Judgment Aggregation

Umberto Grandi
IRIT, University of Toulouse

COST IC1205 Summer School on Computational Social Choice
San Sebastian, 18-22 July 2016

A game!

A new area has been recently acquired by your home town. After a first round of discussion during which all parties have made their proposal for the destination of this area, the local council needs to take a decision. You are one of the three parties equally represented in the council: the Pink party, the Yellow party and the Green party.

- Three parties with a different goal each
- Three projects proposed: Park (PAR), House (HOU), Hospital (HOS)
- At least one project needs to be approved and funded

Now play! (see blackboard)

History I: Some scattered observations

Looking backward it is possible to find a number researchers identifying problems when voting on multiple inter-connected issues:

- In 1837, Poisson obviously had seen it coming (from a paper by Elster)
- In 1921, Italian legal theorist Vacca recognises a problem when studying collective decisions on three correlated issues
- In 1952, Guilbaud published a paper on the logical interpretation of preference aggregation, generalising the Condorcet's paradox
- In 1975, Wilson develops a complex generalisation of Arrow's theorem based on binary vectors, without generalising further than preferences

J. Elster. Excessive Ambitions (ii). *Capitalism and Society*, 2013.

R. Vacca. Opinioni individuali e deliberazioni collettive. *Rivista Internazionale di Filosofia del Diritto*, 1921.

G. T. Guilbaud. Les théories de l'intérêt général et le problème logique de l'agrégation. *Economie appliqué*, 1952.

R. Wilson. On the theory of aggregation. *Journal of Economic Theory*, 1975.

History II: The discursive dilemma

The turning point is when L. A. Kornhauser and L. G. Sager discovered the following situation happening in an American court:

There is a court with three judges. Suppose legal doctrine stipulates that the defendant is liable if and only if there has been a valid contract and that contract has been breached. The judgment is made by majority.

Doctrinal Paradox

	Valid?	Breached?	Liable?
Judge 1:	Yes	Yes	Yes
Judge 2:	No	Yes	No
Judge 3:	Yes	No	No
Majority:	Yes	Yes	No

Should we vote on the premises ("valid and breach") and claim the defendant guilty or on the conclusion ("liable") and claim her innocent?

L. A. Kornhauser and L. G. Sager, Unpacking the court. *Yale Law Journal*, 1986.

L. A. Kornhauser and L. G. Sager, The one and the many: Adjudication in Collegial Courts. *California Law Review*, 1993.

History III: Judgment aggregation and social choice theory

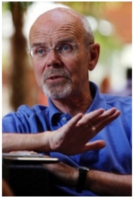  

In 2001, Philip Pettit generalised Kornhauser and Sager's observation, and in 2004 List and Pettit published the first systematic study of the aggregation of judgements and the first impossibility theorem.

Since then:

- 20+ papers by Dietrich and List on judgement aggregation
- *5th Social Choice and Welfare Prize* to Dietrich and List in 2010

P. Pettit. Deliberative Democracy and the Discursive Dilemma. *Philosophical Issues*, 2001.

C. List and P. Pettit. Aggregating Sets of Judgments: An Impossibility Result. *Economics and Philosophy*, 2002.

History IV: Judgement aggregation and computer science

Social choice recently became the paradigm for studying collective decision-making in multiagent systems and AI in general, raising among other things the problem of combinatorial voting:

- multiple referenda (yes/no to a number of questions)
- product configuration (product space of variable domains)
- separable preferences and multiple election paradox

To which we add the classical use of propositional logic in knowledge representation and AI in general:

- a compact way of expressing information (beliefs, knowledge bases...)
- belief merging is a similar framework to that of judgement aggregation

Algorithmic problems in JA had studied, plenty of problems to investigate!

D. Grossi and G. Pigozzi. Judgment Aggregation: A Primer. Morgan and Claypool, 2014.

U. Endriss. Judgment Aggregation. In *Handbook of Computational Social Choice*, 2016.

Outline of the course

- Part I: Paradoxes of aggregation
- Minitutorial on computational complexity
- Part II: Aggregation rules and winner determination
- Part III: Impossibility results and winning coalitions
- Part IV: Manipulation
- Part V: “Applications”

Part I:
Paradoxes of Aggregation

What is a paradox (of aggregation)?

Condorcet Paradox

$$\triangle \succ_1 \bigcirc \succ_1 \square$$
$$\square \succ_2 \triangle \succ_2 \bigcirc$$
$$\bigcirc \succ_3 \square \succ_3 \triangle$$

$$\triangle \succ \bigcirc \succ \square \succ \triangle$$

Discursive Dilemma

	Valid?	Breached?	Liable?
Judge 1	Y	Y	Y
Judge 2	N	Y	N
Judge 3	Y	N	N
Majority	Y	Y	N

To analyse the common structure of paradoxes of aggregation we need:

- A general framework for the study of aggregation problems: binary aggregation
- An explicit representation of rationality assumptions: integrity constraints

Disclaimer: The framework I will be using is equivalent to the classic formula-based JA, and the two settings will be used interchangeably.

