Combining imperfect knowledge of maybe corrupted sources

Gérald Gavin
Laboratory ERIC
University of Lyon, France
Email: gerald.gavin@univ-lyon1.fr

Stéphane Bonnevay
Laboratory ERIC
University of Lyon, France
Email: stephane.bonnevay@univ-lyon1.fr

***Abstract*—Many data management applications, such as setting up Web portals, managing enterprise data, managing community data, and sharing scientific data, require integrating data from multiple sources. Each of these sources provides a set of values and different sources can often provide conflicting values. To discover the true values, data integration systems should resolve conflicts. In this paper, we present a formal probabilistic framework in the expert/authority setting. Each expert has a partial and maybe imperfect view of a binary target tuple $\boldsymbol{b}$ that an authority wishes recovering. The goal of this paper consists of proposing a multi-party aggregating function of experts' views to recover $\boldsymbol{b}$ with an error rate as small as possible. In addition, it is assumed that some of the experts are corrupted by an adversary $\mathcal{A}$. This adversary controls and coordinates the behavior of the corrupted experts and can thus perturb the aggregating process. In this paper, we present a simple aggregating function and we provide a formal upper-bound over of the output tuple error expectation in the worst case, i.e. whatever the behavior of the adversary is.**

I. Introduction

Fusion of conflicting data, when for instance several experts have very different ideas about the same phenomenon, has long been identified as a challenging task in the data fusion community. The inherent imperfection of data is the most fundamental challenging problem of data fusion systems, and thus the bulk of research work has been focused on tackling this issue. In [3], the authors distinguish two kinds of data conflict: (a) uncertainty about the attribute value, caused by missing information; and (b) contradictions, caused by different attribute values.

There are a number of mathematical theories [7] available to represent data imperfection [11], such as probability theory [5], fuzzy set theory [12], possibility theory [8], rough set theory [10], and Dempster-Shafer evidence theory (DSET) [6]. Most of these approaches are capable of representing specific aspect(s) of imperfect data. For example, a probabilistic distribution expresses data uncertainty, fuzzy set theory can represent vagueness of data, and evidential belief theory can represent uncertain as well as ambiguous data.

In [2], the authors present a novel approach that considers dependence between data sources in truth discovery. Intuitively, if two data sources provide a large number of common values and many of these values are rarely provided by other sources (e.g., particular false values), it is very likely that one copies from the other. They apply bayesian analysis to decide dependence between sources and design an algorithm that iteratively detects dependence and discovers truth from conflicting information. They also extend their model by considering accuracy of data sources and similarity between values.

In this paper, each source/expert has a partial and maybe imperfect view of a target binary database $\boldsymbol{b}$ which can be seen as a binary tuple. For concreteness, each expert only knows a subset of components of $\boldsymbol{b}$ with maybe some errors. In our model, the views of the experts are drawn according to an arbitrary probability distribution D belonging to a given family of probability distributions $\mathcal{D}$. In addition, we consider the existence of an unique entity, called an adversary, totally controlling a minority of sources/experts: in particular, it knows each corrupted expert view but not the views of uncorrupted experts (called honest experts). This adversary can be seen as an active noise generator. This paper aims to build a way to aggregate expert views in order to recover $\boldsymbol{b}$ whatever the behavior of the adversary is. As far as we know, it is the first time that such an assumption is considered in data fusion. The novelty of our approach makes difficult comparisons with other existing solutions. In our opinion, the main interest of the paper is that the proposed solution is totally formalized making clear the assumptions and the results.

II. Problem statement

Just to illustrate this problem from a police investigation perspective, let us consider T witnesses which have partially seen a crime scene and have identified a set $\mathcal{N}$ of n suspects. One of these witnesses is assumed to be the police. We define the tuple $\boldsymbol{b} = (b_1, ..., b_n) \in \{-1, 1\}^n$ by $b_i = 1$ if and only if Suspect i is guilty. Witness j knows (maybe erroneously) the culpability or non-culpability of a subset $S_j \subset \mathcal{N}$ of the suspects, i.e., Witness j knows $(b_i)_{i \in S_j}$. An adversary $\mathcal{A}$ has corrupted some of the witnesses. The adversary $\mathcal{A}$ can change the testimony of any corrupted witness. Those witnesses which have not been corrupted are said to be honest. The challenge consists in elaborating a multi-party aggregating function of testimonies allowing the police to recover $\boldsymbol{b}$ (or something close).

[0]978-1-5090-5942-3/17/$31.00 2017 IEEE

Let us consider now an other example in the user/server setting. The server wishes recovering a binary tuple $\boldsymbol{b} = (b_1, \ldots, b_n) \in \{-1, 1\}^n$. For instance, the server could be an online encyclopedia. In order to index its pages, some questions could be asked to the users. For instance, does the i^{th} page contains pornographic pictures? Each user sends this information for a subset of pages. One can imagine that some users are corrupted by an adversary $\mathcal{A}$. How can the server fight against any adversary $\mathcal{A}$ to recover the truth about its pages?

Typically, an authority wishes recovering a binary tuple $\boldsymbol{b}$ partially and imperfectly known by T experts. The knowledge of the j^{th} expert can be represented by a tuple $c_j \in \{-1, 0, 1\}^n$ where $c_{ji} = 0$ means that Expert j does not know b_i. Moreover, if $b_i c_{ji} < 0$ means that b_i is erroneously known by Expert j. In this paper, we assume that $c = (c_1, \ldots, c_T)$ is randomly drawn according to a probability distribution D. Each expert j sends its tuple c_j to the authority. The objective of the authority is to recover $\boldsymbol{b}$ from the values c_{ji}. At this step, the problem is trivial for some probability distributions D. Indeed, if the experts do not input too many incorrect values c_{ji} (i.e. $b_i c_{ji} < 0$) then the majority vote is relevant, i.e. $(\mathsf{sign}(\sum_j c_{ji}))_{i=1,\ldots n} \approx b$.

