Aggregation of probabilistic logically related judgments

Magdalena Ivanovska* Marija Slavkovik†

Abstract

Information aggregation is at the core of many problems in computer science. Judgment aggregation models multi-agent decision making by aggregating individual opinions from various sources. It does however assume that the sources provide Boolean opinions which are subject to the same logical constraints. We relax both of these assumptions and build a more general framework with uncertain information that we model in probabilistic logic. We also propose aggregation functions for this new framework.

Introduction

	p	$p \rightarrow q$	q
Minister 1	true	true	true
Minister 2	true	false	false
Minister 3	false	true	false
Majority	true	true	false

Table 1: An example of a judgment aggregation.

Judgment aggregation [2] is concerned with aggregating sets of truth valuations assigned to logically related issues. Consider an example from [5] with the following issues: Current CO_2 emissions lead to global warming (p); If current CO_2 emissions lead to global warming, then we should reduce CO_2 emissions ($p \rightarrow q$); We should reduce CO2 emissions (q). Each of these issues can be assigned values "true" or "false" as in Table 1. Taking the majority of truth valuations on an issue as the collective judgment for that issue does not always lead to a consistent set of collective judgments, as illustrated with this example.

Various collective decision making problems in artificial intelligence (AI) can be modelled as judgment aggregation problems such as the problem of finding a collective goal in multi agent systems [6] or when one agent needs to find the truth valuations for each of a set of issues by merging information from various sources. We give a simple example.

	$s \vee t$	h	x	$\neg e$
IS 1	0.6	1	1	1
IS 2	0.7	0.5	0.6	1
IS 3	0.1	0.2	0.4	1
IS 4	0.8	0.9	0.8	1
IS 5	0.7	0.4	0.7	1
IS 6	0.5	0.3	0.6	1

Table 2: An example of a source aggregation.

Consider a hotel recommender agent that compiles recommendation for a user by analysing and merging information on hotels from online sources. The user can specify what are the features of a good hotel such as: either close to the centre or well connected with public transport ($s \vee t$); the hotel is a unique experience (x), the hotel is a good value for money (a). The agent should recommend the hotel (h) iff $s \vee t$ and a are assigned a judgment true, i.e. $((s \vee t) \wedge a) \leftrightarrow h$. The user considers a hotel is a good value

*magdalei@ifi.uio.no, University of Oslo

†marija.slavkovik@uib.no, University of Bergen

This paper was presented at the NIK-2018 conference; see http://www.nik.no/.

for money (a) if it is not more than 80 Euro per night ($\neg e$) or it is a unique experience (x), so $(\neg e \vee x) \rightarrow a$. The reviewers, and the hotels, are typically not available for questioning and specifying their information for each recommender agent, but the price of the service would most likely be available online. Ideally each information source (IS) will be processed automatically, meaning that the agent will have the likelihood that a judgment is confirmed by an information source rather than a Boolean truth value assignment [1]. An example of an input from six sources is given in Table 2. We need to be able to aggregate these probabilistic judgment from Table 2 into a set of Boolean judgments for each of the propositions in $\{s \vee t, h, x, a\}$ so that they are also consistent with $\{(\neg e \vee x) \rightarrow a, ((s \vee t) \wedge a) \leftrightarrow h\}$. It is unreasonable to expect that all the sources would satisfy the user set constraints. To model more realistically the aggregation problem, we need to extend the existing framework into a probabilistic one and design aggregation functions for it. This is our object of study.

Probabilistic judgment aggregation

We distinguish between an agenda setter and information sources (agents). The *agenda setter* identifies the set of *issues*, *i.e.*, the agenda for which judgments need to be made. The *agenda* Φ is a finite set of propositional logic formulas,

$$\Phi = \{\varphi_1, \ldots, \varphi_n\},$$

such that φ_i is neither a tautology nor a contradiction. We define the set of *propositional constraints* as a set of propositional formulas Γ representing special relations that should hold among the agenda issues. Γ should be satisfiable, and we allow $\Gamma = \{\top\}$. The agenda setter is interested in aggregating collections of judgments on the agenda issues from various information sources into a set of Boolean judgments that is consistent with Γ.

A *Boolean judgment* [4] on $\varphi \in \Phi$ is either φ or $\neg\varphi$. A *(crisp) set of judgments* J on the agenda Φ is a set containing only Boolean judgments on the elements of Φ. For example, $J = \{\varphi_1, \neg\varphi_2\}$. The presence of φ in J is interpreted as issue φ being assigned the value “true" and the presence of $\neg\varphi$ in J is interpreted as issue φ being “false" according to the judgment set J. The judgment set J is *complete* if it contains one Boolean judgment for each of the issues in the agenda. If the crisp judgment set J is consistent and complete, we say that it is *rational*.

