

Aggregation in Multi-Agent Systems and the Problem of Truth-Tracking

Gabriella Pigozzi
Computer Science and Communications (CSC)
University of Luxembourg
L-1359 Luxembourg
Gabriella.Pigozzi@uni.lu

Stephan Hartmann
Department of Philosophy
London School of Economics
London WC2A 2AE - UK
S.Hartmann@lse.ac.uk

ABSTRACT

One of the major problems that artificial intelligence needs to tackle is the combination of different and potentially conflicting sources of information. Examples are multi-sensor fusion, database integration and expert systems development. In this paper we are interested in the aggregation of propositional logic-based information, a problem recently addressed in the literature on *information fusion*. It has applications in multi-agent systems that aim at aggregating the distributed agent-based knowledge into an (ideally) unique set of propositions. We consider a group of autonomous agents who individually hold a logically consistent set of propositions. Each set of propositions represents an agent's beliefs on issues on which the group has to make a collective decision. To make the collective decision, several aggregation procedures have been proposed in the literature. Assuming that all propositions in question are factually right or wrong, we ask how good belief fusion is as a truth tracker. Will it single out the true set of propositions? And how does information fusion compare with other aggregation procedures? We address these questions in a probabilistic framework and show that information fusion does especially well for agents with a middling competence of hitting the truth of an individual proposition.

Categories and Subject Descriptors

F.4.3 [Mathematical Logic and Formal Languages]: Formal Languages—*Decision problems*; I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*Multi-agent systems*

General Terms

Reliability

Keywords

Information fusion, belief merging, group decision-making, truth-tracking, judgment aggregation, doctrinal paradox.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'07 May 14–18 2007, Honolulu, Hawaii, USA.
Copyright 2007 IFAAMAS.

1. INTRODUCTION

The problem of combining information from equally reliable sources arises in artificial intelligence in several contexts. An example is the definition of an artificial expert system from the knowledge of human experts on a specific topic (a disease, a portfolio investment, etc.). The pieces of information of the human experts are represented in a formal language and the goal is to merge them in an (ideally) unique knowledge base. The sources of information can, of course, also be artificial agents: for instance, a robot receiving information from its sensors needs to take into account all the (possibly conflicting) data to make a decision. The recent field of *information (or belief) fusion* studies how to aggregate individual belief bases into a collective one.¹

In many realistic cases the individual belief bases are in conflict with each other. In this paper we will consider how to combine the beliefs or the knowledge of several agents. More specifically, we will focus on three approaches to aggregate conflicting information from equally reliable sources.

We will assume that the resulting belief base is factually right or wrong. Thus, a natural question is how good a specific fusion operator is in tracking the truth, and how it compares with the other, more traditional, procedures. We will address this question in a probabilistic framework.

The combination of finite sets of logically interconnected propositions has recently drawn much attention in disciplines other than logic, especially decision theory and the new field called *judgment aggregation*. Judgment aggregation studies how to combine consistent individual judgments on logically interconnected propositions into a collective judgment on the same propositions. A *judgment* is an assignment of yes/no to a proposition. The difficulty lies in the fact that a seemingly reasonable aggregation procedure, such as propositionwise majority voting, cannot ensure a consistent collective outcome.

In the original problem of judgment aggregation, a court has to make a decision on whether a person is liable of breaching a contract (proposition R , or conclusion). The judges have to reach a verdict following the legal doctrine. This states that a person is liable if and only if she did a certain action X (first premise P) and had contractual obligation not to do X (second premise Q). The legal doctrine can be formally expressed as the rule $(P \wedge Q) \leftrightarrow R$. Each member of the court expresses her judgment on the propositions P , Q and R such that the rule $(P \wedge Q) \leftrightarrow R$ is satisfied. Suppose now that the seven members of the court

¹See [2] for a survey on logic-based approaches to information fusion.

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

Table 1: The doctrinal paradox

make their judgments according to Table 1.

We see that although each judge expresses a consistent opinion, propositionwise majority voting (consisting in the separate aggregation of the votes for each proposition P , Q and R via the majority rule) results in a majority for P and Q and yet a majority for $\neg R$. This is clearly an inconsistent collective result. The paradox (called the *doctrinal paradox*) lies in the fact that majority voting can lead a group of rational agents to endorse an irrational collective judgment. Clearly, the relevance of such aggregation problems goes beyond the specific court example and it applies to all situations in which individual binary evaluations need to be combined into a group decision.

