

A hierarchical fusion of expert opinion in the TBM

Minh Ha-Duong

CIREC, CNRS

haduong@centre-cired.fr

Abstract

We define a hierarchical method for expert opinion aggregation that combines consonant beliefs in the Transferable Belief Model. Experts are grouped into schools of thought, then opinions are aggregated using the cautious conjunction operator within groups and the non-interactive disjunction across. This method is illustrated with a real-world dataset including 16 experts.

Keywords: Evidence theory, information fusion, expert opinion.

1 Introduction

The aggregation of expert opinion [1] is a challenging application of information fusion. Scientists interact and share evidence, so assuming that opinions are independent overestimates the precision of the results. There can be contradiction among experts, so aggregation methods based on the intersection of opinions do not work well (the intersection is empty). The disagreement between experts may not be a balanced opposition, but rather a dissent minority situation, so aggregation methods like averaging that weight views proportionally to the number of proponents are arguably unbalanced (scientific theories should be evaluated on their own merits). And calibrating the reliability of experts is difficult, so one can not simply assume that some experts are less reliable than others.

To address these challenges, we propose a hierarchical method for the fusion of expert opinion that uses recently published operators for information fusion [2]. Experts are not combined symmetrically, but grouped into schools of thought. Within groups, beliefs are combined using the cautious conjunction rule, and across with the non-interactive disjunction. This is illustrated with a dataset derived from a real-world expert elicitation study [3].

Section 2 reminds the basic mathematical definitions of Transferable Belief Model relevant for information fusion [2]. Section 3 presents the dataset, the implementation, and discusses the proposed hierarchical approach. Section 4 compares theoretically and numerically the hierarchical approach with other fusion procedures, then section 5 concludes.

2 The Transferable Belief Model

The Transferable Belief Model represents and combines uncertain beliefs elaborating on Dempster-Shafer evidence theory [2].

2.1 Basic Belief Assignments

Let us denote Ω a frame of reference, that is a finite set of mutually exclusive states of the world. Uncertainty is represented by allocating the unit mass of belief among *subsets* of the frame of reference Ω . Let 2^Ω denote the power set of Ω . Its elements will be denoted with upper case letters such as $A \subseteq \Omega$ or $X \subseteq \Omega$. The empty subset $\{\}$ will be denoted \emptyset . A basic belief assignment (BBA) is a function $m : 2^\Omega \rightarrow [0, 1]$ such that:

$$\sum_{A \subseteq \Omega} m(A) = 1 \quad (1)$$

The mass $m(A)$ is the portion of the total belief supporting A which do not support more precisely any specific subset of A . As an example, consider a drawing from an urn containing white, black, and red marbles ($\Omega = \{\text{white, black, red}\}$). Knowing only that there is $1/3$ of white marbles would lead to the basic belief assignment defined as: $m(\{\text{white}\}) = 1/3$, $m(\{\text{black, red}\}) = 2/3$. This is not the same as drawing from an urn known to have $1/3$ of each color, which would be represented with the BBA defined as: $m(\{\text{white}\}) = m(\{\text{black}\}) = m(\{\text{red}\}) = 1/3$.

For any subset $A \subseteq \Omega$, the BBA that represents the certain belief that the state of the world is in A is the indicator function $\mathbf{1}_A : 2^\Omega \rightarrow [0, 1]$ defined by:

$$\begin{cases} \mathbf{1}_A(A) = 1 \\ \mathbf{1}_A(X) = 0 \quad \text{if } X \neq A \end{cases} \quad (2)$$

The Transferable Belief Model allows non-zero belief mass to the empty set. The number $m(\emptyset)$, called weight of conflict, represents the internal contradiction arising when beliefs result from information sources pointing in different directions. The extreme case $\mathbf{1}_\emptyset$ represents being completely confounded by contradictory information sources. The vacuous BBA $\mathbf{1}_\Omega$ (no information) can be opposed to the completely confounded BBA $\mathbf{1}_\emptyset$ (too much information), but expert opinion aggregation methods that lead to them are equally unhelpful for decision-making. One way to cancel the weight of contradiction in a BBA m is to *renormalize*, that is replace it by the BBA m^* defined as:

$$\begin{cases} m^*(\emptyset) = 0 \\ m^*(A) = \frac{m(A)}{1-m(\emptyset)} \quad \text{if } A \neq \emptyset \end{cases} \quad (3)$$

Measures of *belief* and *plausibility* are associated with a BBA m as follows. The level of belief in an event $X \subseteq \Omega$, denoted $bel(X)$, is the strength of conviction that X must happen.

