Judgment Aggregation under Issue Dependencies

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

[joint work with Marco Costantini and Carla Groenland]

Talk Outline

- Introduction to Judgment Aggregation
- Binomial Rules: An Attempt to Account for Hidden Dependencies
- Theoretical Analysis: Axiomatics and Computational Complexity
- Experimental Analysis: Aggregating Hotel Reviews

M. Costantini, C. Groenland, and U. Endriss. Judgment Aggregation under Issue Dependencies. Proc. AAAI-2016.

Example: Robots and Air Conditioning

Five robots, equipped with heat sensors and cameras, provide basic services at a congress centre. Should they switch on the A/C?

	Warm?	*Busy?*	*A/C?*
Robot 1:	Yes	Yes	Yes
Robot 2:	Yes	Yes	Yes
Robot 3:	Yes	No	No
Robot 4:	No	Yes	No
Robot 5:	Yes	No	No

Integrity Constraint: $Warm \wedge Busy \rightarrow A/C$

Judgment Aggregation

Studied in Law, Philosophy, Logic, Economics, AI, Computer Science.
Much of it going back to seminal paper by List and Pettit (2002).

Set of *issues* $\mathcal{I} = \{1, \ldots, m\}$. *Ballots* $B = (b_1, \ldots, b_m) \in \{0,1\}^m$.
Integrity constraint Γ over $\{p_1, \ldots, p_m\}$. B is *rational* if $B \in \mathrm{Mod}(\Gamma)$.

Set of *agents* $\mathcal{N} = \{1, \ldots, n\}$. *Profile* $\boldsymbol{B} = (B_1, \ldots, B_n) \in \mathrm{Mod}(\Gamma)^n$.
An *aggregation rule* maps every rational profile to a consensus ballot:

$$F : \mathrm{Mod}(\Gamma)^n \to \mathcal{P}(\{0,1\}^m) \setminus \{\emptyset\}$$

This specific variant of JA is known as *binary aggregation* with IC's.
Could all be done in other frameworks as well.

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

U. Endriss. Judgment Aggregation. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A.D. Procaccia (eds.), *Handbook of Computational Social Choice*. CUP, 2016.

Aggregation Rules

Various rules have been proposed in the literature. Examples:

- A (uniform) *quota rule* accepts an issue if at least k individuals do (e.g., *majority rule* for $k = \lceil \frac{n}{2} \rceil$). Not *collectively rational*.
- The *Kemeny rule* returns the rational ballot(s) minimising the sum of the Hamming distances to the individual ballots.
- A *representative-voter rule* returns the "most representative" input ballot (e.g., *average-voter rule* or *plurality-voter rule*).

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

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

U. Endriss and U. Grandi. Binary Aggregation by Selection of the Most Representative Voter. Proc. AAAI-2014.

Example: Choosing a Common Meal for a Party

A group of 23 gastro-entertainment professionals need to decide on the meal (1 dish + 1 drink) to be served at a party. What to choose?

	Chips?	*Beer?*	*Caviar?*	*Champagne?*
11 individuals:	Yes	Yes	No	No
10 individuals:	No	No	Yes	Yes
2 individuals:	No	Yes	Yes	No

Integrity Constraint: $(\textit{Chips} \text{ XOR } \textit{Caviar}) \wedge (\textit{Beer} \text{ XOR } \textit{Champagne})$

Binomial Rules

Goal: design rules to capture *hidden dependencies* between issues.