Binary Aggregation

Ingredients:

- A finite set $\mathcal{N} = \{1, ..., n\}$ of an odd number of individuals
- A finite set $\mathcal{I} = \{1, \ldots, m\}$ of issues
- The space of all possible opinions/judgment is $D = \{0, 1\}^{\mathcal{I}}$

Definition

An aggregation procedure is a function $F : \mathcal{D}^{\mathcal{N}} \to \mathcal{D}$ mapping each profile of ballots $\boldsymbol{B} = (B_1, \ldots, B_n)$ to an element of the domain $\mathcal{D}$.

Example: Austerity measures

- $\mathcal{N} = \{1, 2, 3\}$
- $\mathcal{I} = \{\text{Cut pensions}, \text{Increase taxes}, \text{Stimulate growth}\}$
- Individuals submit ballots in $\mathcal{D} = \{0, 1\}^3$

$B_1 = (1, 0, 1)$ the first individual is probably young.

Wilson (1975), Dokow and Holzman (2008), Grandi and Endriss (AAAI-2010, IJCAI-2011)

An important assumption

For ease of presentation, $|\mathcal{N}|$ is odd.

Propositional Logic

Start from a number of propositional variables $PS = \{p_1, \ldots, p_m\}$

Construct formulas inductively:

- a propositional variable in *PS* is a formula
- if φ is a formula then $\neg\varphi$ is a formula
- if φ and ψ are formulas then $\varphi \vee \psi$ are formulas

Define $\varphi \wedge \psi := \neg(\neg\varphi \vee \neg\psi)$, and $\varphi \to \psi := \neg\varphi \vee \psi$.

An assignment over variables in *PS* is a function $\rho : PS \to \{0, 1\}$ (=a ballot).
An assignment ρ is a model of a formula φ if:

- $\rho \models p$ iff $\rho(p) = 1$
- $\rho \models \neg\varphi$ iff $\rho \not\models \varphi$
- $\rho \models \varphi \vee \psi$ iff $\rho \models \varphi$ or $\rho \models \psi$

We denote with $\mathrm{Mod}(\varphi)$ all the models of a formula φ. For example
$\mathrm{Mod}((p \vee q) \to r) = \{(1,0,1), (0,1,1), (1,1,1), (0,0,1), (0,0,0)\}$

Integrity Constraints

Integrity constraints are propositional formulas
that express subsets of $\mathcal{D} = \{0,1\}^{\mathcal{I}}$ as admissible ballots.

- One propositional symbol for every issue: $PS = \{p_1, \ldots, p_m\}$
- $\mathcal{L}_{PS}$ closing under connectives $\wedge$, $\vee$,$\neg$, $\rightarrow$ the set of atoms PS

Given an integrity constraint $\mathrm{IC} \in \mathcal{L}_{PS}$, an admissible ballot is $B \in \mathrm{Mod}(\mathrm{IC})$

Example: Austerity Measures

Economists say that $\mathrm{IC} = (p_P \wedge p_I) \rightarrow p_G$

Individual 1 submits $B_1 = (1,0,0)$: B_1 satisfies IC ✓
Individual 2 submits $B_2 = (0,1,0)$: $B_2 \models \mathrm{IC}$ ✓
Individual 3 submits $B_3 = (1,1,1)$: $B_3 \models \mathrm{IC}$ ✓

Majority aggregation outputs $(1,1,0)$: IC not satisfied

Paradoxes of Aggregation

Every individual satisfies the same integrity constraint IC...
...what about the collective outcome?

Definition

A paradox is a triple $(F, \boldsymbol{B}, IC)$, where:

- *F is an aggregation procedure*
- *$\boldsymbol{B} = (B_1, \ldots, B_n)$ a profile*
- *$IC \in \mathcal{L}_{PS}$ an integrity constraint*

such that $B_i \models IC$ for all $i \in \mathcal{N}$ but $F(\boldsymbol{B}) \not\models IC$.

F is collectively rational wrt. IC if there is no $\boldsymbol{B}$ s.t. $(F, \boldsymbol{B}, \mathrm{IC})$ is a paradox.

Preference Aggregation and Binary Aggregation

Linear order $<$ over alternatives $\mathcal{X}$	$\Leftrightarrow$	Ballot $B_{\leqslant}$ over issues $\mathcal{I} = \{ab \mid a \neq b \in \mathcal{X}\}$
ex: $b > c > a$		*ex*: $(0, 1, 1, 1, 0, 0, 0, 0, 0)$

IC$_<$ encodes the integrity constraint for linear orders:

Irreflexivity: $\neg p_{aa}$ for all $a \in \mathcal{X}$

Completeness: $p_{ab} \vee p_{ba}$ for all $a \neq b \in \mathcal{X}$

Transitivity: $p_{ab} \wedge p_{bc} \rightarrow p_{ac}$ for $a, b, c \in \mathcal{X}$ pairwise distinct

The Condorcet Paradox

In 1785 Monsieur le Marquis de Condorcet pointed out that...

$$
\begin{array}{c}
\triangle \succ_{\mathbf{1}} \bigcirc \succ_{\mathbf{1}} \square \\
\square \succ_{\mathbf{2}} \triangle \succ_{\mathbf{2}} \bigcirc \\
\bigcirc \succ_{\mathbf{3}} \square \succ_{\mathbf{3}} \triangle \\
\hline
\triangle \succ \bigcirc \succ \square \succ \triangle
\end{array}
$$

...pairwise majority aggregates linear order into cyclic preferences.