The level of plausibility, denoted $pl(X)$, is the strength of conviction that X could happen. With the special case $bel(\emptyset) = pl(\emptyset) = 0$, these functions are defined when $X \neq \emptyset$ as :

$$bel(X) = \sum_{\substack{A \subseteq X \\ A \neq \emptyset}} m(A) \quad (4)$$

$$pl(X) = \sum_{\substack{A \subseteq \Omega \\ A \cap X \neq \emptyset}} m(A) \quad (5)$$

An intuitive interpretation of the theory of evidence sees $m(X)$ as a mass of belief that can flow to any subset of X . In this view, $bel(X)$ represents the minimal amount of belief that is constrained to stay within X , while $pl(X)$ the maximal amount of belief that could flow into X .

2.2 Non-interactive fusion operators

The two basic combination rules of the transferable belief model will be denoted \oplus and \odot . They provide a way to compute the “intersection” or the “union” of two experts’ opinions.

Before turning to the formal definitions, these rules will be illustrated on a special case: the fusion of two experts holding certain beliefs. Expert 1 views that the state of the world is in $A \subseteq \Omega$, and expert 2 views that the state of the world is in $B \subseteq \Omega$. Their beliefs are respectively represented with $\mathbf{1}_A$ and $\mathbf{1}_B$.

To start with \odot , consider what the result of the fusion should be when one thinks that either expert 1 or expert 2 is a reliable information source. In this case, one is led to believe that the state of the world is in A or B , that is in $A \cup B$. The \odot combination rule is precisely such that $\mathbf{1}_A \odot \mathbf{1}_B = \mathbf{1}_{A \cup B}$. It is called the non-interactive disjunction rule.

The non-interactive conjunction rule \oplus is meant to be used when one thinks that both expert 1 and expert 2 are reliable information sources. The fusion of the two opinions should be the belief that the state of the world is in $A \cap B$. In the transferable belief model, this is possible even if the experts have no common ground, that is $A \cap B = \emptyset$, since $\mathbf{1}_\emptyset$ is a BBA.

For reasons that will become apparent with equation 9, we define these two combination rules with more general functions than BBAs, using the letter μ to denote a real-valued subset function $\mu : 2^\Omega \rightarrow \mathfrak{R}$ which verifies equation 1, but may take values outside of $[0,1]$. The non-interactive conjunction of μ_1 and μ_2 is defined as the function $\mu_1 \odot \mu_2 : 2^\Omega \rightarrow \mathfrak{R}$ such that, for any subset X :

$$(\mu_1 \odot \mu_2)(X) = \sum_{\substack{A \subseteq \Omega \\ B \subseteq \Omega \\ A \cap B = X}} \mu_1(A) \times \mu_2(B) \quad (6)$$

In the same way, \oplus is defined by:

$$(\mu_1 \oplus \mu_2)(X) = \sum_{\substack{A \subseteq \Omega \\ B \subseteq \Omega \\ A \cup B = X}} \mu_1(A) \times \mu_2(B) \quad (7)$$

These operators are commutative, associative and if μ_1 and μ_2 are two BBAs then the result is also a BBA. These properties allow to treat the experts symmetrically when combining their opinions.

Consider $\Omega = \{a, b\}$ and the BBA m defined by $m(\{a\}) = m(\{b\}) = 1/2$. Then $(m \odot m)(\{a\}) = (m \oplus m)(\{a\}) = 1/4$, and $(m \odot m)(\emptyset) = 1/2$. That much weight of conflict may seem surprising. One way to avoid this is to renormalize, see equation 3. The renormalized non-interactive conjunction is usually denoted \oplus and called Dempster's combination rule:

$$m_1 \oplus m_2 = (m_1 \odot m_2)^* \quad (8)$$

In some situations the surprising result is meaningful, consider for example a setting in which two players simultaneously replicate a large number of fair coin tosses. Both conclude that $p(\text{Head}) = p(\text{Tail}) = 1/2$. But if the experiments are independent, then the players disagreed half the time. The non-interactive conjunction \odot is relevant to combine information sources only when some kind of independence relation can be assumed between information sources.