The first two escape-routes that have been suggested are the *premise-based procedure* (PBP) and the *conclusion-based procedure* (CBP). The first procedure is to let each member vote on each premise and to declare the defendant liable only if a majority of the court believes that she did the action X and that she was under contract obligation not to do X . The second procedure requires the judges to privately decide about P and Q and to publicly express their opinions on R only. The defendant will be declared liable if and only if a majority of the judges actually believes that she is liable.

In [5] it has been argued that PBP and CBP are not satisfactory methods for group decision-making and proposed that belief fusion is a superior procedure from a theoretical point of view. However, situations like the court example do not only require that the group makes the *right* decision. The defendant factually is (or is not) guilty: There exists an objective truth that the court is trying to reach. And therefore the question that arises is how good a specific fusion operator is as a truth-tracker. Bovens and Rabinowicz in [1] and List in [4] addressed this question already for PBP and CBP. We follow their account and ask how well fusion does compared to these two procedures.

2. THE FUSION PROCEDURE

As shown in [5], the application of a fusion operator defined in [3] to judgment aggregation problems allows to define consistent group decisions and to avoid paradoxical outcomes. This section summarizes the approach and the results of [5]. The reader is referred to that paper for more details.

One of the key points in the literature on information fusion is that the aggregation of consistent knowledge bases does not guarantee a consistent collective outcome. To overcome this problem, *integrity constraints* (IC) are imposed on the final base.

We consider a finite set N of individuals making their judgments on a given finite set X of propositions (*agenda*). A base K_i of an agent i is a consistent and complete finite set of atomic propositions and compound propositions expressing the agent's judgments.

	K_1	K_2	K_3	K_4	K_5	K_6	K_7	$\mathcal{F}_{IC}(E)$
(1,1,1)	0	0	0	2	2	2	2	8
(1,0,0)	2	2	2	0	0	2	2	10
(0,1,0)	2	2	2	2	2	0	0	10
(0,0,0)	3	3	3	1	1	1	1	13

Table 2: The fusion operator applied to the doctrinal paradox

In a model-based fusion approach, each individual base is interpreted as a set of models. Given a multi-set $E = \{K_1, K_2, \dots, K_n\}$ and IC , a fusion operator \mathcal{F} assigns a belief base to E and IC . Let $\mathcal{F}_{IC}(E)$ denote the resulting collective base.

A majority fusion operator will select the (eventually more than one) model that minimizes the *Hamming distance* to the individual bases. The Hamming distance is defined as the number of propositional letters on which two models differ. For example, the Hamming distance between $\omega = (0, 0, 0, 1)$ and $\omega' = (1, 1, 0, 0)$ is $d(\omega, \omega') = 3$.

The first step is to determine the distance between the models of IC and the models of each K_i in E . The next step is to assign a distance value to each model of IC and E . This is defined by the sum of the Hamming distances defined before.

To illustrate how the majority fusion operator works, we apply it to our initial court example. In the new terminology, the agenda is $X = \{P, Q, R\}$ with $IC = \{(P \wedge Q) \leftrightarrow R\}$. The models for each belief base are the following:

$$\begin{aligned} \text{Mod}(K_1) &= \text{Mod}(K_2) = \text{Mod}(K_3) = \{(1, 1, 1)\} \\ \text{Mod}(K_4) &= \text{Mod}(K_5) = \{(1, 0, 0)\} \\ \text{Mod}(K_6) &= \text{Mod}(K_7) = \{(0, 1, 0)\} \end{aligned}$$

Table 2 shows the result of the IC majority fusion operator on $E = \{K_1, \dots, K_7\}$. The row with a shaded background correspond to the selected collective outcome.

The possible collective outcomes are chosen among the interpretations that are models of IC . Thus, no paradox arises by using this fusion operator. We should mention that the fusion operator does not necessarily select a unique group decision. In some cases, the operator selects a set of models, i.e. the result is a tie between some belief bases.

3. THE FRAMEWORK

Bovens and Rabinowicz [1] investigate the case $(P \wedge Q) \leftrightarrow R$ and make the following assumptions: (i) the prior probability that P and Q are true are equal (q), (ii) all voters have the same competence to assess the truth of P and Q (p), (iii) P and Q are (logically and probabilistically) independent, and (iv) each individual judgment set is logically consistent. Hence, only four situations are possible: $S_1 = \{P, Q, R\} = (1, 1, 1)$, $S_2 = \{P, \neg Q, \neg R\} = (1, 0, 0)$, $S_3 = \{\neg P, Q, \neg R\} = (0, 1, 0)$ and $S_4 = \{\neg P, \neg Q, \neg R\} = (0, 0, 0)$. We can now calculate the probability that fusion ranks the right judgment set first (let us denote this proposition by F) by observing that $\mathcal{P}(F) = \sum_{i=1}^4 \mathcal{P}(F|S_i) \cdot \mathcal{P}(S_i)$. Thus, we have to calculate the prior probabilities $\mathcal{P}(S_i)$ and the conditional probabilities $\mathcal{P}(F|S_i)$ for $i = 1, \dots, 4$. From assumption (i), we obtain (with $\bar{x} := 1 - x$) that $\mathcal{P}(S_1) = q^2$, $\mathcal{P}(S_2) = \mathcal{P}(S_3) = q\bar{q}$ and $\mathcal{P}(S_4) = \bar{q}^2$. To calculate the conditional probabilities, suppose that S_1 is the right judgment