Condorcet Paradox Revisited

	$\triangle\bigcirc$	$\bigcirc\square$	$\triangle\square$
Agent 1	1	1	1
Agent 2	1	0	0
Agent 3	0	1	0
Maj	1	1	0

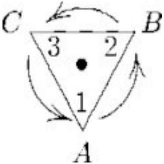

Our definition of paradox:

- F is issue by issue majority rule
- the profile is the one described in the table
- IC that is violated is $p_{\triangle\bigcirc} \wedge p_{\bigcirc\square} \rightarrow p_{\triangle\square}$

Formula-Based Judgment Aggregation

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

Definition

An agenda $\Phi = \{\varphi_1, \neg\varphi_1, \ldots, \varphi_n, \neg\varphi_n\}$ is a finite subset of propositional formulas 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.

Consistent: there exists an assignment that makes all formulas in J true.

Example: the discursive dilemma:

- The agenda: $\Phi = \{\alpha, \neg\alpha, \beta, \neg\beta, \alpha \wedge \beta, \neg(\alpha \wedge \beta)\}$
- Individuals $\mathcal{N}$
- Second judgement set: $J_2 = \{\neg\alpha, \beta, \neg(\alpha \wedge \beta)\}$

Formula-Based Judgment Aggregation and Binary Aggregation

Judgment sets J over agenda Φ	$\Leftrightarrow$	Ballot B_J over issues $\mathcal{I} = \Phi$
ex: $\{\neg\alpha, \beta, \neg(\alpha \wedge \beta)\}$		*ex.:* $(0, 1, 1, 0, 0, 1)$

IC_Φ encodes the rationality assumption of judgment aggregation:

Completeness: $p_\alpha \vee p_{\neg\alpha}$ for all $\alpha \in \Phi$

Consistency: $\neg(\bigwedge_{\alpha \in S} p_\alpha)$ for every inconsistent subset $S \subseteq \Phi$

Doctrinal Paradox/Discursive Dilemma

	α	β	$\alpha \wedge \beta$
Agent 1	1	1	1
Agent 2	0	1	0
Agent 3	1	0	0
Majority	1	1	0

Our definition of paradox:

- F is issue by issue majority rule
- profile described in the table
- IC that is violated is $\neg(p_\alpha \wedge p_\beta \wedge p_{\neg(\alpha\wedge\beta)})$

Common feature: Three issues

Ostrogorski Paradox

	Eco	EnV	Int	$Party$
Voter 1	0	1	0	0
Voter 2	0	1	0	0
Voter 3	1	0	0	0
Voter 4	1	1	1	1
Voter 5	1	1	0	1
Majority	1	1	0	0

A voter votes for a party if she agrees on a majority of the issues: but a majority of the voters disagrees with the elected party on a majority of issues!

Our definition of paradox:

- F is issue by issue majority rule
- IC that is violated is $P \leftrightarrow [(E \wedge V) \vee (V \wedge I) \vee (I \wedge E)]$

After some calculation IC is equivalent to a conjunction of clauses of size 3:

$$(P \vee \neg E \vee \neg V) \wedge (P \vee \neg E \vee \neg I) \wedge (P \vee \neg I \vee \neg V) \wedge (\neg P \vee E \vee V) \wedge (\neg P \vee E \vee I) \wedge (\neg P \vee I \vee V)$$

M. Ostrogorski. La démocratie et l'organisation de partis politiques, 1902.

Multiple Election Paradox

The outcome of the majority rule is the acceptance of all three issues, even if this combination was not voted for by any of the individuals:

	A	B	C
Voter 1	1	0	1
Voter 2	0	1	1
Voter 3	1	1	0
Maj	1	1	1

- Every paradox of aggregation by our definition is an instance of the MEP - no agent votes for irrational outcomes -
- Every instance of MEP can be seen as a paradox of aggregation by devising suitable integrity constraints -e.g., disjunction of individual ballots -

Brams, Kilgour, and Zwicker. The paradox of multiple elections. *Social Choice and Welfare*, 1998.