2.3 Cautious conjunction

BBAs can be factorized as follows. For any proper subset $A \subsetneq \Omega$ and real number¹ s , denote A^s the function $\mu : 2^\Omega \rightarrow \mathfrak{R}$ such that:

$$\begin{cases} \mu(\Omega) = e^{-s} \\ \mu(A) = 1 - e^{-s} & \text{(we assumed } A \neq \Omega) \\ \mu(X) = 0 & \text{if } X \neq A \text{ and } X \neq \Omega \end{cases}$$

These functions A^s can be used to factorize any BBA m such that $m(\Omega) > 0$ [5]. More precisely, for any such BBA m there is a unique function² $s : 2^\Omega \setminus \Omega \rightarrow \mathfrak{R}$ (negative values are allowed) such that:

$$m = \bigodot_{\substack{A \subseteq \Omega \\ A \neq \Omega}} A^{s(A)} \quad (9)$$

If $s < 0$, then A^s is not a BBA. We suggest to interpret A^s as the change in one's beliefs realized by giving confidence s to a new piece of evidence stating that the state of the world is in A . Positive infinity for s represents a perfectly convincing proof, a limit case excluded in the above definition. Negative weight $s < 0$ have an algebraic justification comparable to negative numbers: considering A with weight s exactly counterbalances considering A with weight $-s$. It has been suggested that negative values of s should be used for information sources known to be deliberately lying [5].

If two BBA m_1 and m_2 admit corresponding weight functions s_1 and s_2 , then:

¹The letter s stands for "Shafer's weight of evidence". This was previously denoted w [4, Chapter 5], but in the recent literature w denotes the "weight of evidence" defined by $w = e^{-s}$ [2].

²The weights can be computed as follows, where $|X|$ denotes the number of elements (cardinality) of a subset $X \subseteq \Omega$:

$$q(X) = (m \odot \mathbf{1}_X)(X) = \sum_{X \subseteq A} m(A) \quad (10)$$

$$s(X) = \sum_{X \subseteq A} (-1)^{|X|-|A|} \ln(q(A)) \quad (11)$$

q is called the *commonality* function. It can be interpreted as the amount of belief that can flow to every point of X .

$$m_1 \odot m_2 = \bigcirc_{\substack{A \subset \Omega \\ A \neq \Omega}} A^{s_1(A)+s_2(A)} \quad (12)$$

This highlights the intuition behind the \odot operator. The non-interactive conjunction adds up distinct pieces of evidence. For example, when combining two experts which point exactly in the same direction A with the same weight s , the result is $A^s \odot A^s = A^{2s}$. It models a situation in which two pieces of distinct evidence pointing out in the same direction add up to stronger beliefs.

To combine experts who share evidence, [2] defined the cautious conjunction operator, denoted \otimes . It combines any two BBAs such that $m_1(\Omega) > 0$ and $m_2(\Omega) > 0$ by taking the maximum of their weight functions as follows:

$$m_1 \otimes m_2 = \bigcirc_{\substack{A \subset \Omega \\ A \neq \Omega}} A^{\max(s_1(A), s_2(A))} \quad (13)$$

If m_1 and m_2 are BBAs, then $m_1 \otimes m_2$ is also a BBA. The combination rule \otimes is commutative and associative, it treats experts symmetrically. It is also idempotent, that is $m \otimes m = m$, and distributes over the noninteractive rule ($m_1 \odot m_2$) \otimes ($m_1 \odot m_3$) = $m_1 \odot$ ($m_2 \otimes m_3$).

Distributivity has an interesting interpretation related to the fusion of beliefs. Consider two experts in the following scenario. Expert 1's belief results from the noninteractive conjunction of two pieces of evidence, $m_1 = A^s \odot B^t$. Expert 2 shares one piece of evidence with expert 1, and has an independent piece, so that $m_2 = A^s \odot C^u$. Then distributivity implies that in the fusion, the shared evidence A^s is not counted twice $m_1 \otimes m_2 = A^s \odot (B^t \otimes C^u)$.

3 An expert aggregation situation

3.1 Dataset and implementation

Climate sensitivity is a measure of the climate change problem. It is denoted $\Delta T_{2\times}$, and defined as the equilibrium global mean surface temperature change following a doubling of

atmospheric CO₂ concentration, compared to pre-industrial levels. The value of this parameter, critical for climate policy, is not known precisely: for a long time, the [1.5°C, 4.5°C] interval has been regarded as the canonical uncertainty range of $\Delta T_{2\times}$.