However, the majority vote strategy could be not relevant anymore if some experts misbehave by sending malicious values c_{ji}. Worse, one can assume the existence of a coalition of experts aiming to perturb the recovering of $\boldsymbol{b}$. For concreteness, these experts decide to collaborate by elaborating a common malicious strategy. A simple way to represent this scenario consists of assuming the existence of an adversary $\mathcal{A}$ which has corrupted a minority $\mathcal{C}$ of experts. $\mathcal{A}$ can be seen as an external unique entity which totally controls and coordinates the behavior of the corrupted experts[1]. In our problem, its power consists in arbitrarily modifying the tuples c_j sent by the corrupted experts. The challenge consists in elaborating an aggregating function Aggregate allowing the authority to recover $\boldsymbol{b}$ (or something close) whatever the behavior of $\mathcal{A}$ is.

Clearly, for some probability distributions D, our problem cannot be solved. Indeed, let us assume that the sets $S_j = \{i \in \{1, \ldots n\} | c_{ji} \neq 0\}$ do not overlap. In this case, each component b_i is known by at most one expert and it is not possible to distinguish a corrupted value from an honest one. Consequently, the adversary $\mathcal{A}$ could generate $\sum_{j \in \mathcal{C}} |S_j|$ errors in the recovering of $\boldsymbol{b}$ and nothing can be done to prevent this (without any other assumption). Thus, only probability distributions D ensuring overlapping makes our problem relevant.

In the next section we propose a formalization of this problem. In Section IV, we propose a function Aggregate exploiting the redundancy of the knowledge of the experts. The principle of this function is very simple. The experts which disagree on too many instances i with too many other experts are eliminated. It follows that a corrupted expert which inputs too many incorrect values is eliminated. Consequently, it should behave almost honestly to not be eliminated. In Section V, we provide an upper-bound for the expectation error of the tuple output by Aggregate in the worst case, i.e., independent of the adversary $\mathcal{A}$. In Section VI, numerical values dealing with simple probability distributions D are given. The analysis of these results shows that Aggregate dramatically outperforms the naive approach consisting of taking a majority vote.

[1]The definition and the properties of the adversaries considered in this paper are directly inspired from the Secure Multiparty Computation framework see ([9], [1])

III. Formalization

Let n, T, ϑ be positive integers s.t. $\vartheta < T/2$, let $\mathcal{N} = \{1, \ldots, n\}$ and let $\Delta = \{-1, 0, 1\}^{T \times n}$. The target tuple is denoted by $\boldsymbol{b} \in \{-1, 1\}^n$ and $\mathcal{J} = \{1, \ldots, T\}$ refers to the set of the T experts. The set of the subsets $\mathcal{C} \subset \mathcal{J}$ s.t. $|\mathcal{C}| \leq \vartheta$ is denoted by $P_\vartheta(\mathcal{J})$.

A. Definition of $\mathcal{D}_\perp$

Definition 1. *Let $\mathcal{V}$ be the set of probability distributions defined over $\{-1, 0, 1\}^n$ and let $\mathcal{D}_\perp$ denote the family of probability distributions over Δ defined by*

$$\mathcal{D}_\perp = \{D_1 \times \cdots \times D_T | D_i \in \mathcal{V}\}$$

Throughout this paper, we will only consider probability distributions $D \in \mathcal{D}_\perp$ (defined over Δ). Let $(c_1, \ldots, c_T) \in \Delta$ be randomly drawn according to $D \in \mathcal{D}_\perp$. By definition of $\mathcal{D}_\perp$, it is ensured that the T tuples $c_1, \ldots, c_T$ are independent. In the setting of this paper, it means that each expert j has generated its tuple c_j independently of the tuples of the other experts. For instance, one can imagine that the experts are anonymous and do not know each other. It can make sense in an open network such as the Web. This will be used to simplify the adversary model by reducing its power. Conversely, one can easily imagine some settings where this assumption is not relevant and further investigations should be done to remove or at least to restrict it.

B. Overview

Let $\mathcal{D} \subseteq \mathcal{D}_\perp$ be a family of probability distributions over Δ, let $D \in \mathcal{D}$ and let $c = (c_1, \ldots, c_T) \in \Delta$ be drawn according to D. The objective is to elaborate a function Aggregate : $\Delta \to \{-1, 1\}^n$ (computing by the authority) inputting c and outputting a binary tuple $\boldsymbol{o} \in \{-1, 1\}^n$ as close as possible to $\boldsymbol{b}$. More precisely, it is desired to minimize the error rate $er(\boldsymbol{o})$ defined as the Hamming distance between $\boldsymbol{b}$ and $\boldsymbol{o}$ divided by $2n$, i.e.,

$$er(\boldsymbol{o}) = \frac{1}{2n} \|\boldsymbol{o} - \boldsymbol{b}\|_1$$

In order to elaborate this function, the authority is not assumed to know D but only $\mathcal{D}$. As suggested in the introduction, we assume the existence of an adversary $\mathcal{A}$ able to corrupt a chosen subset $\mathcal{C}$ of experts. The set of uncorrupted experts (also called honest) is denoted by $\mathcal{H} = \mathcal{J} \setminus \mathcal{C}$. We propose to overestimate the real-life adversary power by assuming that it knows everything except the tuples c_j of uncorrupted experts. In particular, it is assumed that $\mathcal{A}$ knows the target tuple $\boldsymbol{b}$,

the function Aggregate, the probability distribution $D \in \mathcal{D}$ and can replace corrupted experts' inputs by arbitrary values in $\{-1, 0, 1\}$. In other words, $\mathcal{A}$ can arbitrarily modify at most ϑ arbitrarily chosen tuples c_j. Roughly speaking, $\mathcal{A}$ can be seen as an active noise generator. The tuple $\boldsymbol{o}$ outputs by Aggregate should be as close as possible to $\boldsymbol{b}$ regardless of the behavior of $\mathcal{A}$. The authority wishes building a function Aggregate robust against any adversary $\mathcal{A}$ for families $\mathcal{D}$ as large as possible.