A *probabilistic judgment* on the issue $\varphi \in \Phi$ is a simple likelihood formula of the type

$$\ell(\varphi) \geq^* a,$$

where $\geq^* \in \{\geq, =\}$ and $a \in [0,1]$.[1] The judgment $\ell(\varphi) \geq a$ expresses that the likelihood (probability) of the statement φ being true is at least a. It immediately implies that $\ell(\neg\varphi) \leq 1 - a$. $\ell(\varphi) = a$ is a stronger statement expressing that the likelihood of φ being true is exactly a. (And $\ell(\neg\varphi) = 1 - a$ correspondingly). $\ell(\varphi) = a$ implies $\ell(\varphi) \geq a$. In general, if $\ell(\varphi) \geq^* a$ implies $\ell(\varphi) \geq^* b$ ($a \geq b$) we will say that $\ell(\varphi) \geq^* a$ is a *stronger judgment* then $\ell(\varphi) \geq^* b$.

A *probabilistic judgment set* $\hat{J}$ is a set consisting of one probabilistic judgment on each of the issues in the agenda:

$$\hat{J} = \{\ell(\varphi_1) \geq^* a_1, \ldots, \ell(\varphi_n) \geq^* a_n\}. \tag{1}$$

[1] The likelihood formulas that we use here are taken from the language for reasoning about likelihood [3]. The general form of a linear likelihood formula is $a_1\ell(\varphi_1) + \ldots + a_n\ell(\varphi_n) \geq a$. This language is interpreted over measurable probability structures and the corresponding axiomatic system is based on the properties of probability functions.

Defined like this, a probabilistic judgment set is always *complete* since, in the case of "abstention" on an issue φ, we assume the tautology $\ell(\varphi) \geq 0$.

In the probabilistic case, consistency and completeness are not enough of conditions for rationality. For example, $\hat{J} = \{\ell(p_1) \geq 0.3, \ell(p_1 \wedge p_2) \geq 0.4, \ell(p_1 \wedge \neg p_2) \geq 0.1\}$ is a consistent judgment set. However, the second formula in it implies that $\ell(p_1) \geq 0.4$, which is a stronger judgment then the existing $\ell(p_1) \geq 0.3$ and, as such, is a more valuable judgment. To ensure that we always have the strongest possible judgments in the consistent judgment sets, we introduce the notion of a *final judgment*. A consistent probabilistic judgment set is *final* if it does not imply stronger judgmnts than the ones it contains. A probabilistic judgment set $\hat{J}$ is *rational* if it is consistent and final.

Probabilistic judgments can be subject to *probabilistic constraints* $\hat{\Gamma}$ to denote that certain combinations of issues must have a certain likelihood. For example for agenda $\Phi = \{p_1, p_2, p_3\}$, where p_1, p_2, and p_3 represent the three possible states of a random variable, we can have the integrity constraint $\ell(p_1) + \ell(p_2) + \ell(p_3) = 1$. Unlike the constraints Γ which are given by the agenda setter, the integrity constraints $\hat{\Gamma}$ describe facts of the world and we assume that all information sources provide probabilistic judgment sets that satisfy the integrity constraints when these are given. The probabilistic constraints $\hat{\Gamma}$ apply only to probabilistic, but not to crisp judgment sets.

Given the above, we can define our problem of interest as follows: Given m rational probabilistic judgment sets on the issues of an agenda Φ coming from m different information sources, to aggregate this judgment sets in a meaningful way in order to obtain a rational crisp judgment set as a result. In the following section, we propose a couple of different functions (aggregators) for this purpose.

Aggregators

A *(crisp) profile* P for m agents is a finite sequence of rational judgment sets $P = (J_1, \ldots, J_m)$, where J_k, $k = 1, \ldots, m$ is the judgment set of the agent k. A *probabilistic profile* is a sequence of probabilistic judgment sets:

$$\hat{P} = (\hat{J}_1, \ldots, \hat{J}_m), \text{where } \hat{J}_k = \{\ell(\varphi_1) \geq^* a_1^k, \ldots, \ell(\varphi_n) \geq^* a_n^k\}$$

is the judgment set coming from the kth information source.

Given a probabilistic judgment $\ell(\varphi) \geq^* a$, we can obtain a crisp judgment by choosing a coefficient $c \in (0, 1]$:

$$\text{crisp}(\ell(\varphi) \geq^* a, c) = \begin{cases} \varphi & \text{iff } a \geq c \\ \neg\varphi & \text{iff } a < c \end{cases} \tag{2}$$