[3] conducted structured interviews using expert elicitation methods drawn from decision analysis with 16 leading U.S. climate scientists. The authors obtained judgments about a number of key climate variables, including subjective PDFs for the climate sensitivity parameter. In the dataset, no probability is allocated to climate sensitivity lower than -6°C , or larger than 12°C . For the sake of numerical tractability, this range was subdivided in seven ranges at the subdivision points $\{-6, 0, 1.5, 2.5, 3.5, 4.5, 6, 12\}$, and PDFs were discretized to obtain, for each expert, a probability distribution p_i on $\Omega = \{\omega_1, \dots, \omega_7\}$. Then we transformed each p_i into an implicit consonant belief function m_i , using the procedure described in [6]³. Figure 1 represents the beliefs of the 16 experts.

3.2 Discounting beliefs

Most experts BBA m_i verify $m_i(\Omega) = 0$, and cannot be factorized as described above. But various reasons justify to take basic beliefs assignments such that $m(\Omega) = 0$ with a grain of salt. No information source is 100% reliable, especially human ones. Many philosophers consider that fundamentally, scientific knowledge can never be absolute and definitive. On the contrary, it is necessarily based on a possibly large but finite number of human observations, and is always open to revision in front

³Given a probability distribution p , a consonant belief function m is defined as follows. Order the states of the world from most to least probable, that is $p(\omega_{n_1}) > \dots > p(\omega_{n_{|\Omega|}})$. Consider the sets $A_k = \{\omega_{n_1}, \dots, \omega_{n_k}\}$ and assign to A_k the belief mass $m(A_k) = |A| \times (p(\omega_{n_k}) - p(\omega_{n_{k+1}}))$ with the convention that $p_{n_{|\Omega|+1}} = 0$.

This transformation is used as a contingency because our dataset was given as probabilities: possibilistic expert survey datasets are hard to find. This allows to illustrate the hierarchical fusion of consonant beliefs. But to draw substantive conclusions from the numerical results, one would need to discuss this procedure's legitimacy more. That is out of the scope of this paper.

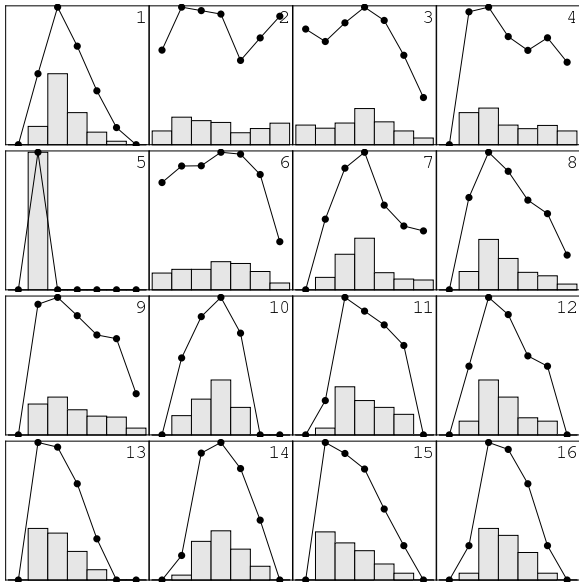


Figure 1: Beliefs of the 16 experts in [3]. The vertical axis goes from 0 to 1. The horizontal axis discretizes the $[-6^{\circ}\text{C}, 12^{\circ}\text{C}]$ climate sensitivity range into seven intervals using a non-uniform subdivision at $-6, 0, 1.5, 2.5, 3.5, 4.5, 6$ and 12°C . On each graph, the grey histogram represents the actual elicited probability distribution discretized from the elicited probability density function P_i (i.e. $p_i(\omega_1) = P_i(-6 \leq \Delta T_{2 \times \text{CO}_2} < 0) \dots$). The dotted lines represents the possibility distribution corresponding to the implicit consonant belief function (see Footnote 3). Four qualitatively different groups of distributions can be seen: Experts in group $G_1 = \{2, 3, 6\}$ allow cooling. Those in group $G_2 = \{4, 7, 8, 9\}$ allow high outcomes but no cooling, $G_3 = \{1, 10 \dots 16\}$ disallow extreme cases, and the opinion of the single expert outlying in $G_4 = \{5\}$ is concentrated on $[0^{\circ}\text{C}, 1^{\circ}\text{C}]$. Better ways to group experts together are discussed in section 3.3, but this heuristic is sufficient to illustrate the hierarchical method discussed here.

of new experimental evidence. And finally, the elicitation of expert's opinions, for example by asking them probability density functions, was necessarily coarse. Experts who allocated no significant probability weight to extreme outcomes might have agreed that there was a very small possibility.