Idea: award 1 point to potential outcome $B^\star$ for every ballot B_i and issue set $I \subseteq \mathcal{I}$ with $|I| \in K$ such that B_i and $B^\star$ fully agree on I. So K is the range of subset sizes we choose to care about.

$$\sum_{I \subseteq \mathcal{I} \text{ s.t. } |I| \in K} |\{i \in \mathcal{N} \mid \forall j \in I. b_{ij} = b^\star_j\}| \;=\; \sum_{B \in \boldsymbol{B}} \sum_{k \in K} \binom{\mathrm{Agr}(B, B^\star)}{k}$$

As those binomial coefficients can get very large, our most general definition also includes a *weight function* $w : K \to \mathbb{R}^+$.

$$F_{K,w} : \boldsymbol{B} \;\mapsto\; \operatorname*{argmax}_{B^\star \in \mathrm{Mod}(\Gamma)} \sum_{B \in \boldsymbol{B}} \sum_{k \in K} w(k) \cdot \binom{\mathrm{Agr}(B, B^\star)}{k}$$

Interesting special cases: $K = \{k\}$ (in which case w is irrelevant).
Note: this is *Kemeny rule* for $k = 1$ and *plurality-voter rule* for $k = m$.

Theoretical Results

Nice axiomatic properties (but full characterisation is open):

Theorem 1 *Binomial rules are amongst the very few rules discussed in the literature that satisfy both collective rationality and reinforcement:*

$$F(\boldsymbol{B}) \cap F(\boldsymbol{B'}) \neq \emptyset \;\; \textit{implies} \;\; F(\boldsymbol{B} \oplus \boldsymbol{B'}) = F(\boldsymbol{B}) \cap F(\boldsymbol{B'})$$

Binomial rules cover the range from the trivial to the highly intractable:

Theorem 2 *Winner determination for $F_{\{k\}}$ is in P if $(m-k) \in O(1)$.*

Theorem 3 *But the same problem is $\mathrm{P}^{\mathrm{NP}}[\log]$-complete if $k \in O(1)$.*

Experiment: Aggregating Hotel Reviews

Ratings on 1–5 star scale for 6 features (*location*, *cleanliness*, etc.) of 1850 hotels collected on TripAdvisor (68 reviews/hotel on average).

Translation to binary judgments: *accept* (4–5) or *reject* (1–3) issue.
So: 6 issues, no IC, 1850 profiles, with an average of 68 ballots each.

Presence of at least some hidden dependencies intuitively very likely.

H. Wang, Y. Lu, and C. Zhai. Latent Aspect Rating Analysis on Review Text Data: A Rating Regression Approach. Proc. KDD-2010. Data available on `PrefLib.org`.

The Compliant Reviewer Problem

What makes for a good meta review (the result of the aggregation)?

> *You are writing a hotel review for an online magazine and you want to please as many of your readers as possible (to maximise the number of like's received). Suppose a reader will like your review if she agrees with you on $\geqslant k$ issues.*

We will use this *compliant-reviewer score* to evaluate our results (measuring the quality of outcome $B^\star$ for profile $\boldsymbol{B}$ and threshold k):

$$\frac{1}{n} \cdot |\{i \in \mathcal{N} \mid \mathrm{Agr}(B_i, B^\star) \geqslant k\}|$$

Results for Full Data Set

Not much difference between majority rule and two binomial rules with with full index set $K = \{1, \ldots, m\}$ and two different weight functions.

