Aggregation in multi-agent systems and the problem of truth-tracking

Gabriella Pigozzi[1] and Stephan Hartmann[2]
[1]University of Luxembourg, Luxembourg; [2]Tilburg University, Tilburg, The Netherlands
[1]gabriella.pigozzi@uni.lu; [2]s.hartmann@uvt.nl

Judgment aggregation

How can individual judgments (yes/no) on logically interconnected propositions be aggregated into corresponding collective judgments on the same propositions?

The discursive paradox [Kornhauser & Sager 1986]

P = the defendant did some action X
Q= the defendant had a contractual obligation not to do action X
R = the defendant is liable

$$(P \wedge Q) \leftrightarrow R$$

	P	Q	R
Judges 1,2,3	Yes	Yes	Yes
Judges 4,5	Yes	No	No
Judges 6,7	No	Yes	No
Majority	Yes	Yes	No

Proposition-wise majority voting can give an inconsistent collective decision.

Two escape routes: Premise-based (PBP) and conclusion-based procedures (CBP)
PBP: aggregate on P and Q
CBP: aggregate on R

PBP and CBP avoid the paradox but may lead to two different results ⇒ The agenda-setter can manipulate the outcome.

Fusion operators and MAS

Belief fusion [Lin & Mendelzon, Konieczny & Pino Pérez] studies the aggregation of (possibly conflicting) finite sets of information into a collective one.

- Similar problems in CS: combination of different and potentially conflicting sources of information, e.g. multi-sensor fusion, database integration and expert systems development.
- Applications in MAS that aim at aggregating the distributed agent-based knowledge into an (ideally) unique set of propositions.
- Application of information fusion operators to paradoxes of aggregation like the *discursive paradox* ⇒ No paradox

Belief fusion

- A *fusion operator* Δ maps a multiset of information bases $\{K_1, K_2, \ldots, K_n\}$ and a set of IC into a new (collective) information base $\Delta_{IC}(E)$.
- Merging operators require the satisfaction of *integrity constraints* (IC).
- Model-based fusion operators: models of $\Delta_{IC}(E)$ are models of IC that *differ minimally* from the models of each K_i.
- *Hamming distance*: the distance between $w = (1, 0, 0, 1)$ and $w' = (0, 1, 0, 1)$ is dist(w, w') = 2.

What happens when we apply methods from information fusion to collective decision problems?

Belief fusion applied to the discursive paradox [Pigozzi 2006]

Agenda $X = \{P, Q, R\}$ with
$IC = \{(P \wedge Q) \leftrightarrow R\}$

Mod(K_1)=Mod(K_2)=Mod(K_3)=$\{(1,1,1)\}$
Mod(K_4)=Mod(K_5)=$\{(1,0,0)\}$ and
Mod(K_6)=Mod(K_7)=$\{(0,1,0)\}$

	K_1	K_2	K_3	K_4	K_5	K_6	K_7	Δ^E_{IC}
(1,1,1)	0	0	0	2	2	2	2	8
(1,1,0)	1	1	1	1	1	1	1	7
(1,0,1)	1	1	1	1	1	3	3	11
(1,0,0)	2	2	2	0	0	2	2	10
(0,1,1)	1	1	1	3	3	1	1	11
(0,1,0)	2	2	2	2	2	0	0	10
(0,0,1)	2	2	2	2	2	2	2	14
(0,0,0)	3	3	3	1	1	1	1	13

The fusion operator selects (1,1,1).

A distance-based fusion operator allows to merge a finite number of information bases into a (possibly unique) information base.

The problem of truth-tracking

- We assume that all propositions in question are factually right or wrong.
- How good is information fusion as a truth tracker? Will it single out the true set of propositions?
- How does information fusion compare with other aggregation procedures?
- Assumption: The agents have a competence to judge the truth or falsity of a proposition correctly and vote accordingly.

