi \in J'_{i^\star} \backslash J_{i^\star}$ for some $i^\star$ and $J_i = J'_i$ for all $i \neq i^\star$, then $\varphi \in F(\boldsymbol{J}) \Rightarrow \varphi \in F(\boldsymbol{J'})$.

(Note that the majority rule satisfies all of these axioms.)

Winning Coalitions

Notation: Let $N_\varphi^{\boldsymbol{J}}$ be the set of individuals accepting φ in profile $\boldsymbol{J}$.

An alternative way of interpreting independence:

- F is *independent* iff there exists a family of *winning coalitions* $\mathcal{W}_\varphi \subseteq 2^{\mathcal{N}}$, one for each $\varphi \in \Phi$, such that $\varphi \in F(\boldsymbol{J}) \Leftrightarrow N_\varphi^{\boldsymbol{J}} \in \mathcal{W}_\varphi$.

Suppose F is independent. Then:

- If F is *unanimous*, then $\mathcal{N} \in \mathcal{W}_\varphi$ for any formula $\varphi \in \Phi$.
- If F is *neutral*, then $\mathcal{W}_\varphi = \mathcal{W}_\psi$ for any formulas $\varphi, \psi \in \Phi$.
- If F is *anonymous*, then $C \in \mathcal{W}_\varphi \Rightarrow C' \in \mathcal{W}_\varphi$ for $|C| = |C'|$.
- If F is *monotonic*, then $C \in \mathcal{W}_\varphi \Rightarrow C' \in \mathcal{W}_\varphi$ for $C \subseteq C'$.

We are now ready to prove some simple *characterisation results* . . .

Quota Rules

A *quota rule* F_q is defined by a function $q : \Phi \to \{0, 1, \ldots, n{+}1\}$:

$$F_q(\boldsymbol{J}) \quad = \quad \{\varphi \in \Phi \mid |N_\varphi^{\boldsymbol{J}}| \geqslant q(\varphi)\}$$

A quota rule F_q is called *uniform* if q maps any given formula to the same number k. Examples:

- The *unanimous rule* F_n accepts φ *iff* everyone does.
- The *constant rule* F_0 (F_{n+1}) accepts all (no) formulas.
- The *(strict) majority rule* F_{maj} is the quota rule with $q = \lceil \frac{n+1}{2} \rceil$.
- The *weak majority rule* is the quota rule with $q = \lceil \frac{n}{2} \rceil$.

Observe that for *odd* n the majority rule and the weak majority rule coincide. For *even* n they differ (and only the weak one is complete).

Characterisation of Quota Rules

Proposition 1 (Dietrich and List, 2007) *An aggregation procedure is anonymous, independent and monotonic iff it is a quota rule.*

Proof: Clearly, any quota rule has these properties (right-to-left).

For the other direction (proof sketch):

- Independence means that acceptance of φ only depends on the coalition $N_\varphi^{\boldsymbol{J}}$ accepting it.
- Anonymity means that it only depends on the cardinality of $N_\varphi^{\boldsymbol{J}}$.
- Monotonicity means that acceptance of φ cannot turn to rejection as additional individuals accept φ.

Hence, it must be a quota rule. ✓

F. Dietrich and C. List. Judgment Aggregation by Quota Rules: Majority Voting Generalized. *Journal of Theoretical Politics*, 19(4)391–424, 2007.

More Characterisations

Clearly, a quota rule F_q is uniform *iff* it is neutral. Thus:

Corollary 2 *An aggregation procedure is anonymous, neutral, independent and monotonic (= ANIM) iff it is a uniform quota rule.*

Now consider a uniform quota rule F_q with quota q. Two observations:

- For F_q to be *complete*, we need $q \leqslant \max_{0 \leqslant x \leqslant n}(x, n-x) \Rightarrow q \leqslant \lceil \frac{n}{2} \rceil$.
- For F_q to be *compl.-free*, we need $q > \min_{0 \leqslant x \leqslant n}(x, n-x) \Rightarrow q > \lfloor \frac{n}{2} \rfloor$.

For n *even*, no such q exists. Thus:

Proposition 3 *For n even, no aggregation procedure is ANIM, complete and complement-free.*

For n *odd*, such a q does exist, namely $q = \lceil \frac{n}{2} \rceil = \frac{n+1}{2}$. Thus:

Proposition 4 *For n odd, an aggregation procedure is ANIM, complete and complement-free iff it is the (strict) majority rule.*

Impossibility Theorem

We have seen that the majority rule is *not consistent*. Is there another "reasonable" aggregation procedure that does not have this problem?