The Common Structure of Paradoxes

Integrity constraints formalising classical paradoxes:

Condorcet Paradox: $\neg p_{\triangle\bigcirc} \vee \neg p_{\bigcirc\square} \vee p_{\triangle\square}$

Discursive Dilemma: $\neg p_{\alpha} \vee \neg p_{\neg\beta} \vee \neg p_{(\alpha\rightarrow\beta)}$

Ostrogorski Paradox: $(P \vee \neg E \vee \neg V) \wedge (P \vee \neg E \vee \neg I) \wedge$
$\wedge (P \vee \neg I \vee \neg V) \wedge (\neg P \vee E \vee V) \wedge (\neg P \vee E \vee I) \wedge (\neg P \vee I \vee V)$

Multiple Election Paradox: can be related to our definition

Integrity constraints formalising classical paradoxes
of aggregation all feature a clause of size 3

Characterisation of Paradoxes of the Majority Rule

Theorem

The majority rule (for an odd number of individuals) is collectively rational wrt. IC if and only if IC is equivalent to a conjunction of clauses of size $\leqslant 2$.

$$\mathcal{IC}(Maj) = 2\mathrm{CNF}^{\sim}$$

Proof (sketch - see blackboard if time permits)

- A clause $\ell_1 \vee \ell_2$ of size 2 cannot generate a paradox: if ℓ_1 and ℓ_2 are both rejected then there exists an individual who rejects both literals.
- Every formula is equivalent to the conjunction of its *prime implicates*. If there exists a prime implicate of size $\geqslant 3$ then we can devise a discursive-dilemma-like situation in which the majority rule rejects it.

This thm is a syntactic counterpart of a result by Nehring and Puppe that we will see tomorrow.

Characterisation of Non-Paradoxical Aggregation Rules

A generalised dictatorship, or rule based on the most representative voter, always copies the ballot of an individual, possibly a different one in every profile:

Proposition

F is collectively rational with respect to all IC in $\mathcal{L}_{PS}$ if and only if it is a generalised dictatorship.

Proof idea: write a formula that forces F to vote as one of the individuals.

This class of rules includes:

- Classical dictatorships $F(B_1, \ldots, B_n) = B_i$ for $i \in \mathcal{N}$
- Average Voter Rule: map $(B_1, \ldots, B_n)$ to the ballot B_i that minimises a notion of distance to the profile (the "average voter").

Some references

Apart from the two general references on judgement aggregation given at the beginning, here are references for the specific results in this part of the lecture.

- An easy introduction to binary aggregation with constraints:

  U. Grandi and U. Endriss. Binary Aggregation with Integrity Constraints. In *Proceedings of IJCAI-2011*.

- An analysis of aggregation paradoxes and a full proof of the theorem characterising paradoxical ICs for the majority rule:

  U. Grandi. The Common Structure of Paradoxes in Aggregation Theory. ArXiv 1406.2855, 2012.

- A proof of the result characterising the class of generalised dictatorships:

  U. Endriss and U. Grandi. Lifting Integrity Constraints in Binary Aggregation. AIJ, 2013.

Intermezzo:

Mini-Tutorial on
Computational Complexity

Computational Complexity: what is a problem?

Computational complexity provides a measure of how difficult a *problem* is, by measuring the amount of resources needed to solve it: space or time.

What is a *decision problem*?

Instance/Input (size)

$\Downarrow$

Question

$\Downarrow$

Answer: Yes/No (in a certain time)

Important remarks:

- Complexity is not a property of a particular procedure to solve a problem.
- Complexity is not a property of an instance of a problem: it is a function that links the size of an input to the time required to compute the answer.

Easy and hard problems

Easy problems

A decision problem is in **P** *if there exists an algorithm that replies YES/NO in polynomial time.*

MODEL CHECKING
Instance: Integrity constraint IC, ballot B
Question: Does $B \models$ IC?

Difficult problems

A problem is in **NP** *if I can guess a certificate for YES in a large search space and then verify it in polynomial time.*

SATISFIABILITY
Instance: Integrity constraint IC
Question: Is there a ballot B such that $B \models$ IC?

PS. I am oversimplifying things.

Even harder problems: The second layer of the polynomial hierarchy

An **NP-oracle** is a machine that solves an NP-problem in one unit of time.

*A problem is in Σ_2^p (**sigma-2-p**) if I can guess a certificate for YES from an exponential search space, and then verify it with one query to an NP-oracle.*

Classical example: "there is an X such that for all Y we have that..."

*A problem is in Π_2^p (**pi-2-p**) if I can guess a certificate for NO from an exponential search space, and then verify it by asking one query to an NP-oracle.*

Classical example: "for all X there exists a Y such that..."

*A problem is in Θ_2^p (**theta-2-p**) if it can give a YES/NO answer by asking a logarithmic number of queries to the NP-oracle.*

Classical example: binary search (see tomorrow's lecture)

How many complexity classes are there?

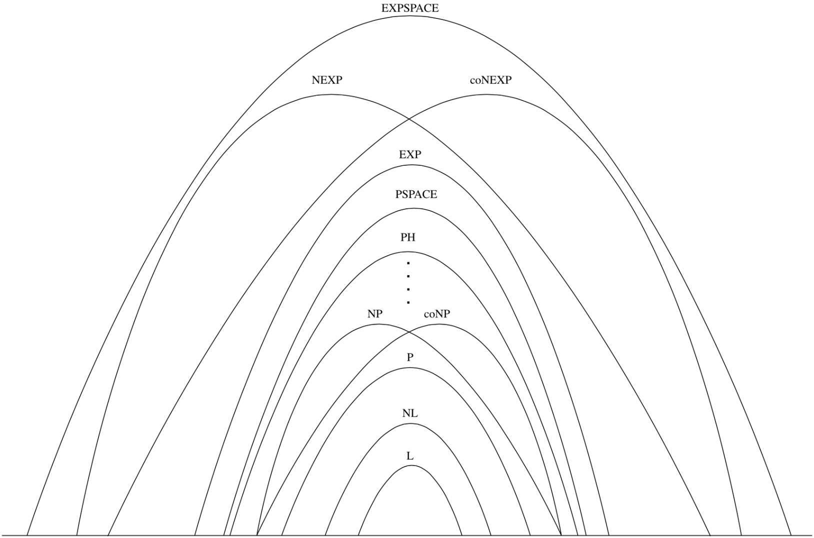

Membership and completeness

We have a decision problem $\textsc{Problem}$ and a complexity class CLASS:

- We prove membership of by devising an algorithm that solves $\textsc{Problem}$ in the time described by CLASS.
- We prove completeness by reducing a known CLASS-complete problem $\textsc{Hard}$ to $\textsc{Problem}$: we find a polynomial translation τ such that

$$x \in \textsc{Hard} \Leftrightarrow \tau(x) \in \textsc{Problem}$$

Warning and references

When studying a problem $\textsc{Problem}$ always specify clearly what is the input and what is the question, and reason about the size of the input.

A difficult problem on an exponential input becomes an easy problem:

- Input: numbers $X = \{x_1, \ldots, x_n\} \in \mathbb{Z}$. Is there $X' \subset X$ s.t. $\sum x' = 0$?
- If $x_1, \ldots, x_n$ are encoded in binary then the problem is NP-complete.
- If $x_1, \ldots, x_n$ are encoded in unary $(5 = 11111)$ then the problem is in P!

If you want to know more, these are two classical textbooks:

R. Garey and S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. 1979.

Michael Sipser. *Introduction to the Theory of Computation*. 3rd edition, 2012.

List of NP-complete problems:
`https://en.wikipedia.org/wiki/List_of_NP-complete_problems`

Part II:

Aggregation Rules and Winner Determination

Aggregation procedures

Aggregation procedures are the protagonists of JA. Most definitions take the form of non-resolute functions:

Definition

A non-resolute aggregation procedure is a function $F : \mathcal{D}^{\mathcal{N}} \to 2^{\mathcal{D}}$ mapping each profile of ballots $\boldsymbol{B} = (B_1, \ldots, B_n)$ to a subset of the domain $\mathcal{D}$.

A suitable tie-breaking rule can be used to obtain a resolute definition.

Later we will learn to evaluate aggregation procedures axiomatically, now we will start an algorithmic analysis.

Winner determination

What is the first problem to study when evaluating algorithmically an aggregation procedure?

$\textsc{WinDet}(F)$	
Instance:	Integrity constraint IC, admissible profile $\boldsymbol{B}$, subset $I \subseteq \mathcal{I}$, partial ballot $\rho : I \to \{0, 1\}$
Question:	Is there a $B^\star \in F(\boldsymbol{B})$ s.t. $\forall j \in I,\ B^\star(j) = \rho(j)$?

Some comments:

- What is the size of an instance?
- Can we compute a winning ballot? (hint: the partial ballot in the input)
- What if we had a "for all $B^* \in F(\boldsymbol{B})$" instead of "is there $B^* \in F(\boldsymbol{B})$"?

Quota Rules

Definition

For a given quota $k \leqslant n$, the **uniform quota rule** Q_k *is defined as* $Q_k(\boldsymbol{B})_j = 1$ *iff* $|\{i \in \mathcal{N} \mid b_{i,j} = 1\}| \geqslant k$.

- A uniform quota rule accepts an issue iff the number of people accepting it exceeds a given quota k.
- Special case: the majority rule. Which quota?

Majority or Nothing Rule

Definition

The **Majority or Nothing Rule** *outputs the outcome of the majority rule if this is consistent, otherwise it outputs all consistent ballots.*

NB. This rule does not exist.

Kemeny Rule

Definition

The **Kemeny rule** *is defined as follows:*

$$\text{Kemeny}^{IC}(\boldsymbol{B}) \quad = \quad \underset{B \models IC}{\text{argmin}} \sum_{i \in \mathcal{N}} H(B, B_i)$$

The Kemeny rule (aka many different names) selects the consistent ballot that minimises the sum of the Hamming distance from the individual ballots.

Dictatorship and AVR

Definition

F is called a **dictatorship** *if there exists $i \in \mathcal{N}$ such that for all profiles $\boldsymbol{B}$ we have that $F(\boldsymbol{B}) = B_i$.*

Definition

The **average-voter rule** *is the aggregation rule that selects those individual ballots that minimise the Hamming distance to the profile:*

$$\mathrm{AVR}(\boldsymbol{B}) \quad = \quad \underset{\{B_j \mid j \in \mathcal{N}\}}{\operatorname{argmin}} \sum_{i \in \mathcal{N}} H(B, B_i)$$

Maximum Subagenda Rule

Definition

The **maximum subagenda rule** *is defined as follows:*

$$\mathrm{MSA}^{\Gamma}(\boldsymbol{B}) \quad = \quad \underset{B \models \Gamma}{\mathrm{argmax}^{\subseteq}} \{ j \in \mathcal{I} \mid b_j = \mathrm{Maj}(\boldsymbol{B})_j \}$$

Where $\mathrm{argmax}^{\subseteq}$ return arguments that produce a result that is maximal with respect to set-inclusion. MSA takes a maximally consistent subset of the majority and completes it to a consistent ballot (it is not the ballot that agrees on the highest number of issues with the majority, this is the Slater rule).