The framework

- We consider the case of $P \wedge Q \leftrightarrow R$.
- P and Q are logically and probabilistically independent.
- All agents are equally competent (or reliable) and independent.
- The chance that an agent correctly judges the truth or falsity of P (her **competence**) is $\boldsymbol{p}$. The same for Q.
- The **prior probability** that P is true is $\boldsymbol{q}$. The same for Q.
- There are four possible situations:
  - $S_1 = \{P, Q, R\} = (1, 1, 1)$
  - $S_2 = \{P, \neg Q, \neg R\} = (1, 0, 0)$
  - $S_3 = \{\neg P, Q, \neg R\} = (0, 1, 0)$
  - $S_4 = \{\neg P, \neg Q, \neg R\} = (0, 0, 0)$

Method:

- We want to calculate the probability that fusion ranks the right situation first ($\boldsymbol{F}$).
- Note that

$$\mathcal{P}(F) = \sum_{i=1}^{4} \mathcal{P}(F|S_i) \cdot \mathcal{P}(S_i),$$

- The prior probabilities of the situations are (with $\bar{x} := 1 - x$): $\mathcal{P}(S_1) = q^2$; $\mathcal{P}(S_2) = \mathcal{P}(S_3) = q\bar{q}$; $\mathcal{P}(S_4) = \bar{q}^2$
- Let's assume that S_1 is the right situation. Then information fusion ranks the right situation first if

$$d_1 \leq \min(d_2, d_3, d_4).$$

- The distances d_i can be expressed in terms of the numbers n_i of agents for each S_i:

$$d_1 = 2n_2 + 2n_3 + 3n_4 \ , \ d_2 = 2n_1 + 2n_3 + n_4$$
$$d_3 = 2n_1 + 2n_2 + n_4 \ , \ d_4 = 3n_1 + n_2 + n_3$$

- We now calculate $\mathcal{P}(F|S_1)$:

$$\sum_{n_1,\ldots,n_4=0}^{N} \binom{N}{n_1, \ldots, n_4} p^{2n_1} (p\bar{p})^{n_2+n_3} \bar{p}^{2n_4} \, \mathcal{C}(n_1, \ldots, n_4)$$

- The sum is constrained: $\mathcal{C}(n_1, \ldots, n_4) = 1$ if (i) $\sum_{i=1}^{4} n_i = N$ and (ii) $d_1 \leq \min(d_1, \ldots, d_4)$. Otherwise $\mathcal{C}(n_1, \ldots, n_4) = 0$.

Results

1. **Comparison of the procedures for the right situation**

Fusion ranks the right judgment set first (R) compared with PBP (G) and CBP (B) for $N = 3$ and $q = .5$

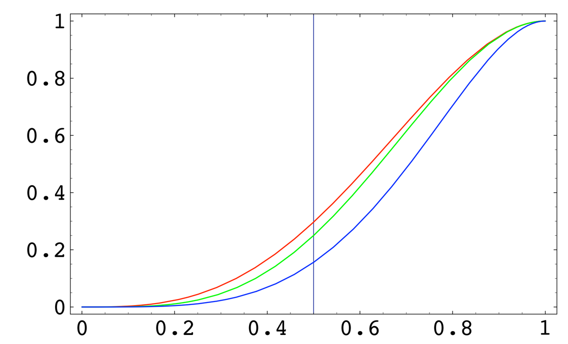

Same for $N = 21$

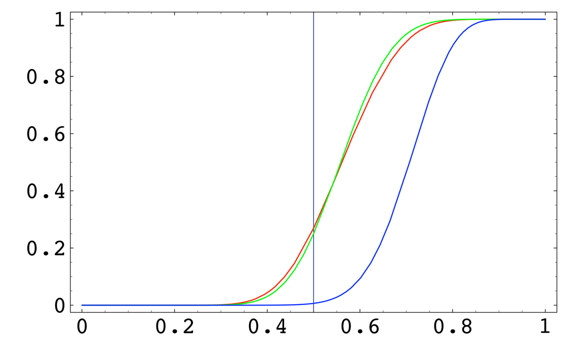

2. **Comparison of the procedures for the right result with $q = .5$**

We now calculate the probability that F ranks the right result first (proposition $\boldsymbol{G}$). The calculation proceeds analogously to $\mathcal{P}(F)$.

Fusion ranks a judgment set with the right result (not necessarily for the right reasons) first (R) compared with PBP (G) and CBP (B) for $N = 3$ and $q = .5$.