Theorem 5 (List and Pettit, 2002) *No judgment aggregation procedure for an agenda Φ with $\{p, q, p \wedge q\} \subseteq \Phi$ that satisfies anonymity, neutrality, and independence is complete and consistent.*

Remark 1: Note that the theorem requires $|\mathcal{N}| > 1$.

Remark 2: Similar impossibilities arise for other agendas with some minimal structural richness. To be discussed later on.

C. List and P. Pettit. Aggregating Sets of Judgments: An Impossibility Result. *Economics and Philosophy*, 18(1):89–110, 2002.

Proof

From anonymity, neutrality and independence: collective acceptance of φ can only depend on the *number* $\#[\varphi]$ of individuals accepting φ.

- Case where the number n of individuals is *even:*

  Consider a scenario where $\#[p] = \#[\neg p]$.

  As argued above, we need to accept either both or neither:

  – Accepting both contradicts consistency. ✓

  – Accepting neither contradicts completeness. ✓

- Case where the number n of individuals is *odd* (and $n > 1$):

  Consider a scenario where $\frac{n-1}{2}$ accept p and q; 1 each accept exactly one of p and q; and $\frac{n-3}{2}$ accept neither p nor q.

  That is: $\#[p] = \#[q] = \#[\neg(p \wedge q)]$. But:

  – Accepting all three formulas contradicts consistency. ✓

  – But if we accept none, completeness forces us to accept their complements, which also contradicts consistency. ✓

What next?

List and Pettit's impossibility theorem raises (at least) two questions:

- Which of the assumptions in the theorem should we relax to turn the impossibility into a possibility?
- Maybe a result that is specific to agendas including $\{p, q, p \wedge q\}$ is not satisfactory. Can we do better?

Circumventing the Impossibility

If we are prepared to relax some of the assumptions of Theorem 5, we may be able to circumvent the impossibility and successfully aggregate judgements. Next, we will explore some such possibilities:

- Relaxing the *input* conditions: drop the (implicit) *universal domain* assumption and design rules for restricted domains [not today]
- Relaxing the *output* conditions: drop the *completeness* requirement (dropping consistency works but is unattractive)
- Giving up *anonymity*: dictatorships will surely work, but maybe we can do a little better than that
- Weakening *neutrality* (which may after all be inappropriate for logical propositions?)
- Weakening *independence* (which is known to be the source of trouble also in other areas of social choice theory)

Quota Rules with a High Quota

We may *drop completeness* from our list of requirements.

Clearly, a (uniform) quota rule with a sufficiently high quota will be consistent. Dietrich and List (2007) give lower bounds for the quota to ensure consistency as a function of n and the size of the largest *minimally inconsistent subset* of the agenda Φ. Example:

> Let $\Phi = \{p, \neg p, q, \neg q, p \wedge q, \neg(p \wedge q)\}$. The largest mi-subset is $\{p, q, \neg(p \wedge q)\}$. Any quota $> \frac{2}{3}$ will ensure consistency.

F. Dietrich and C. List. Judgment Aggregation by Quota Rules: Majority Voting Generalized. *Journal of Theoretical Politics*, 19(4):391–424, 2007.

Oligarchic Rules

As we have seen, quota rules with sufficiently high quotas can circumvent the impossibility, if we are prepared to give up completeness.

Instead, we may try replacing completeness by *deductive closure:*
$\varphi \in \Phi$ and $J \models \varphi$ imply $\varphi \in J$ for the (collective) judgement set J

The *oligarchic rule* for the set of individuals $X \subseteq \mathcal{N}$ is the rule that accepts φ *iff* everyone in X does. Special cases:

- dictatorial rule: $|X| = 1$
- unanimous rule: $|X| = n$

It is easy to check that any oligarchic rule satisfies:

- *consistency* and *deductive closure* (if individuals do);
- the *universal domain* assumption, *neutrality*, and *independence*;
- but *not anonymity* (unless $|X| = n$) *nor completeness* (unless $|X| = 1$).

Gärdenfors (2006) gives a more precise axiomatic characterisation.