C. The adversary model

First, we consider an adversary $\mathcal{A}$ which can control at most $\vartheta < T/2$ experts: we do not see how to fight against adversaries corrupting a majority of experts. It is a quite restricting assumption for some realistic applications but it seems difficult to overcome it.

Moreover, in this paper we only consider probability distributions $\mathcal{D} \subset \mathcal{D}_\perp$. By definition of $\mathcal{D}_\perp$, the tuples $c_1, \cdots, c_T$ are independent. It means that the knowledge of the tuples $(c_j)_{j\in\mathcal{C}}$ is not informative about the "honest" tuples $(c_j)_{j\in\mathcal{H}}$. This can be used to restrict the power of the adversary $\mathcal{A}$ by assuming that it chooses the tuples[2] $(c_j^*)_{j\in\mathcal{C}}$ *a priori*, i.e. before to know the tuples $(c_j)_{j\in\mathcal{C}}$. In other words, $\mathcal{A}$ can be seen as a pair $(\mathcal{C}, \delta) \in P_\vartheta(\mathcal{J}) \times \Delta$ where $\mathcal{C}$ refers to the set of corrupted experts and δ contains the malicious values, i.e. for any $j \in \mathcal{C}$, Expert j inputs the j^{th} row δ_j of δ.

D. The objective

Consider a family of probability distributions $\mathcal{D} \subset \mathcal{D}_\perp$, $D \in \mathcal{D}$, a function Aggregate : $\Delta \to \{-1, 1\}^n$ and an adversary $\mathcal{A} = (\mathcal{C}, \delta)$. We define

$$er_{\mathcal{A}}^{\mathsf{Aggregate}}(D) \stackrel{\mathsf{def}}{=} E(er(\boldsymbol{o}))$$

as the expectation of $er(\boldsymbol{o}) = \frac{1}{2n}\|\boldsymbol{o}-\boldsymbol{b}\|_1$ where $\boldsymbol{o}$ is the tuple output by the following protocol:

Protocol 1.

// In this protocol, it is assumed that all the communications are done via secure channels.

1) *$\mathcal{A}$ corrupts the subset $\mathcal{C}$ of experts*,
2) *Let $(c_1, \ldots, c_T)$ be drawn according to D,*
3) *For each $j \in \mathcal{J}$, Expert j receives*[3] *c_j and sends a tuple c_j^* to the authority defined as follows:*

$$\begin{cases} c_j^* = c_j \text{ if } j \in \mathcal{H} \\ c_j^* = \delta_j \text{ if } j \in \mathcal{C} \end{cases}$$

4) *The authority outputs $\boldsymbol{o} = \mathsf{Aggregate}(c_1^*, \ldots, c_T^*)$*

The authority is interested in building a function Aggregate minimizing

$$er^{\mathsf{Aggregate}}(\mathcal{D}) = \sup_{\mathcal{A}\in P_\vartheta(\mathcal{J})\times\Delta; D\in\mathcal{D}} er_{\mathcal{A}}^{\mathsf{Aggregate}}(D) \quad (1)$$

[2]that it will send to the authority

[3]for instance from real-life.

c_1^*	c_2^*	c_3^*	c_4^*	c_5^*	b	MV	Agg
-1	0	1	1	0	-1	1	-1
0	-1	1	1	0	-1	1	-1
1	1	1	0	0	1	1	1
1	0	1	0	1	1	1	1
1	1	0	0	1	1	1	1
0	1	0	0	1	1	1	1

Fig. 1. Table illustrating the Majority Vote (MV) and Aggregate (Agg) in the case $n = 5, T = 5, \theta = 2$. The sets $\alpha(j)$ computed in $\mathsf{Aggregate}(c_1^*, \ldots, c_5^*)$ are equal to $\alpha(1) = \alpha(2) = \alpha(5) = \{1, 2, 5\}$ and $\alpha(3) = \alpha(4) = \{3, 4\}$. As $|\alpha(3)| = |\alpha(4)| < T-\theta = 3$, Expert 3 and Expert 4 are eliminated, i.e. $w_3 = w_4 = 0$. It follows that $\mathsf{Aggregate}(c_1^*, \ldots, c_5^*)$ consists of taking a majority vote over c_1^*, c_2^* and c_5^*.

for families $\mathcal{D}$ as large and realistic as possible[4].

IV. A proposal for Aggregate

A. Case of perfect honest experts

We first assume that the experts do not receive incorrect values, i.e. $c_{ji}b_i \geq 0$ for any $(i, j) \in \mathcal{N} \times \mathcal{J}$. Consequently, honest experts do not input incorrect values. Because of this assumption, we can define a natural elimination strategy where the honest experts cannot be eliminated and the corrupted experts cannot input too many incorrect values without being eliminated. Our solution exploits the fact that honest experts only input correct values and that the corrupted experts are in a minority. We say that two experts j and j' are **compatible** if they do not disagree on at least one instance $i \in \mathcal{N}$. For any $j \in \mathcal{J}$, $\alpha(j)$ refers to the subset of experts $j' \in \mathcal{J}$ which are compatible with j.

$$\alpha(j) = \{j' \in \mathcal{J} \mid \forall i \in \mathcal{N}\ c_{ji}^* c_{j'i}^* \geq 0\}.$$

Clearly, for each $j \in \mathcal{H}$

$$|\alpha(j)| \geq |\mathcal{H}| \geq T - \vartheta.$$

The definition of the function Aggregate is based on this fact. It simply consists of eliminating experts $j \in \mathcal{J}$ verifying $|\alpha(j)| < T - \vartheta$ before estimating c_i by a majority vote. In other words, the weight w_j of Expert j in the majority vote is defined by

$$w_j = \begin{cases} 1 & \text{if } |\alpha(j)| \geq T - \vartheta \\ 0 & \text{otherwise.} \end{cases}$$

According to the previous discussion, honest experts are not eliminated. Intuitively, if a corrupted expert inputs many incorrect values, then it becomes incompatible with almost all honest experts, implying that it is eliminated because $|\alpha(j)| \approx \vartheta < T - \vartheta$. As a corollary, a corrupted expert should behave "almost honestly" in order to avoid elimination. This naturally leads to the following definition of Aggregate.