Binomial Rule

Definition

For every $k \leqslant n$, the **binomial-k rule** *is defined as:*

$$\mathrm{Bin}_k^{IC}(\boldsymbol{B}) \quad = \quad \underset{B \models IC}{\mathrm{argmax}} \sum_{i \in \mathcal{N}} \binom{n - H(B, B_i)}{k}$$

Ok, this is hard to get. Consider the easier example of k=m-1:

- guess a consistent ballot
- one point for each subset of $m-1$ issues on which the ballot coincides with the ballot of one individual in the profile
- the ballot with the maximal score wins

If you are reading the slides, here are the answers

1. Computing the result of a quota rule is easy/polynomial
   Just read the input and count how many voters accept a given issue.
2. The Majority-or-Nothing Rule is NP-complete: it corresponds to SAT.
3. Winner determination for the Kemeny rule is Θ_2^p-complete
   Membership: use an oracle *Kemeny-score* which given a ballot and an integer K tells you whether there is a winning ballot at distance at most K. The oracle is in NP since it can be modelled as an integer program. Then do binary search over the K (bounded by maximal distance $m \times n$) to find the exact distance of the winner in a logarithmic number queries to the oracle. Then a last call gives you the winner. **Completeness**: reduction from the voting version of the Kemeny rule.
4. Dictatorships as well as the AVR are easy/polynomial.
5. The MSA is Σ_2^p complete. **Membership**: Define an oracle that given a partial assignment tells me if it can be extended to an assignment satisfying IC. Guess a subset of the majority: decide with a polynomial number of queries to the oracle if it is maximally consistent (one query per subset), and then complete it to an assignment satisfying IC. **Completeness**: reduction from SKEPTICAL INFERENCE.
6. The special case of the binomial rule with $k = m - 1$ is polynomial: only ballots at distance 1 from an individual can win. It is Θ_2^p in general.

References and lots of other rules

- Quota rules have been defined and studied in this paper:

  F. Dietrich and C. List. Judgment aggregation by quota rules: majority voting generalized. *Journal of Theoretical Politics*, 2007

- Rules based on minimisation (Kemeny and MSA):

  J. Lang, G. Pigozzi, M. Slavkovik, L. van der Torre: Judgment aggregation rules based on minimization. In *Proceedings of TARK 2011*. Extended version on the authors' webpage.

- The binomial rule in this paper:

  M. Costantini, C. Groenland, U. Endriss: Judgment Aggregation under Issue Dependencies. In *Proceedings of AAAI 2016*.

- Papers on the complexity of judgement aggregation (there are others):

  J. Lang, M. Slavkovik: How Hard is it to Compute Majority-Preserving Judgment Aggregation Rules? In *Proceedings of ECAI 2014*.

  U. Endriss, R. de Haan: Complexity of the Winner Determination Problem in Judgment Aggregation: Kemeny, Slater, Tideman, Young. In *Proceedings of AAMAS 2015*.

  U. Endriss, U. Grandi, R. de Haan, J. Lang: Succinctness of Languages for Judgment Aggregation. In *Proceedings of KR 2016*.

Part III:

Axiomatic Approach and Winning Coalitions

Formula-Based Judgment Aggregation

Let the complement of φ be $\neg\varphi$ if φ is non-negated and ψ if $\varphi = \neg\psi$.

Definition

An agenda is a finite subset of propositional formulas closed under complementation and not containing double negations. We assume that there are at least two logically independent formulas in Φ (non-triviality).

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.

Complement-free: α and its complement are never both in J.

Consistent: there exists an assignment that makes all formulas in J true.

$\mathcal{N}$ is a set of judges and $\mathcal{J}(\Phi)$ all consistent and complete judgment sets.

Definition

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

The discursive dilemma - take III

Discursive Dilemma

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

Corresponds to:

- The agenda $\Phi = \{p, \neg p, q, \neg q, p \wedge q, \neg(p \wedge q)\}$
- The judges $\mathcal{N} = \{1, 2, 3\}$
- Judgement set $J_1 = \{p, q, p \wedge q\}$
- Judgement set $J_2 = \{\neg p, q, \neg(p \wedge q)\}$
- Judgement set $J_3 = \{p, \neg q, \neg(p \wedge q)\}$
- $\mathrm{Maj}(\boldsymbol{J}) = \{p, q, \neg(p \wedge q)\}$: complete, complement-free but not consistent.

Axiomatic Properties

Aggregation procedures have been studied using the axiomatic method:

Unanimity (U): If $\varphi \in J_i$ for all $i \in \mathcal{N}$ implies $\varphi \in F(\boldsymbol{J})$.