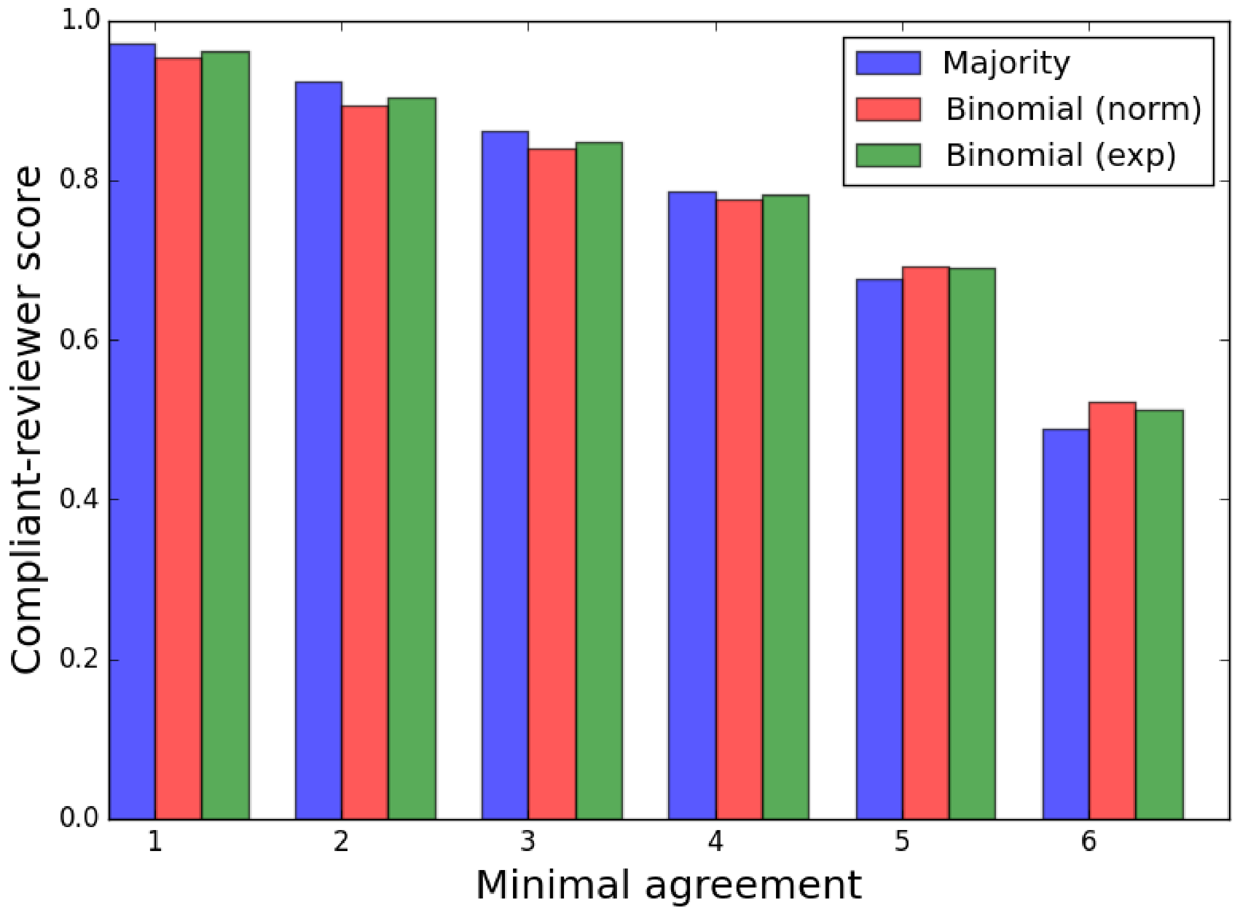

Polarisation

In the paper, we develop a formal measure of *polarisation* of a profile, defined as the product of a *correlation* and an *uncertainty coefficient:*

- correlation = average strength of dependencies between issue pairs
- uncertainty = average disagreement on individual issues

A subset of *31 profiles* is "*highly polarised*" (avg. of 32 ballots each).

Results for Highly Polarised Profiles

Positive effect of our rules much clearer for this dataset. Cases of agreement on $\geqslant 4$ issues with the compliant reviewer most realistic.