[4]As explained in the introduction, $\mathcal{D}$ cannot contain all probability distributions but only *redundant* ones.

$$\mathsf{Aggregate}\,(c_1^*,\ldots,c_T^*) = \left(\mathsf{sign}\left(\sum_{j\in\mathcal{J}} w_j \;\; c_{ji}^*\right)\right)_{i\in\mathcal{N}}$$

B. General case

In this section, we extend the previous study by assuming that honest experts can input incorrect values c_{ji}, i.e. $c_{ji}b_i < 0$. The function Aggregate should be adapted to not exclude too many honest experts. To achieve this, it suffices to strengthen the definition of incompatibility between two experts and to relax the strategy of elimination. The new function Aggregate is parameterized by two new positive integers $0 \leq \sigma \leq n$ and $0 \leq \tau \leq T - 2\vartheta$. Two experts j and j' are **compatible** if they do not disagree on more than σ instances $i \in \mathcal{N}$ and an expert will be eliminated if it is compatible with less than $T - \vartheta - \tau$ experts.

$\mathsf{Aggregate}_{\sigma,\tau}\,(c_1^*,\ldots,c_T^*)$

---

1) $I(j,j') := \{i \in \mathcal{N} \text{ — } c_{ji}^* c_{j'i}^* < 0\}$
   //$I(j,j')$ is the set of components of $\boldsymbol{b}$ where the experts j and j' disagree

2) $\alpha(j) := \{j' \in \mathcal{J} \mid |I(j,j')| \leq \sigma\}$
   //$\alpha(j)$ is the set of the experts compatible with Expert j

3) $w_j := \begin{cases} 1 & \text{if } |\alpha(j)| \geq T - \vartheta - \tau \\ 0 & \text{otherwise.} \end{cases}$
   // w_j is the weight of Expert j in the final vote

4) Output $\left(\mathsf{sign}\left(\sum_{j\in\mathcal{J}} w_j \;\; c_{ji}^*\right)\right)_{i\in\mathcal{N}}$

The previous section deals with the case $\sigma = \tau = 0$. In practice, σ, τ are chosen as large as possible ensuring that the honest experts are eliminated with a very small probability.

V. Analysis

Let $\mathcal{D} \subset \mathcal{D}_\perp$. In this section, we propose an upper bound of $er^{\mathsf{Aggregate}_{\sigma,\tau}}(\mathcal{D})$ which can be efficiently computed for some families of probability distributions $\mathcal{D}$. Given $D \in \mathcal{D}$ and a subset $\mathcal{H} \subset \mathcal{J}$ of honest experts, we consider the two quantities $\Gamma_{D,\mathcal{H}}(u,v)$ and $\rho_{D,\mathcal{H}}^{\sigma,\tau}$ defined as follows (these quantities are formally defined in Appendix A):

- $\Gamma_{D,\mathcal{H}}(u,v)$ is the probability that (strictly) more than u components of $\boldsymbol{b}$ are (correctly) known[5] by (strictly) less than v honest experts.
- $\rho_{D,\mathcal{H}}^{\sigma,\tau}$ is the probability there is at least one honest expert incompatible with more than τ other honest experts. Note that $\rho_{D,\mathcal{H}}^{\sigma,\tau}$ upper-bounds the probability that at least one honest expert is eliminated.

We then consider the suprema $\Gamma_{\mathcal{D}}(u,v)$, $\rho_{\mathcal{D}}^{\sigma,\tau}$ of these quantities over the choices of $D \in \mathcal{D}$ and $\mathcal{H} \subset \mathcal{J}$ s.t. $|\mathcal{H}| \geq T - \vartheta$.

$$\Gamma_{\mathcal{D}}(u,v) = \max_{D\in\mathcal{D};H\subset\mathcal{J}:|\mathcal{H}|\geq T-\vartheta} \Gamma_{D,\mathcal{H}}(u,v)$$
$$\rho_{\mathcal{D}}^{\sigma,\tau} = \max_{D\in\mathcal{D};H\subset\mathcal{J}:|\mathcal{H}|\geq T-\vartheta} \rho_{D,\mathcal{H}}^{\sigma,\tau}$$

Moreover, the number of incorrect values c_{ji}^* sent to the authority by the corrupted experts which are not eliminated[6] is denoted by $\Omega_{\mathcal{A},D}^{\sigma,\tau}$ (see Appendix A to get a formal definition). The supremum of the expectation of this quantity is denoted by $\Omega_{\mathcal{D}}^{\sigma,\tau}$, i.e.

$$\Omega_{\mathcal{D}}^{\sigma,\tau} = \max_{\mathcal{A}\in P_\vartheta(\mathcal{J})\times\Delta, D\in\mathcal{D}} E(\Omega_{\mathcal{A},D}^{\sigma,\tau})$$

These three suprema can be used to upper-bound $er^{\mathsf{Aggregate}_{\sigma,\tau}}(\mathcal{D})$.

Proposition 1. *We have,*

$$er^{\mathsf{Aggregate}_{\sigma,\tau}}(\mathcal{D}) \leq \rho_{\mathcal{D}}^{\sigma,\tau} + \min_{0\leq u\leq n;0\leq v\leq T}\left(\Gamma_{\mathcal{D}}(u,v) + \frac{u}{n} + \frac{\Omega_{\mathcal{D}}^{\sigma,\tau}}{nv}\right)$$

Proof. (Sketch). See Appendix B for details. Let u, v be integers arbitrarily chosen. Let us consider the event "$E =$ no honest expert is eliminated and there are at less than u components $I = \{i_1,\ldots,i_{t\leq u}\}$ such that $\sum_{j\in\mathcal{H}} c_{ji_k} \leq v$ for any $k = 1,\ldots,t$". E is not satisfied with a probability smaller that than $\rho_{\mathcal{D}}^{\sigma,\tau} + \Gamma_{\mathcal{D}}(u,v)$. In this case, we upper-bound the error rate by 1. Assume now that E is satisfied. In this case, the error can be upper-bounded by $\frac{u}{n} + \frac{\Omega}{nv}$ where Ω is the number of input malicious values: it suffices to assume that the components $b_{i\in I}$ are erroneously predicted and to notice that at most $\frac{\Omega}{v}$ components of $\mathcal{N}\setminus I$ can be erroneously predicted. We conclude by using the fact that $E(\Omega) \leq \Omega_{\mathcal{D}}^{\sigma,\tau}$. □

For some families of probability distributions $\mathcal{D}$, there exists σ, τ such that this upper-bound is small and can be efficiently computed. This is the object of the next section.