Discounting is a simple way to add doubt to a basic belief assignment. Let r be a number in $[0, 1]$ called a *reliability factor*. Discounting the BBA m means replacing it by:

$$\text{dis}(m, r) = rm + (1 - r)\mathbf{1}_{\Omega} \quad (14)$$

Discounting can be seen as a technical fix to ensure that beliefs can be factorized and combined using the cautious operator. In this case, a reliability factor close to 1 such as $r = 0.999$ is legitimate. Or discounting can be seen as justified in principle. In this case a lower reliability factor such as $r = 0.8$ can be used, with the caveat that an exact value of this parameter is difficult to justify.

3.3 Fusion procedures

A simple way to combine expert opinions is to pool all beliefs together symmetrically using a fusion operator like the non-interactive conjunction or disjunction, the cautious conjunction or Dempster's combination rule. Beyond this, we propose a hierarchical fusion procedure. It aims to be relevant when science is not yet stabilized, and the notion of 'competing theories' can be used. Sociology of science suggests that at some moments in the progress of science, in front of a big unexplained problem, scientists tend to group into schools of thought, which correspond to alternative candidate theories.

Within each group, contradiction should be low because experts share an explanation of the way the world works. Experts can all be considered reliable information sources, but the independence assumption does not hold. Beliefs will be combined with a cautious conjunction operator.

To combine across groups, we assume that only one theory will be adopted in the end

and use the non-interactive disjunction operator⁴. This deals with the challenge of representing equally minority views because all theories are treated equally, regardless of the number of experts in the group.

Ways to summarize the opinion of a variety of experts are often divided in behavioral and mathematical methods. In behavioral approaches, also called interactive expert aggregation methods, experts exchange information with each other. In mathematical approaches, each expert is interviewed separately in a first phase, and then opinions are combined afterward according to some algorithmic aggregation method. Behavioral approaches have interesting advantages: the group judgment is more legitimate since it comes from the experts themselves and collective deliberation is a natural social process. But there are known human biases towards conservatism and overconfidence in group-thinking, and managing all the interactions between experts is complicated, time consuming and thus costly. Mathematical methods are cheaper, simpler and aim to rationalize the procedure.

The hierarchical approach is a mathematical method based on a qualitative behavioral analysis: a partition of experts into a small numbers of schools of thought. Although in this paper we determined the groups using the elicited beliefs themselves for illustrative purposes, there are better procedures. For example, the network of experts can be analyzed with catalogues of publications since experts who have published together have seen the same data, they are more likely to share evidence. Another classical method to determine how a group of people is organized is to analyze the content of the semi-structured face-to-face interviews conducted in the expert elicitation. Finally, the experts themselves know their community, they can help to discover how it is organized, and they can validate the results of the sociological analysis.

⁴Using the \odot operator also assumes that schools of thought are non-interactive. This is not essential to the approach: the bold disjunctive combination rule could be used instead [2].

Representing the diversity of viewpoints by a small numbers of schools of thought is admittedly a strong simplification of complex social reality, but treating all experts symmetrically is even simpler. If it were clear from the start what the different schools of thought are, one could select a single expert to represent each position, and then pool the opinions symmetrically. Otherwise, it is only after a formal sociological study of the experts community that the population of experts can be organized around a small number of archetypes.

4 Comparing fusion methods

This section critically assesses the different ways to combine opinions defined above, theoretically and numerically. How do they interact with discounting? Can they be expected to give mathematically interesting results (not $\mathbf{1}_\emptyset$ or $\mathbf{1}_\Omega$) for expert opinion fusion with consonant beliefs? How do they deal with non-independence, complete contradiction, minority views, and discounting?