Anonymity (A): $F(J_1, \ldots, J_n) = F(J_{\sigma(1)}, \ldots, J_{\sigma(n)})$ for all $\sigma : \mathcal{N} \to \mathcal{N}$.

Neutrality (N): For all $\varphi, \psi \in \Phi$ if $\varphi \in J_i \Leftrightarrow \psi \in J_i$ for all $i \in \mathcal{N}$ then $\varphi \in F(\boldsymbol{J}) \Leftrightarrow \psi \in F(\boldsymbol{J})$.

Independence (I): For all $\varphi \in \Phi$ and any two profiles $\boldsymbol{J}, \boldsymbol{J}' \in \mathcal{J}(\Phi)^n$, if $\varphi \in J_i \Leftrightarrow \varphi \in J_i'$ for all agents $i \in \mathcal{N}$, then $\varphi \in F(\boldsymbol{J}) \Leftrightarrow \varphi \in F(\boldsymbol{J}')$.

We also specify what properties of the judgment sets are lifted from individual to collective level:

- F is **complete** if $F(\boldsymbol{J})$ is complete for all $\boldsymbol{J} \in \mathcal{J}(\Phi)^{\mathcal{N}}$.
- F is **complement-free** if $F(\boldsymbol{J})$ is complement-free for all $\boldsymbol{J} \in \mathcal{J}(\Phi)^{\mathcal{N}}$.
- F is **consistent** if $F(\boldsymbol{J})$ is consistent for all $\boldsymbol{J} \in \mathcal{J}(\Phi)^{\mathcal{N}}$.

Winning coalitions

Let $N_\varphi^{\boldsymbol{J}}$ be the set of agents accepting formula φ in profile $\boldsymbol{J}$.

Lemma

F is independent iff for each $\varphi \in \Phi$ there is a family of sets of judges $\mathcal{W}_\varphi$ such that $\varphi \in F(\boldsymbol{J})$ iff $N_\varphi^{\boldsymbol{J}} \in \mathcal{W}_\varphi$.

$\mathcal{W}_\varphi$ are the winning coalitions for φ. We can reinterpret some of the axioms:

Unanimity: $\mathcal{N} \in \mathcal{W}_\varphi$ for all $\varphi \in \Phi$;

Anonimity: $\mathcal{W}_\varphi$ are closed under equinumerosity;

Neutrality: $\mathcal{W}_\varphi = \mathcal{W}_\psi$ for all $\varphi, \psi \in \Phi$;

Completeness: $\mathcal{W}_\varphi$ is maximal.

Winning coalitions are used in impossibility theorems to show that the only rules satisfying a certain set of axioms are dictatorships: the ultrafilter technique.

An Impossibility Theorem

Theorem [List and Pettit, 2002]

No aggregator for agenda $\Phi = \{p, q, p \wedge q\}$ can be anonymous, neutral, independent, complete and consistent.

Proof.

Suppose F satisfies all these axioms. Use the power of winning coalitions: If formulas φ and ψ are accepted by the same number of judges, then F either accept or rejects both of them (independence + neutrality + anonymity).

Consider a profile $\boldsymbol{J}$ such that:

- $\frac{n-1}{2}$ judges accept both p and q
- one judge accepts p but not q
- one judge accepts q but not p
- $\frac{n-3}{2}$ judges accept neither p nor q

The same number of judges in $\boldsymbol{J}$ accepts p, q and $\neg(p \wedge q)$. F has to either accept all or reject all formulas, in both cases contradicting consistency. $\square$

A slightly more complex one

The characterisation of paradoxes for the majority rule seen yesterday corresponds to the following result:

Theorem [Nehring and Puppe, 2007]

The majority rule is consistent for a given agenda Φ iff any inconsistent subset of Φ does itself have an inconsistent subset of size at most 2 (median property).

See exercises at the end for a proof.

What else can you do with axioms

The axiomatic rule is a powerful tool to characterise and analyse aggregators:

- Prove characterisation results such as
  "F is a quota rule iff F satisfies AXIOMS"
- Prove existential agenda characterisation results:
  "If Φ is in a certain class then all F satisfying AXIOMS are dictatorial"
- Prove universal agenda characterisation results (aka safety):
  "F is consistent on a certain class of agendas iff F satisfies AXIOMS"

Refer to the two textbooks for more results of this kind:

D. Grossi and G. Pigozzi. Judgment Aggregation: A Primer. Morgan and Claypool, 2014.

U. Endriss. Judgment Aggregation. In *Handbook of Computational Social Choice*, 2016.

Part IV:
Manipulation

Where are the preferences?

Judgments are not preferences, the incentive structure they induce is minimal:

- $F(\boldsymbol{B}) = J_i$ then agent i is happy
- $F(\boldsymbol{B}) \neq J_i$ then agent i is not happy

Is the majority rule strategy-proof with this incentive structure?

It is plausible to assume that "the more formulas in J_i are accepted, the more agent i is happy" (Hamming preferences). But even then:

Theorem [Dietrich and List, 2007]

All independent and monotonic rules are strategy-proof for Hamming preferences.

But independent aggregators are rarely consistent...

Many different approaches to manipulation

How hard is to manipulate non-independent procedures?