VI. Numerical application

Notation. *$\mathcal{B}_{p,m}$ denotes the cdf of the binomial distribution with parameters p, m. Let $X_1,\ldots,X_m$ be m independent random variables belonging to $\{-1,0,1\}$ drawn according to the same probability distribution, i.e. $\mathsf{Pr}(X_i = k) = p_k$. The cdf of the probability distribution of $Z = X_1 + \cdots + X_m$ is denoted by $\mathcal{B}_{p_{-1},p_0,m}$.*

Let $p, p_e \in [0,1]$ such that $p_e < p$. In this section, we consider a very simple family $\mathcal{D}_{p,p_e} \subset \mathcal{D}_\perp$ of probability distributions. Each component of $\boldsymbol{b}$ is known by each (honest) expert with a

[5] We say that a component b_i is correctly known by v honest experts if $b_i \sum_{j\in\mathcal{H}} c_{ji} = v$.

[6] we say that Expert j is eliminated when $w_j = 0$. An eliminated expert does not participate in the majority vote

probability larger than p. Some of them are erroneously known with a probability smaller than p_e.

Definition 2. *$\mathcal{D}_{p,p_e} \subset \mathcal{D}_\perp$ is the family of probability distributions D defined over Δ satisfying the following requirements:*

- *c_{ji} and $c_{j'i'}$ are independent if $(i,j) \neq (i',j')$.*
- *$\mathsf{Pr}(c_{ji} \neq 0) \geq p$*
- *$\mathsf{Pr}(b_i c_{ji} < 0) \leq p_e$*

The quantities $\Gamma_{\mathcal{D}_{p,p_e}}(u,v)$, $\rho^{\sigma,\tau}_{\mathcal{D}_{p,p_e}}$ and $\Omega^{\sigma,\tau}_{\mathcal{D}_{p,p_e}}$ can be easily computed or at least upper-bounded.

Lemma 1. *We have,*

1) $\Omega^{\sigma,\tau}_{\mathcal{D}_{p,p_e}} \leq \vartheta \cdot \max_{i \in U} \left(i\mathcal{B}_{1-\mathcal{B}_{p-p_e,i}(\sigma),T-\vartheta}(\vartheta+\tau)\right)$
2) $\Gamma_{\mathcal{D}_{p,p_e}}(u,v) = 1 - \mathcal{B}_{\mathcal{B}_{p_e,1-p,T-\vartheta}(v-1),n}(u)$
3) $\rho^{\sigma,\tau}_{\mathcal{D}_{p,p_e}} \leq T \cdot \left(1 - \mathcal{B}_{1-\mathcal{B}_{2p_e(p-p_e),n}(\sigma),T-\vartheta-1}(\tau)\right)$

Proof. To prove this lemma, we consider the worst[7] probability distribution $D \in \mathcal{D}_{p,p_e}$ defined by $\mathsf{Pr}(c_{ji} \neq 0) = p$ and $\mathsf{Pr}(b_i c_{ji} < 0) = p_e$. Moreover, we assume that the adversary controls exactly ϑ experts.

1 - The set S_j of correct values c_{ji} received by each honest expert $j \in \mathcal{H}$ is denoted by

$$S_{j\in\mathcal{H}} = \{i \in \mathcal{N} | b_i c_{ji} > 0\}$$

The set M_j of incorrect values input by a corrupted expert $j \in \mathcal{C}$ is denoted by

$$M_{j\in\mathcal{C}} = \{i \in \mathcal{N} | b_i c^*_{ji} < 0\}$$

Let us upper-bound the probability p_j that Expert j is not eliminated, i.e. $w_j = 1$. It suffices that $|S_k \bigcap M_j| > \sigma$ to ensure that an honest Expert k and Expert j are incompatible. For each $i \in M_j$, the probability that $b_i c_{ki} > 0$ is equal to $p - p_e$. Consequently, as c_{ji} and $c_{ki'}$ are independent, the probability that Expert j and Expert k are incompatible is smaller than

$$\rho_{|M_j|} \stackrel{\mathsf{def}}{=} \mathsf{Pr}_D(|S_k \bigcap M_j| > \sigma) = 1 - \mathcal{B}_{p-p_e,|M_j|}(\sigma)$$

and as Expert j is eliminated if it is incompatible with more than $\vartheta + \tau$ honest experts,

$$\mathsf{Pr}_D(w_j = 1) \leq \mathcal{B}_{\rho_{|M_j|},T-\vartheta}(\vartheta+\tau)$$

It follows that the number of incorrect values input by Expert j is upper-bounded, in mean, by

$$|M_j|(\mathcal{B}_{\rho_{|M_j|},T-\vartheta}(\vartheta+\tau)) \leq \max_{i\in U}(iB_{\rho_i,T-\vartheta}(\vartheta+\tau))$$

Thus,

$$\Omega^{\sigma,\tau}_{\mathcal{A},D} \leq \vartheta \cdot \max_{i\in U}(iB_{\rho_i,T-\vartheta}(\vartheta+\tau))$$

[7]It is the probability distribution where the probability that an honest expert inputs correct values is the smallest and the probability that it inputs incorrect values is the largest.