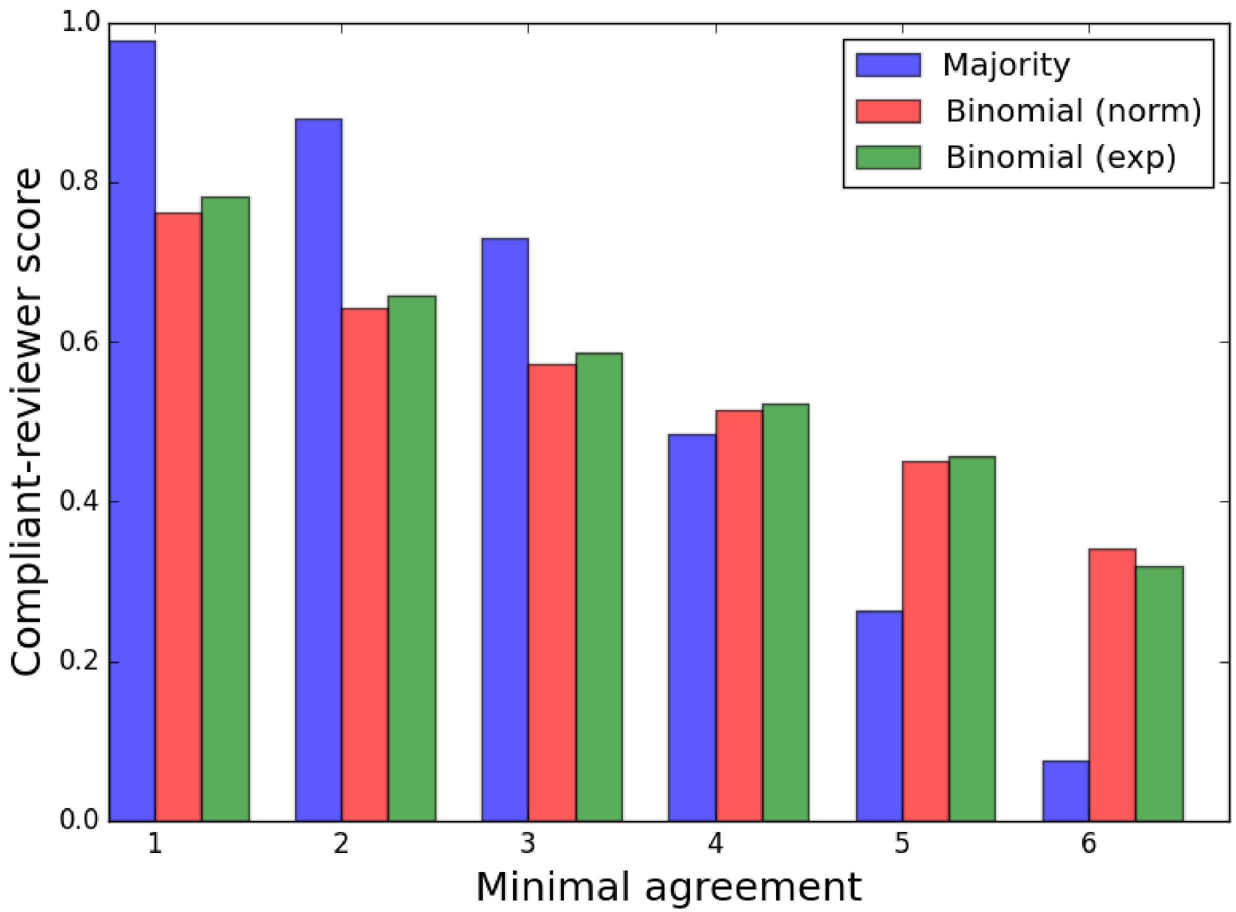

Last Slide

Proposal for a *new family of judgment aggregation rules:*

- Attempt to account for hidden dependencies between issues
- Score agreement of outcome with ballots on subsets of issues
- Parameters: subset sizes to consider + weight function

Initial results for these so-called *binomial rules:*

- Includes *spectrum of rules* from Kemeny to plurality-voter rule
- Complexity: *winner determination* ranges from P to $\mathrm{P}^{\mathrm{NP}}[\log]$
- Axiomatics: both *collective rationality* and *reinforcement* ok
- Experiments: *good performance* for highly polarised hotel reviews

New concepts of potentially independent interest:

- Notion of *polarisation* of a profile in judgment aggregation
- *Compliant Reviewer Problem* (related to *Ostrogorski Paradox*)

M. Costantini, C. Groenland, and U. Endriss. Judgment Aggregation under Issue Dependencies. Proc. AAAI-2016. *Paper will be on my website next week.*