The discussion will follow the results presented on Figure 2, which allows to compare 8 alternative ways to perform the fusion of expert opinion. On each plot, the vertical axis goes from 0 to 1, and horizontally the numbers (from 1 to 7) denote the states of the world ω_1 to ω_7 . There are three series of points on each plot. The top one is labelled *pl*, while the middle one is labelled *p* and the bottom *bel*. They display respectively the plausibility of singletons $pl(\{\omega_i\})$, the pignistic probability $p^m(\omega_i) = \frac{1}{1-m(\emptyset)} \sum_{\omega \in X} \frac{m(X)}{|X|}$, and the belief $bel(\{\omega_i\})$.

The top left plot presents the result obtained with Dempster's rule, $\bigoplus_{i=1..n} \text{dis}(m_i, 0.8)$. We examine only $r = 0.8$ and not $r = 0.999$ because Dempster's rule \bigoplus without discounting is known to give counter-intuitive results. Consider for example three states of the world, $\Omega = \{A, B, C\}$, and the problem of combining bayesian beliefs corresponding to the two probability distributions p_1 and p_2 , defined respectively by $p_1(A) = 0.9$, $p_1(B) = 0$, $p_1(C) = 0.1$, and $p_2(A) = 0$, $p_2(B) = 0.9$, $p_2(C) = 0.1$. The result ac-

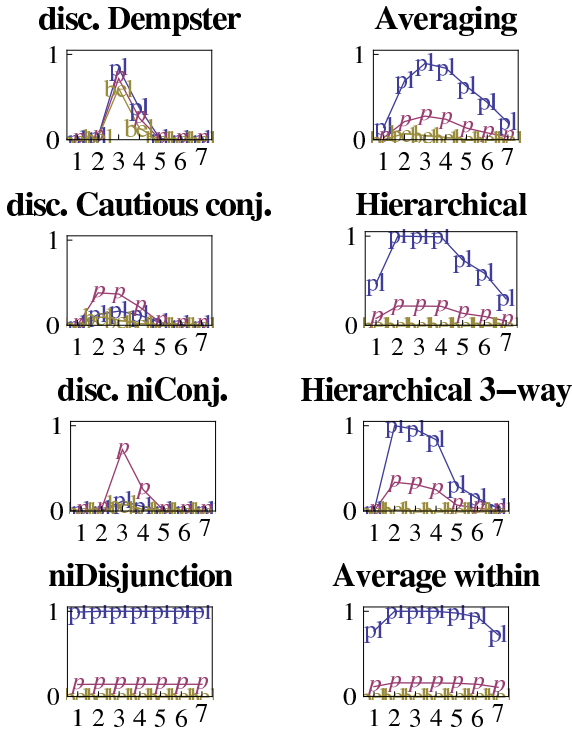


Figure 2: Comparing 8 alternative procedures to fusion expert opinion.

According to Dempster’s rule has a belief weight 0.85 to the state of the world C , which is paradoxical since both information sources agree that this is the least probable outcome. In the same example, if opinions are taken with a reliability factor $r = 0.8$ before combination, the weight going to state of the world C is only 0.105, which is much more intuitive.

Dempster’s rule has many drawbacks for pooling expert opinion. Because of contradiction, it requires discounting, but there is little evidence to determine the reliability factor. It assumes independence, reducing the plausibility of states of the world that are outside the central range (on the figure, beliefs are very focused around ω_3), but much of this precision is unwarranted if experts are not independent.

Consider now the second and third cases in the left column. They show respectively the results obtained with cautious conjunction $\bigwedge_{i=1\dots n} \text{dis}(m_i, 0.8)$ and with non-interactive conjunction $\bigodot_{i=1\dots n} \text{dis}(m_i, 0.8)$. Since these rules produce the trivial result $\mathbf{1}_\Omega$ when the information sources conflict completely, dis-

counting is needed to recover informative results. This is hard to justify, when the whole point of the transferable belief model is to accept $\mathbf{1}_\Omega$ as a theoretically correct result. The precision of the non-interactive conjunction is also highly questionable given that experts interact.

The left bottom plot shows the result of the non interactive disjunction $\bigodot_{i=1\dots n} \text{dis}(m_i, 0.999)$. This operator produces uninformative beliefs close to $\mathbf{1}_\Omega$. Discounting could only make the result even less informative. This operator has potential to combine bayesian, but not consonant, beliefs.