2 - Let $h_i = b_i \sum_{j\in\mathcal{H}} c_{ji}$. As c_{ji} and $c_{j'i'}$ are independent,

$$\mathsf{Pr}(h_i < v) = \mathcal{B}_{p_e,1-p,T-\vartheta}(v-1) = \rho_v$$

and the random variables h_i are independent. Thus, the probability that the number of instances $i \in \mathcal{N}$ satisfying $h_i < v$ is strictly larger than u is equal to

$$\Gamma_{D,\mathcal{H}}(u,v) = 1 - \mathcal{B}_{\rho_v,n}(u)$$

3 - Let j and j' be two honest experts and let $i \in \mathcal{N}$.

$$\mathsf{Pr}_D(c_{ji}c_{j'i} < 0) = 2p_e(p-p_e)$$

It follows that the probability that Expert j is incompatible with Expert j' is equal to

$$\rho = 1 - \mathcal{B}_{2p_e(p-p_e),n}(\sigma)$$

As the values input by honest experts are independent, the probability that Expert j is incompatible with more than τ other honest experts is equal to $1 - \mathcal{B}_{\rho,T-\vartheta-1}(\tau)$ implying that

$$\rho^{\sigma,\tau}_{D,\mathcal{H}} \leq T \cdot (1 - \mathcal{B}_{\rho,T-\vartheta-1}(\tau))$$

□

By injecting these upper-bounds in the inequality of Proposition 1, we get an upper-bound $\mathsf{UB}_{p,p_e,\sigma,\tau}$ of $er^{\mathsf{Aggregate}_{\sigma,\tau}}(\mathcal{D}_{p,p_e})$. Evaluating $\mathsf{UB}_{p,p_e,\sigma,\tau}$ requires computing a "min" over a finite set. To achieve this, we propose a brute force computation by considering all the possible cases. For the parameters used in our experiments, no optimizations are needed[8]. Let us recall that σ, τ are parameters of Aggregate and thus they can be arbitrarily chosen. Let σ^*, τ^* minimizing $\mathsf{UB}_{p,p_e,\sigma,\tau}$, i.e.

$$(\sigma^*,\tau^*) \stackrel{\mathsf{def}}{=} \underset{0\leq\sigma\leq n;0\leq\tau\leq T-2\vartheta}{\mathsf{argmin}} \mathsf{UB}_{p,p_e,\sigma,\tau}$$

Recovering σ^*, τ^* requires computing a "min" over a finite set. To achieve this, we propose a brute force computation by considering all the possible cases.

Results. Computations of $\mathsf{UB}_{p,0,0,0}$ are proposed (see Fig 3) for several values of n, T, ϑ, p. For instance when $n = 1000$, $T = 1000$, $p = 0.1$, $p_e = 0$ and $\vartheta = T/4 = 250$, the error rate $er^{\mathsf{Aggregate}_{0,0}}(\mathcal{D}_{p,0})$ is less than 2% on average against any adversary $\mathcal{A}$. These results could be compared to the naive approach consisting of a simple majority vote. In (almost) all our experiments, $\vartheta \geq p(T-\vartheta)$ ensuring that this naive strategy leads to an error rate larger than 50% in the worst case, i.e. each corrupted experts sends n incorrect values. Indeed, in this case, the number of honest inputs is, in mean, smaller than the number of corrupted inputs making the majority vote fail.

Moreover, for several pairs (n,T), we fixed $\vartheta = T/5$ and we searched[9] p ensuring that $\mathsf{UB}_{p,0,0,0} \approx 1\%$ (see Fig. 4). We observe that p decreases with both n and T. For instance,

[8]In fact, $\Gamma_{D,\mathcal{H}}(u,v)$ converges quickly to 0 when u, v grow. Thus, only "small" values of u and v need considered.

[9]With a dichotomic search.

TABLE I
COMPUTATION OF $\mathsf{UB}_{p,0,0,0}$ FOR SEVERAL VALUES OF n, T, ϑ, p.

$(n,T,p) \setminus \vartheta/T$	0.10	0.20	0.25	0.30	0.35	0.40
$(10^2, 10^3, 0.2)$	0.0%	1.3%	2.1%	4.8%	9.0%	18.8%
$(10^3, 10^3, 0.1)$	0.2%	1.0%	1.8%	3.5%	6.2%	10.0%
$(10^3, 10^2, 0.1)$	0.8%	3.8%	6.3%	9.7%	15.2%	23.0%
$(10^4, 10^2, 0.1)$	0.1%	0.4%	0.7%	1.2%	2.0%	3.4%
$(10^4, 50, 0.1)$	0.9%	1.7%	2.1%	3.0%	3.9%	5.8%
$(10^4, 50, 0.2)$	0.0%	0.1%	0.1%	0.3%	0.4%	0.9%

TABLE II
GIVEN $n, T, \vartheta = T/5$, WE GIVE A VALUE OF p ENSURING THAT $\mathsf{UB}_{p,0,0,0} = 1 \pm 0.1\%$. THIS VALUE WAS OBTAINED BY A DICHOTOMIC SEARCH.

$n \setminus T$	200	400	800	1600	3200	6400
200	27.3%	24.1%	24.1%	24.1%	24.1%	24.1%
400	23.0%	17.5%	14.2%	13.4%	13.3%	13.3%
800	18.6%	12.6%	12.0%	10.9%	9.8%	9.3%
1600	14.2%	9.8%	8.5%	7.7%	7.7%	7.1%
3200	12.3%	7.7%	6.3%	5.5%	5.5%	5.2%
6400	9.8%	6.0%	4.9%	4.4%	3.8%	3.8%
12800	6.0%	4.4%	3.8%	3.2%	3.0%	2.8%

TABLE III
FIX $n = 2000$ AND $T = 500$, WE COMPUTED $\mathsf{UB}_{p,p_e,\sigma^*,\tau^*}$ FOR DIFFERENT VALUES OF ϑ, p, p_e. THE OPTIMAL PARAMETERS τ^*, σ^* ARE GIVEN IN SUBSCRIPT.