P. Gärdenfors. A Representation Theorem for Voting with Logical Consequences. *Economics and Philosophy*, 22(2):181–190, 2006.

The Premise-Based Procedure

Another option is to sacrifice *neutrality* ...

Suppose we *can* divide the agenda into *premises* and *conclusions:*

$$\Phi \quad = \quad \Phi_p \uplus \Phi_c$$

The *premise-based procedure* PBP for Φ_p and Φ_c is this function:

$$\begin{aligned} \mathrm{PBP}(\boldsymbol{J}) \quad = \quad & \Delta \cup \{\varphi \in \Phi_c \mid \Delta \models \varphi\}, \\ & \text{where } \Delta = \{\varphi \in \Phi_p \mid |\{i \mid \varphi \in J_i\}| > \frac{n}{2}\} \end{aligned}$$

If we assume that

- the set of premises is the set of literals in the agenda,
- the agenda Φ is is closed under propositional letters, and
- the number n of individuals is odd,

then $\mathrm{PBP}(\boldsymbol{J})$ will always be *consistent* and *complete*.

Example

Note that the premise-based procedure violates *unanimity*:

	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
PBP:	No	No	No	No

Distance-Based Procedures

Finally, we might be willing to sacrifice *independence* ...

The *distance-based procedure* is defined as follows:

$$\mathrm{DBP}(\boldsymbol{J}) \quad = \quad \operatorname*{arg\,min}_{J \in \mathcal{J}(\Phi)} \sum_{i=1}^{n} H(J, J_i)$$

Here the *Hamming distance* $H(J, J')$ between judgment sets J and J' is the number of positive agenda formulas on which they differ.

Remark: The DBP may return a set of tied winners.

The DBP behaves like the majority rule in case that is consistent, and makes a "reasonable" (consistent) choice otherwise. Variants are possible.

G. Pigozzi. Belief Merging and the Discursive Dilemma: An Argument-based Account of Paradoxes of Judgment. *Synthese*, 152(2):285–298, 2006.

M.K. Miller and D. Osherson. Methods for Distance-based Judgment Aggregation. *Social Choice and Welfare*, 32(4):575–601, 2009.

Agenda Characterisations

- The *impossibility result* we have seen showed that consistent aggregation is impossible under certain assumptions—but only for agendas including $\{p, q, p \wedge q\}$.

  Instead we might ask: *for which agendas* does this impossibility arise? That is, we are after *agenda charcterisations*.

- Recall that we have seen several *characterisation results* already (for quota rules). They only use *choice-theoretic axioms* (independence, etc.) and *syntactic conditions* on the outcome (completeness and complement-freeness). No logic so far.

Safety of the Agenda under Majority Voting

Previously we saw that the majority rule can produce an inconsistent outcome for *some* (not all) profiles based on agendas $\Phi \supseteq \{p, q, p \wedge q\}$. How can we *characterise* the class of agendas with this problem?

An *agenda* Φ is said to be *safe* for an aggregation procedure F if the *outcome* $F(\boldsymbol{J})$ is *consistent* for *any* admissible profile $\boldsymbol{J} \in \mathcal{J}(\Phi)^n$.

Theorem 6 (Nehring and Puppe, 2007) *An agenda Φ is safe for the (strict) majority rule iff Φ has the median property $[$for $|\mathcal{N}| \geqslant 3]$.*

A set of formulas Φ satisfies the *median property* if every inconsistent subset of Φ does itself have an inconsistent subset of size $\leqslant 2$.

K. Nehring and C. Puppe. The Structure of Strategy-proof Social Choice. Part I: General Characterization and Possibility Results on Median Space. *Journal of Economic Theory*, 135(1):269–305, 2007.

Proof

Claim: Φ is safe [$F_{\text{maj}}(\boldsymbol{J})$ is consistent] $\Leftrightarrow$ Φ has the median property

($\Leftarrow$) Let Φ be an agenda with the median property. Now assume that there exists an admissible profile $\boldsymbol{J}$ such that $F_{\text{maj}}(\boldsymbol{J})$ is *not* consistent.

- $\rightsquigarrow$ There exists an inconsistent set $\{\varphi, \psi\} \subseteq F_{\text{maj}}(\boldsymbol{J})$.
- $\rightsquigarrow$ Each of φ and ψ must have been accepted by a strict majority.
- $\rightsquigarrow$ One individual must have accepted both φ and ψ.
- $\rightsquigarrow$ Contradiction (individual judgment sets must be consistent). ✓

($\Rightarrow$) Let Φ be an agenda that violates the median property, i.e., there exists a minimally inconsistent set $\Delta = \{\varphi_1, \ldots, \varphi_k\} \subseteq \Phi$ with $k > 2$.