Turn now to the right column and averaging, also called the linear opinion pool: $\frac{1}{n} \sum_{i=1\dots n} \text{dis}(m_i, 0.999)$. It is mathematically equivalent to discount the opinions before averaging, or to discount after averaging. But there is no reason to discount the average opinion, that only means adding unjustified imprecision to the result. So only $r = 0.999$ is interesting.

The theoretical criticism of averaging is that the weight of an opinion increases with the number of experts holding it. Yet scientific arguments should be evaluated on their own merits, not by *argumentum ad populum* (Latin: “appeal to the people”). It is only at the social decision-making stage that the quality and number of people behind each view should matter. Groupthink and bandwagon effects are known dangers when pooling opinions. Thus, a fusion method that gives equal attention to the minority and the majority views would be preferable.

This brings us to the results of hierarchical fusion. Denoting G_1, \dots, G_4 the groups of experts, the second and third plots on the right column show $\bigodot_{k=1\dots N} \bigwedge_{i \in G_k} \text{dis}(m_i, 0.999)$. In the third plot, we merged G_1 and G_2 together for a 3-way hierarchical fusion. This shows that results are significantly sensitive to the clustering of experts: the plausibility of the ‘above 4.5°C’ case, drops from 0.61 to 0.15. The procedure requires only technical discounting ($r = 0.999$) because groups

were constructed so that the degree of conflict among experts in the same group is low.

Lastly, we examined a hierarchical fusion where the first step is averaging, rather than the cautious conjunction. The plausibility levels increase: the result is more ambiguous.

5 Concluding remarks

This paper compared procedures to fusion consonant beliefs in the Transferable Belief Model. Procedures that combine opinions symmetrically have problems to aggregate expert opinion when there is a wide range of competing scientific theories. The non-interactive or the cautious conjunction of all opinions only allow to say ‘Experts contradict each other’. The non-interactive disjunction rule says ‘Everything is possible’. Dempster’s rule of combination lead to overconfidence. And averaging is a way to weight views proportionally to the number of proponents, while scientific theories should be assessed only on their own merit.

The proposed hierarchical fusion procedure is built around a simple model of experts’ social relations: they are divided into schools of thought. Scientific methods are available to determine the structure of epistemic communities. Beliefs are aggregated using the cautious conjunction operator within, and combined using the non-interactive disjunction rule across groups.

This solves several theoretical problems with opinion aggregation. It does not use discounting, thus avoiding the issue of expert calibration. Within groups, cautious conjunction does pool together distinct streams of evidence to make beliefs firmer. But it is not assumed that opinions are, statistically speaking, independent: this would overestimate the precision of actual information. Disjunction allows to deal with complete contradiction among opinions without falling into degenerate results or paradoxes. When several scientific theories compete to explain the same observations, it should not be assumed that both are true at the same time (conjunction), but that at least one will remain

(disjunction). Pooling opinions across schools of thoughts, rather than across individual experts, is arguably a more balanced procedure. And unlike averaging, minority views are equally taken into account.

Even with purely mathematical expert aggregation methods, one has to make sure that no major point of view is omitted when selecting the experts. Therefore sociological considerations on the population of experts cannot really be avoided. The hierarchical approach brings forward transparently that qualitative analysis. Finding out the detailed structure of epistemic communities to explain the differences between theories may be as policy-relevant as computing aggregate beliefs.

Acknowledgements

We gratefully thank David Keith for the data and acknowledge the precious comments of T. Denœux and CIREN colleagues.

References

- [1] R. M. Cooke. *Experts in Uncertainty. Opinion and Subjective Probability in Science*. Oxford U. Press, 1991.
- [2] T. Denœux. Conjunctive and disjunctive combination of belief functions induced by non distinct bodies of evidence. *Art. Intel.*, 172(2–3):234–264, 2007.
- [3] M. Granger Morgan and D. W. Keith. Subjective judgments by climate experts. *Env. Sci. and Tech.*, 29(10):468A–476A, October 1995.
- [4] G. Shafer. *A Mathematical Theory of Evidence*. Princeton U. Press, Princeton (NJ), 1976.
- [5] P. Smets. The canonical decomposition of a weighted belief. In *International Joint Conference on AI*, pages 1896–1901, 1995.
- [6] P. Smets. Quantified epistemic possibility theory seen as an hyper cautious Transferable Belief Model. In *Rencontres Francophones sur la Logique Floue et ses Applications (LFA 2000)*, pages 343–353. Cepadues-Editions, 2000.