$(p,p_e) \setminus \vartheta/T$	0.10	0.20	0.30	0.40
$(5\%, 0.15\%)$	$2.3\%_{100,0}$	$7.0\%_{90,0}$	$18.2\%_{80,0}$	$51.8\%_{50,0}$
$(10\%, 0.30\%)$	$1.2\%_{250,0}$	$3.5\%_{80,1}$	$7.9\%_{70,1}$	$19.2\%_{60,1}$
$(15\%, 0.45\%)$	$0.8\%_{220,1}$	$2.4\%_{200,1}$	$5.3\%_{90,2}$	$12.3\%_{70,2}$
$(20\%, 0.60\%)$	$0.7\%_{240,2}$	$1.9\%_{60,1}$	$4.1\%_{110,3}$	$8.6\%_{200,4}$

$(p,p_e) \setminus \vartheta/T$	0.10	0.20	0.30	0.40
$(5\%, 0.3\%)$	$3.3\%_{160,0}$	$9.9\%_{140,0}$	$25.2\%_{120,0}$	$65.7\%_{30,1}$
$(10\%, 0.6\%)$	$1.9\%_{190,1}$	$5.3\%_{170,1}$	$12.2\%_{70,2}$	$26.7\%_{60,2}$
$(15\%, 0.9\%)$	$1.6\%_{260,2}$	$4.0\%_{150,3}$	$8.4\%_{130,3}$	$17.2\%_{60,4}$
$(20\%, 1.2\%)$	$1.4\%_{270,4}$	$3.3\%_{170,5}$	$6.6\%_{100,6}$	$13.3\%_{50,7}$

$(p,p_e) \setminus \vartheta/T$	0.10	0.20	0.30	0.40
$(5\%, 0.45\%)$	$4.4\%_{200,0}$	$12.9\%_{180,0}$	$32.7\%_{50,1}$	$72\%_{40,1}$
$(10\%, 0.9\%)$	$2.7\%_{270,1}$	$7.1\%_{40,1}$	$15.2\%_{120,1}$	$32.7\%_{50,3}$
$(15\%, 1.35\%)$	$2.1\%_{190,4}$	$5.4\%_{170,4}$	$11.1\%_{90,5}$	$23.0\%_{80,5}$
$(20\%, 1.8\%)$	$1.9\%_{340,5}$	$4.6\%_{180,7}$	$9.1\%_{110,8}$	$17.8\%_{70,9}$

it suffices that each honest expert knows each component of $\boldsymbol{b}$ with a probability larger than 7.7% to allow the authority to recover $\boldsymbol{b}$ with an error smaller (in mean) than 1% when the number of experts is larger than 1600 and the size of $\boldsymbol{b}$ is larger than 1600.

Computations of $\mathsf{UB}_{p,p_e,\sigma^*,\tau^*}$ for different values of ϑ, p, p_e are presented in Fig 5. As expected, we see that σ^*, τ^* grow with p_e and ϑ.

VII. CONCLUSION AND FUTURE WORK

Relevant upper-bounds over $er^{\mathsf{Aggregate}}(\mathcal{D})$ were proposed in this paper. We proposed a numerical application dealing with simple families $\mathcal{D}_{p,p_e}$ of probability distributions ensuring that each component of $\boldsymbol{b}$ is (maybe imperfectly) known by each (honest) expert with a probability larger than p. These families could make sense for some applications. Other families $\mathcal{D}$ could be considered. However, it could be difficult to give an analytic upper-bound of $er^{\mathsf{Aggregate}}(\mathcal{D})$. Nevertheless, approximations can be obtained by sampling worst-case probability distributions $D \in \mathcal{D}$.

Moreover, we are convinced that Aggregate and its analysis (under the specific conditions of this paper) can be improved. Furthermore, the target tuple $\boldsymbol{b}$ is binary but natural extensions to numerical tuples $\boldsymbol{b} \in \mathbb{R}^n$ could be provided by applying threshold ε, i.e. c_{ji} and $c_{j'i}$ are said to be equal if $|c_{ji} - c_{j'i}| < \varepsilon$. The new parameter ε should be carefully chosen in order that most of honest experts input the same values c_{ji} within ε. In [4], Dubois et al. propose that each expert sends an interval I_i containing the true value b_i. Investigations should be done to adapt our work to this setting.

ACKNOWLEDGMENT

The authors would like to thank the BAG members for their helpful discussions always around a coffee.

REFERENCES

[1] R. Cramer, I. Damgård, and J. B. Nielsen. Multiparty computation from threshold homomorphic encryption. In *EUROCRYPT*, pages 280–299, 2001.

[2] Xin Luna Dong, Laure Berti-Equille, and Divesh Srivastava. Integrating conflicting data: The role of source dependence. *Proc. VLDB Endow.*, 2(1):550–561, August 2009.

[3] Xin Luna Dong and Felix Naumann. Data fusion - resolving data conflicts for integration. *PVLDB*, 2(2):1654–1655, 2009.

[4] Didier Dubois and Henri Prade. Possibility theory and data fusion in poorly informed environments. *Control Engineering Practice*, 2(5):811–823, 1994.

[5] Hugh Durrant-Whyte and Thomas C Henderson. Multisensor data fusion. *Springer handbook of robotics*, pages 585–610, 2008.

[6] Terrence L. Fine. Review: Glenn shafer, a mathematical theory of evidence. *Bulletin of the American Mathematical Society*, 83(4):667–672, 07 1977.

[7] Bahador Khaleghi, Alaa Khamis, Fakhreddine O. Karray, and Saiedeh N. Razavi. Multisensor data fusion: A review of the state-of-the-art. *Inf. Fusion*, 14(1):28–44, January 2013.

[8] CV Negoita, LA Zadeh, and HJ Zimmermann. Fuzzy sets as a basis for a theory of possibility. *Fuzzy sets and systems*, 1:3–28, 1978.

[9] O.Goldreich, S.Michali, and A.Wigderson. How to play any mental game or a completeness theorem for protocols with honest majority. In *STOC*, pages 218–229, 1987.

[10] Zdzislaw Pawlak. *Rough Sets: Theoretical Aspects of Reasoning About Data*. Kluwer Academic Publishers, Norwell, MA, USA, 1992.