In the above definition, the coefficient c acts as a "strictness" coefficient for the crispifying. Namely, it determines how probable the issues should be (in the least) in order to consider them as true. Obviously, its value is context dependent. We assume that it is determined by the agenda setter. We can crispify a probabilistic judgment set $\hat{J}$ by crispifying each of its judgments. Let $\mathbf{c} = (c_1, \ldots, c_n)$ be a given vector of coefficients, where each c_i corresponds to one issue $\varphi_i \in \Phi$ and $c_i \in (0, 1]$ for all $i \in \{1, \ldots, n\}$. Then the crisp judgment set obtained from $\hat{J}$ and $\mathbf{c}$ is defined as

$$\text{crisp}(\hat{J}, \mathbf{c}) = \{\text{crisp}(\ell(\varphi_i) \geq^* a_i, c_i) \mid \ell(\varphi_i) \geq^* a_i \in \hat{J}\}. \tag{3}$$

If $c_i = c$, for every $i = 1, \ldots, n$, for some $c \in (0, 1]$, we denote $\text{crisp}(\hat{J}, \mathbf{c})$ by $\text{crisp}(\hat{J}, c)$ and we call it a *uniform crispifying*.

Uniform quota rules Given a profile $\hat{P}$, a crispifying vector $\mathbf{c}$ and a quota $q \in \mathbb{N}$, $1 \leq q \leq m$, let $N_a(\hat{P}, \varphi_i)$ denote the number of agents for which the likelihood of φ_i is greater that a. Then we can define $e_q(\hat{P}, \varphi_i)$ as the maximal rational number $a \in (0,1)$ s.t. $N_a(\hat{P}, \varphi_i) \geq q$. We can now define the *uniform quota function* $\hat{f}_q$:

$$\hat{f}_q(\hat{P}, \mathbf{c}) = \{\varphi_i : e_q(\hat{P}, \varphi_i) > c_i\} \cup \{\neg\varphi_i : e_q(\hat{P}, \varphi_i) < c_i\}.$$

As an illustration, consider the example in Table 2. For uniform $c = 0.6$ and quota $q = 3$ we obtain $\hat{f}_3(\hat{P}, 0.6) = \{s \vee t, x, \neg h, \neg e\}$, which is not consistent with Γ.

If $q = m$, we obtain the unanimous function that selects as collective only those judgments φ_i who are assigned a likelihood $a_i^k > c_i$ by all the agents k. For $q = \frac{m}{2}$ we obtain the *issue-by-issue majority function*, which we denote with M. Under issue-by-issue majority function the profile is aggregated by selecting the judgments that are in the most of the judgment sets in the profile. The set $\hat{f}_{\frac{m}{2}}(\hat{P}, \mathbf{c}) = M(\hat{P}, \mathbf{c})$ is called a *majoritarian set* for $\hat{P}$.

Aggregating without a quota There are two dimensions that can play a role in determining the likelihood of a judgment determined by a profile: which likelihoods are assigned to that judgment in the profile and by how many agents.

Let $\ell^{\hat{P}}(\varphi) = a$ denote that the likelihood of φ induced by profile $\hat{P}$ is a. The simplest way to determine a is to take the average of likelihoods in the profile.

Given $\varphi_i \in \Phi$ and a profile $\hat{P}$, the average likelihood for φ_i in $\hat{P}$ is defined as $E_{\hat{P}}(\varphi_i) = \frac{1}{m}\sum_{k=1}^{m} a_i^k$, if there exist at least one agent k for which $a_i^k > 0$, otherwise $E_{\hat{P}}(\varphi_i) = 0.5$. For $\neg\varphi_i$ we have that $E_{\hat{P}}(\neg\varphi_i) = 1 - E_{\hat{P}}(\varphi_i)$.

Future work

Once having defined the framework for aggregating probabilistic judgments the immediate next step is to inspect the characterisation of aggregation functions. Namely we can look for the most common desirable properties of aggregation and establish whether there exists an aggregation function that satisfies them. Experience from (computational) social choice theory indicates that an ideal aggregate, one that satisfies all desirable properties, does not exist. Consequently, we need to design a selection of aggregation functions, to cover different aggregation problems. Another immediate step is to consider the computational complexity of calculating the collective judgments.

References

[1] C. X. Aggarwal, C. C.and Zhai, editor. *A Survey of Text Classification Algorithms*. Springer US, Boston, MA, 2012.

[2] D. Grossi and G. Pigozzi. *Judgment Aggregation: A Primer*. Morgan and Claypool Publishers, 2014.

[3] Joseph Y. Halpern. *Reasoning about Uncertainty*. MIT Press, 2003.

[4] L. Lang, P. Pigozzi, M. Slavkovik, L. van der Torre, and S. Vesic. A partial taxonomy of judgment aggregation rules, and their properties. *Social Choice and Welfare*, 48:1–30, November 2016.

[5] C. List and C. Puppe. Judgment aggregation: A survey. In P. Anand, C. Puppe, and P. Pattanaik, editors, *Oxford Handbook of Rational and Social Choice*. Oxford, 2009.

[6] M. Slavkovik and G. Boella. Recognition-primed group decisions via judgement aggregation. *Synthese*, 189(1):51–65, 2012.