- Premise-based rule is NP-hard with Hamming preferences: *Endriss et al. Complexity of Judgement Aggregation. JAIR, 2012.*
- Kemeny rule (and non-resolute manipulation): *Ronald's poster*

Study manipulation, bribery and control under Hamming-like preferences:

- *Papers by D. Baumeister, G. Erdelyi, O.J. Erdelyi, and J. Rothe.*

Assign goals to agents (like in the initial game) and study the equilibria induced:

- *Grandi, Grossi and Turrini. Equilibrium Refinement through Negotiation in Binary Voting. IJCAI-2015 (amended and expanded version on ArXiv)*

Define and study group manipulation:

- *Botan, Novaro, and Endriss. Group Manipulation in Judgment Aggregation. AAMAS-2016.*

Part V:

“Applications”

Collective annotations

How to accurately generate annotated data such as images or structured text?

Crowdsourcing is a viable option: ask a number of people to annotate the same object and then aggregate their reply to minimize mistakes.

Judgment aggregation (and social choice in general) provides a number of principled methods to aggregate information:

- Axiomatic and empirical analysis of methods for group annotation
- Development of procedures for human computation

J. Kruger, U. Endriss, R. Fernndez, and C. Qing. Axiomatic Analysis of Aggregation Methods for Collective Annotation. AAMAS-2014.

C. Qing, U. Endriss, R. Fernndez, and J. Kruger. Empirical Analysis of Aggregation Methods for Collective Annotation. COLING-2014.

(voting) A. Mao, A. D. Procaccia, Y. Chen. Better Human Computation Through Principled Voting. AAAI 2013.

Multiagent Argumentation and Judgment Aggregation

(Abstract) argumentation models debates with a number of graphs:

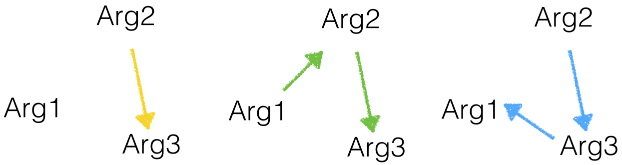

What is a collectively accepted argument? Basic question: should we aggregate the attack graphs or the sets of accepted arguments?

S. Coste-Marquis, C. Devred, S. Konieczny, M.C. Lagasquie-Schiex, P. Marquis. On the merging of Dungs argumentation systems. AIJ, 2007.

M. Caminada, G. Pigozzi. On Judgment Aggregation in Abstract Argumentation. Autonomous Agents and Multi-Agent Systems, 2011.

R. Booth, E. Awad, I. Rahwan. Interval Methods for Judgment Aggregation in Argumentation. KR 2014.

Opinion Diffusion on a Network

Model diffusion (of judgments, of preferences, of opinions in general) on a network with aggregation procedures:

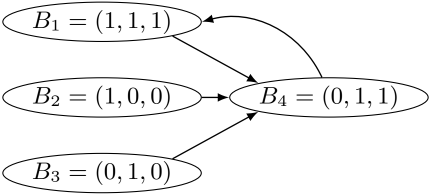

Agent 4 updates her opinion by aggregating the opinions of her influencers 1,2 and 3 using an aggregation procedure F, so $B'_4 = F(B_1, B_2, B_3)$.

Grandi, Lorini and Perrussel. Propositional Opinion Diffusion. AAMAS-2015.

Brill et al. Pairwise Diffusion of Preference Rankings in Social Networks. IJCAI-2016. (preferences)

Conclusions

Judgment aggregation is a general framework to study collective decisions over multiple interconnected issues:

- a general definition of paradox
- all sorts of rules to aggregate information
- axiomatic and algorithmic tools to analyse procedures
- open problems in modelling strategic behaviour

The theoretical framework behind judgment/binary aggregation is solid: plenty of potential "applications". Don't be scared to use propositional logic!

For a more complete and detailed introduction:

D. Grossi and G. Pigozzi. Judgment Aggregation: A Primer. Morgan and Claypool, 2014.

U. Endriss. Judgment Aggregation. In *Handbook of Computational Social Choice*, 2016.

Exercises

1. (lengthy and maybe not too easy) Formula-based JA is equivalent to BA with integrity constraints: Prove it formally.
2. Once you are convinced that every agenda can be represented as an integrity constraint and *viceversa*, show the following:
   - Show (directly) that an agenda Φ has the median property if and only if the integrity constraint that corresponds to it can be expressed with a conjunction of clauses of size 2.
   - Formulate the maximal subagenda rule in formula-based JA.
   - (*one beer problem*) Look up the definition of a *totally blocked agenda*. Translate this condition into binary aggregation as a property of an integrity constraint (possibly a syntactic property).
3. Find at least two names of the Kemeny Rule in the literature on JA.
4. The Slater Rule outputs the consistent ballot that coincides with the majority outcome on the highest number of issues. What is the complexity of winner determination for this rule? (hint: mimic what done for the Kemeny rule)
5. Prove the Lemma relating winning coalitions with the independence axiom.
6. Prove everything you can about aggregation procedures on one single issue (on agendas with a single formula). hint: look up May's Theorem