[11] F.K.J. Sheridan. A survey of techniques for inference under uncertainty. *Artificial Intelligence Review*, 5(1-2):89–119, 1991.

[12] Lotfi A Zadeh. Fuzzy sets. *Information and control*, 8(3):338–353, 1965.

APPENDIX A
FORMAL DEFINITION OF QUANTITIES CONSIDERED IN SECTION 5

- Let u, v be positive integers s.t. $u \leq n$ and $v \leq T$. Let $I_v = \{i \in \mathcal{N} | b_i \sum_{j \in \mathcal{H}} c_{ji} < v\}$. $\Gamma_{D,\mathcal{H}}(u,v)$ denotes the probability under D that the cardinality of I_v is strictly larger than u, i.e.

$$\Gamma_{D,\mathcal{H}}(u,v) = \mathsf{Pr}(|I_v| > u)$$

- Let $I(j,j') := \{i \in \mathcal{N} - c^*_{ji} c^*_{j'i} < 0\}$ and let $\alpha_{D,\mathcal{H}}(j) = \{j' \in \mathcal{H} || I(j,j')| > \sigma\}$.

$$\rho^{\sigma,\tau}_{D,\mathcal{H}} = \mathsf{Pr}(\exists j \in \mathcal{H}, |\alpha_{D,\mathcal{H}}(j)| > \tau)$$

- $\Omega^{\sigma,\tau}_{\mathcal{A},D} = \sum_{j \in \mathcal{C}} w_j |\{i \in \mathcal{N} | c^*_{ji} b_i < 0\}|$

APPENDIX B
PROOF OF PROPOSITION 1

Proof. According to notation of Section 3, $\boldsymbol{o}$ denotes the random tuple output by $\mathsf{Aggregate}_{\sigma,\tau}$. For the sake of simplicity, $er^{\mathsf{Aggregate}_{\sigma,\tau}}_{\mathcal{A}}(D)$ will be denoted by er.

The event G refers to the fact that no honest expert is eliminated[10], i.e. $w_j = 1$ for any $j \in \mathcal{H}$. By definition, $\mathsf{Pr}(\overline{G}) \leq \rho^{\sigma,\tau}_{\mathcal{D}}$.

Let u, v be arbitrary positive integers s.t. $u \leq n$ and $v \leq T$. Let $I_v = \{i \in \mathcal{N} | b_i \sum_{j \in \mathcal{H}} c_{ji} \leq v\}$. The event $|I_v| \geq u$ will be denoted by F. By definition,

$$\mathsf{Pr}(F) = \Gamma_{D,\mathcal{H}}(u,v) \leq \Gamma_{\mathcal{D}}(u,v)$$

In the following of the proof, er' denotes the expectation of the error of $\boldsymbol{o}$ assuming that $G, \overline{F}$ are realized, i.e.,

$$er' = \frac{1}{2n} E(\|\boldsymbol{b} - \boldsymbol{o}\|_1 | \ G, \overline{F})$$

Clearly, $er \leq \mathsf{Pr}(\overline{G}) + \mathsf{Pr}(F,G) + \mathsf{Pr}(\overline{F},G) er'$ implying that

$$er \leq \rho^{\sigma,\tau}_{\mathcal{D}} + \Gamma_{\mathcal{D}}(u,v) + \mathsf{Pr}(\overline{F},G) er'$$

Let us focus on er' by assuming that the events $G, \overline{F}$ are realized. According to the definition of $\overline{F}$, the cardinality of I_v is smaller than u implying that

$$|\{i \in I_v | o_i \neq b_i\}| \leq u$$

By definition, for each $i \in \overline{I_v}$, $b_i \sum_{j \in \mathcal{H}} c_{ji} \geq v$. It follows that

$$|\{i \in \overline{I_v} | o_i \neq b_i\}| \leq \frac{\Omega^{\sigma,\tau}_{\mathcal{A},D}}{v}$$

Consequently, the error of $\boldsymbol{o}$ is smaller than $\frac{1}{n}(u + \frac{\Omega^{\sigma,\tau}_{\mathcal{A},D}}{v})$, implying that

$$er' \leq \frac{u}{n} + \frac{1}{nv} E(\Omega^{\sigma,\tau}_{\mathcal{A},D} | G, \overline{F})$$

[10]The fact that G is realized means that all the values input by the honest experts are considered by $\mathsf{Aggregate}_{\sigma,\tau}$.

As $\Omega^{\sigma,\tau}_{\mathcal{A},D}$ is a positive random variable,

$$E(\Omega^{\sigma,\tau}_{\mathcal{A},D} | G, \overline{F}) \leq \frac{E(\Omega^{\sigma,\tau}_{\mathcal{A},D})}{\mathsf{Pr}(G, \overline{F})} \leq \frac{\Omega^{\sigma,\tau}_{\mathcal{D}}}{\mathsf{Pr}(G, \overline{F})}$$

It follows that

$$\begin{aligned} er &\leq \rho^{\sigma,\tau}_{\mathcal{D}} + \Gamma_{\mathcal{D}}(u,v) + \mathsf{Pr}(G, \overline{F}) \left(\frac{u}{n} + \frac{\Omega^{\sigma,\tau}_{\mathcal{D}}}{nv \mathsf{Pr}(G, \overline{F})} \right) \\ &\leq \rho^{\sigma,\tau}_{\mathcal{D}} + \Gamma_{\mathcal{D}}(u,v) + \frac{u}{n} + \frac{\Omega^{\sigma,\tau}_{\mathcal{D}}}{nv} \end{aligned}$$

As u, v were arbitrarily chosen, $er \leq \rho^{\sigma,\tau}_{\mathcal{D}} + \min_{0 \leq u \leq n; 0 \leq v \leq T} \left(\Gamma_{\mathcal{D}}(u,v) + \frac{u}{n} + \frac{\Omega^{\sigma,\tau}_{\mathcal{D}}}{nv} \right)$. This concludes the proof.

□