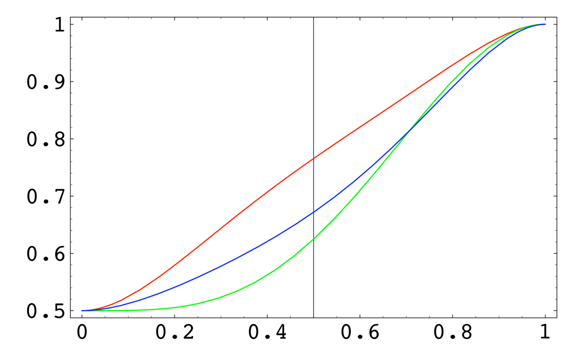

Same for $N = 31$

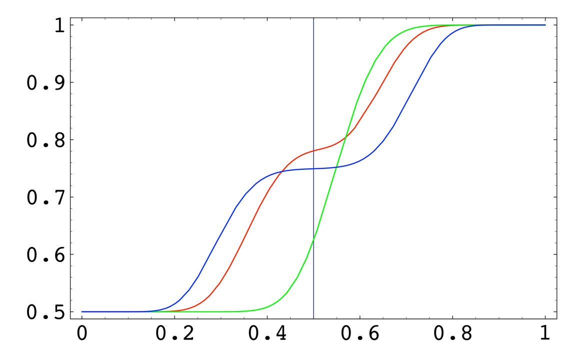

3. **Comparison of the procedures for the right result with $q = .2$**

For $N = 3$ and $q = .2$

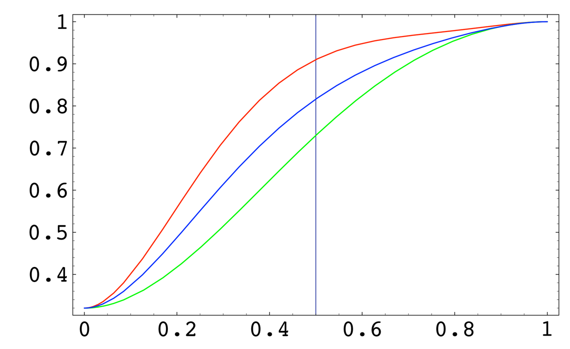

Same for $N = 51$

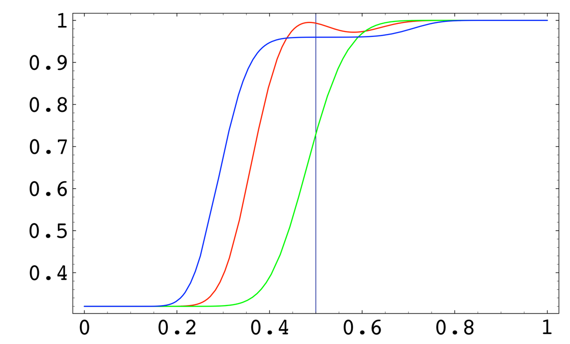

Upshot: Fusion does especially well for middling values of p.

4. **Understanding the dip**

We decompose:

$$\mathcal{P}(G) = \sum_{i=1}^{4} \mathcal{P}(G|S_i) \cdot \mathcal{P}(S_i)$$
$$= \sum_{i=1}^{4} [\mathcal{P}_{<}(G|S_i) + \mathcal{P}_{>}(G|S_i)] \cdot \mathcal{P}(S_i)$$

with
$\mathcal{P}_{<}(G|S_i)$ [$\mathcal{P}_{>}(G|S_i)$]: the probability that fusion selects the right outcome given that the *minority* [*majority*] votes for the right situation S_i.
For $N = 21$:

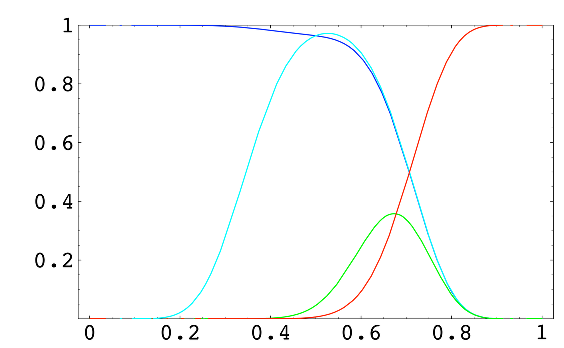

R: majority votes for the right situation, $q = .2$
Curves for the probability that F gets the right outcome given that the minority votes for the right situation, S_1 true (G); the minority votes for the right situation, S_2 or S_3 true (B); the minority votes for the right situation, S_4 true (T).