For simplicity, suppose n (the number of individuals) is divisible by 3. There exists a consistent profile $\boldsymbol{J}$ under which individual i accepts all formulas in Δ except for $\varphi_{1+(i \bmod 3)}$. But then the majority rule will accept all formulas in Δ, i.e., $F_{\text{maj}}(\boldsymbol{J})$ is inconsistent. ✓

Agenda Characterisation for Classes of Rules

Now instead of a single aggregator, suppose we are interested in a *class of aggregators*, possibly determined by a set of *axioms*.

We might ask:

- *Possibility:* Does there exist an aggregator meeting certain axioms that will be consistent for any agenda with a given property?
- *Safety:* Will every aggregator meeting certain axioms be consistent for any agenda with a given property?

Possibility Theorem for Monotonic Rules

Theorem 7 (Nehring and Puppe, 2010) *There exists a unanimous, anonymous, independent and monotonic aggregation procedure for the agenda Φ that is collectively rational iff Φ is not blocked.*

Here an agenda Φ is called *blocked* if there exists a $\varphi \in \Phi$ with a *conditional path* from φ to $\sim\varphi$ and *vice versa*, where a conditional path is a sequence $\varphi_1, \varphi_2, \ldots, \varphi_k$ such that $\varphi_i \neq \sim\varphi_{i+1}$ and $\{\varphi_i, \sim\varphi_{i+1}\}$ is part of some mi-subset of Φ for every $i < k$.

Proof: Omitted.

List and Puppe (2009) give an overview of known possibility theorems.

K. Nehring and C. Puppe. Abstract Arrovian Aggregation. *Journal of Economic Theory*, 145(2):467–494, 2010.

C. List and C. Puppe. Judgment Aggregation: A Survey. In: *Handbook of Rational and Social Choice*, Oxford University Press, 2009.

Safety Theorem for Systematic Rules

Suppose we know that the group will use *some* aggregation procedure meeting certain requirements, but we do not know which procedure exactly. Can we guarantee that the outcome will be consistent?

A typical result (for the majority rule axioms, minus monotonicity):

Theorem 8 (Endriss et al., 2010) *An agenda Φ is safe for any anonymous, neutral, independent, complete and complement-free aggregation procedure iff Φ has the simplified median property .*

An agenda Φ has the *simplified median property* if every inconsistent subset of Φ has itself an inconsistent subset $\{\varphi, \psi\}$ with $\models \varphi \leftrightarrow \neg\psi$.

Note: This is more restrictive than the median property: $\{\neg p, p \wedge q\}$.

U. Endriss, U. Grandi and D. Porello. Complexity of Judgment Aggregation: Safety of the Agenda. Proc. AAMAS-2010.

Proof

Claim: Φ is safe for any ANI/complete/comp-free rule $F \Leftrightarrow \Phi$ has SMP

($\Leftarrow$) Suppose Φ has the SMP. For the sake of contradiction, assume $F(\boldsymbol{J})$ is inconsistent. Then $\{\varphi, \psi\} \subseteq F(\boldsymbol{J})$ with $\models \varphi \leftrightarrow \neg\psi$. Now:

- $\rightsquigarrow$ $\varphi \in J_i \Leftrightarrow {\sim}\psi \in J_i$ for each individual i (from $\models \varphi \leftrightarrow \neg\psi$ together with consistency and completeness of individual judgment sets)
- $\rightsquigarrow$ $\varphi \in F(\boldsymbol{J}) \Leftrightarrow {\sim}\psi \in F(\boldsymbol{J})$ (from neutrality)
- $\rightsquigarrow$ both ψ and ${\sim}\psi$ in $F(\boldsymbol{J}) \rightsquigarrow$ contradiction (with complement-freeness) ✓

($\Rightarrow$) Suppose Φ violates the SMP. Take any minimally inconsistent $\Delta \subseteq \Phi$. If $|\Delta| > 2$, then also the MP is violated and we already know that the majority rule is not consistent. ✓ So can assume $\Delta = \{\varphi, \psi\}$.