Figure 1: The probability that merging (solid line), PBP (dotted line) and CBP (dashed line) identify the right situation as a function of the competence p for $N = 21$ and $q = .2$.

set. Then n_i (of N) voters will vote for profile S_i , with $n_1 + n_2 + n_3 + n_4 = N$. The majority fusion operator selects the model that minimizes the distance to the profiles. This means that, if S_1 is the right judgment set, belief fusion tracks the truth if $d_1 \leq \min(d_1, \dots, d_4)$. The distances d_i can be expressed in terms of the numbers n_i of voters for the situations S_i ($i = 1, \dots, 4$): $d_1 = 2n_2 + 2n_3 + 3n_4$, $d_2 = 2n_1 + 2n_3 + n_4$, $d_3 = 2n_1 + 2n_2 + n_4$ and $d_4 = 3n_1 + n_2 + n_3$. We can now calculate $\mathcal{P}(F|S_1)$:

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

The sum is constrained: $\mathcal{C}(n_1, \dots, n_4) = 1$ if (i) $\sum_{i=1}^4 n_i = N$ and (ii) $d_1 \leq \min(d_1, \dots, d_4)$. Otherwise $\mathcal{C}(n_1, \dots, n_4) = 0$. $\mathcal{P}(F|S_2)$, $\mathcal{P}(F|S_3)$ and $\mathcal{P}(F|S_4)$ can be obtained analogously. Putting everything together, we calculate $\mathcal{P}(F)$.

4. RESULTS

Figure 1 compares the fusion procedure with the PBP and the CBP for the question how good they are in selecting the correct judgment set for $N = 21$ voters and $q = .2$. We see that all three curves converge to 1 if the competence of the voters approaches 1. Moreover, we observe that the fusion operator outperforms PBP and CBP for small and middling values of the competence (i.e. for $p \approx .5$). However, it is no surprise that the second best procedure is the premise-based one. In fact, Bovens and Rabinowicz have shown already that we should prefer the PBP to the CBP if we want to reach the right decision for the right reasons. For high values of competence p , the premise-based procedure turns out to be slightly better than the fusion procedure as a truth-tracker.

We now turn to evaluate how the three approach do when it comes to identify the right result, but not necessarily the right reasons. It turns out that the fusion operator greatly outperforms the PBP and the CBP for small size groups. Yet, as the size of the group increases, both the PBP and the CBP do better than the fusion procedure for low values in competence, and the PBP does better than fusion for high values of p . But, for middling values of p , fusion is always




Figure 2: The probability that merging (solid dotted line), PBP (dotted line) and CBP (dashed line) identify the right outcome as a function of the competence p for $N = 51$ and $q = .5$.

superior. We can also observe that, whenever fusion is not the best procedure, it lies in-between the premise-based and the conclusion-based procedures. Due to space limitations, we show only the result for $N = 51$ and $q = .5$ in Figure 2.

In future work, we want to systematically compare the various procedures for different values of q and for different numbers of voters. We will also test the stability of our main result — that fusion does better for middling values of the competence parameter — for different logical rules (more premises, disjunctions etc.).

5. REFERENCES

- [1] L. Bovens and W. Rabinowicz. Democratic answers to complex questions. An epistemic perspective, *Synthese*, 150: 131–153, 2006.
- [2] S. Konieczny and E. Grégoire. Logic-based approaches to information fusion. *Information Fusion*, 7: 4–18, 2006.
- [3] S. Konieczny and R. Pino-Pérez. Merging with integrity constraints. In *Fifth European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU'99)*, pages 233–244, 1999.
- [4] C. List. The discursive dilemma and public reason. *Ethics*, 116(2): 362–402, 2006.
- [5] G. Pigozzi. Belief merging and the discursive dilemma: an argument-based account to paradoxes of judgment aggregation. *Synthese*, 152(2): 285–298, 2006.