W.l.o.g., must have $\varphi \models \neg\psi$ but $\neg\psi \not\models \varphi$ (otherwise SMP holds).

But now we can find a rule F that is not safe: accept a formula if at most one individual does and take a profile with $J_1 = \{{\sim}\varphi, {\sim}\psi, \ldots\}$, $J_2 = \{{\sim}\varphi, \psi, \ldots\}$, and $J_3 = \{\varphi, {\sim}\psi, \ldots\}$. Then $F(\boldsymbol{J}) = \{\varphi, \psi, \ldots\}$. ✓

Comparing Possibility and Safety Results

Possibility theorems and safety theorems are closely related:

- Possibility: *some* aggregator in the class determined by the given axioms will produce consistent outcomes *iff* the agenda has a given property
- Safety: *all* aggregators in the class determined by the given axioms will produce consistent outcomes *iff* the agenda has a given property

In what situations do we need these results?

- Possibility: a mechanism designer wants to know whether she can design an aggregation rule meeting a given list of requirements
- Safety: a system might know certain properties of the aggregator users will employ (but not all properties) and we want to be sure there won't be any problem (we might want to check this again and again)

For safety problems in particular we might want to develop *algorithms*, i.e., *complexity* plays a role.

Complexity of Safety of the Agenda

Deciding whether a given agenda is safe for the majority rule (as well as several classes of rules we get by relaxing the axioms defining the majority rule) is located at the second level of the polynomial hierarchy.

Proving those results involves the following lemma (and variations):

Lemma 9 (Endriss et al., 2010) *Deciding whether a given agenda has the median property is Π_2^p-complete.*

Proof: Omitted.

Recall that $\Pi_2^p = \mathrm{coNP}^{\mathrm{NP}}$ is the class of problems for which we can verify a certificate for a negative answer in polynomial time if we have access to an NP oracle. A typical problem in the class is deciding truth of formulas of the form $\forall \boldsymbol{x} \exists \boldsymbol{y} \varphi(\boldsymbol{x}, \boldsymbol{y})$. So: very hard.

U. Endriss, U. Grandi and D. Porello. Complexity of Judgment Aggregation: Safety of the Agenda. Proc. AAMAS-2010.

Algorithmic Work in Judgment Aggregation

Endriss, Grandi, Porello (JAIR'12, subsuming AAMAS'10/COMSOC'10): *complexity* of *winner determination*, *manipulation*, and *safety*.

Baumeister, Erdélyi, Erdélyi, Rothe (ADT'11, COMSOC'12): *complexity* of *bribery* and *control*.

Nehama (WINE'11): *approximate JA* (goal is to make inconsistent outcomes unlikely).

Applications in AI: *multiagent systems* (Slavkovik, PhD'12); *abstract argumentation* (Rahwan & Thomé, AAMAS'10, Caminada & Pigozzi, JAAMAS'11). See also: *belief merging* (Konieczny et al., since ~2002).

Recent trend in JA to focus on *design of aggregation rules*. One such paper by computer scientists: Lang, Pigozzi, Slavkovik, van der Torre (TARK'11).

Further Reading

Possible starting points:

- Easy-going tutorial paper on basic JA:

  C. List. The Theory of Judgment Aggregation: An Introductory Review. *Synthese*, 187(1):179–207, 2012.

- More advanced tutorial paper focussing on impossibilities:

  D. Grossi and G. Pigozzi. Introduction to Judgment Aggregation. In: *Lectures on Logic and Computation*, Springer, 2012.

- Paper on complexity of JA:

  U. Endriss, U. Grandi, and D. Porello. Complexity of Judgment Aggregation. *Journal of Artificial Intell. Research*, 45:481–514, 2012.

Or visit the website for my Amsterdam course on computational social choice for more material:

`http://www.illc.uva.nl/~ulle/teaching/comsoc/2012/`

Last Slide

This has been an introduction to judgment aggregation.

We have seen:

- paradox, formalisation, axioms, aggregation procedures
- characterisation results: axioms identifying procedures
- impossibility result: axioms precluding consistency
- circumventing the impossibility by weakening requirements
- agenda characterisation: possibility and safety results
- algorithmic questions, particularly complexity concerns

Take Home Message: This is a young field (only 10 years old). Algorithmic analysis and exploitation for applications in computer science have only just begun. Still plenty of opportunities!