Belief Merging Operators as Maximum Likelihood Estimators

Patricia Everaere[1] , **Sébastien Konieczny**[2] and **Pierre Marquis**[2,3]
[1]CRIStAL – CNRS, Université de Lille, France
[2]CRIL – CNRS, Université d'Artois, France
[3]Institut Universitaire de France, France
patricia.everaere-caillier@univ-lille.fr, konieczny@cril.fr, marquis@cril.univ-artois.fr

Abstract

We study how belief merging operators can be considered as maximum likelihood estimators, i.e., we assume that there exists a (unknown) true state of the world and that each agent participating in the merging process receives a noisy signal of it, characterized by a noise model. The objective is then to aggregate the agents' belief bases to make the best possible guess about the true state of the world. In this paper, some logical connections between the rationality postulates for belief merging (IC postulates) and simple conditions over the noise model under consideration are exhibited. These results provide a new justification for IC merging postulates. We also provide results for two specific natural noise models: the world swap noise and the atom swap noise, by identifying distance-based merging operators that are maximum likelihood estimators for these two noise models.

1 Introduction

The aggregation of pieces of information, encoding preferences, beliefs or judgments, that are owned by a number of agents, is a key challenge in many settings, among social choice, vote, merging, or judgment aggregation. Whatever the setting, the design of aggregation operators is typically guided either by axiomatic concerns, or by epistemic concerns. In the former case, expected properties (postulates) are exhibited, and operators satisfying them are pointed out or impossibility theorems are proved. In the latter case, it is assumed that the preferences, beliefs or judgments that are reported by the agents are a noisy perception of "correct" preferences, beliefs or judgments. Such an epistemic approach have received great attention when dealing with preferences [Young, 2003; Conitzer and Sandholm, 2005; Xia *et al.*, 2010; Conitzer, 2014; Elkind and Slinko, 2016; Pivato, 2012], and can be traced back at least to Condorcet, who justified the use of majoritarian vote in this way (his famous jury theorem [Condorcet, 1785]). Contrastingly, the epistemic interpretation of belief merging operators gave rise to few research work (we know a single paper about it, it is discussed in the related work section).

Propositional belief merging operators aim to associate a belief base with a set of (usually conflicting) belief bases representing the individual beliefs of the agents. So far, belief merging operators have been evaluated with respect to several criteria, such as rationality properties leading to the so-called IC merging operators [Konieczny and Pino Pérez, 2002a; Konieczny and Pino Pérez, 2011], properties inspired by vote [Haret *et al.*, 2016], computational complexity [Konieczny *et al.*, 2004], strategy-proofness [Everaere *et al.*, 2007], etc. In this work, we want to evaluate belief merging operators as *maximum likelihood estimators (MLE)*. One assumes that there exists a true state of the world $\omega^\star$ (represented by a specific propositional interpretation) and that each agent participating in the merging process receives a noisy signal of it (a formula K_i). The objective is then to aggregate the agents' belief bases to make the best possible guess about $\omega^\star$. The key issue we want to address is to determine the extent to which belief merging operators are suited to this objective, i.e., to figure out how the corresponding merged bases and $\omega^\star$ are connected.

In the following, we provide a number of results in this direction. First, we present some logical connections between rationality postulates for merging (IC postulates) and simple conditions (especially about independence) over the noise model P under consideration. We also rationalize a large family of distance-based IC merging operators Δ by showing that for each operator Δ of the family, there exists a noise model P such that Δ is a MLE for P. Then we focus on two specific noise models. For the first one, the world swap noise, we show that a standard distance-based merging operator based on the drastic distance is a MLE. And for the second one, the atom swap noise, we show that the distance-based merging operator based on the sum and the Hamming distance is a MLE. Finally, we report some empirical results showing that those two distance-based merging operators are efficient noise cancellers in practice, meaning that the number of agents needed for identifying the true state of the world with high probability remains small enough.

The rest of the paper is organized as follows. In the next section, some formal preliminaries on belief merging are provided. In Section 3, noise models and maximum likelihood estimators are defined in the belief merging setting. In Section 4, conditions under which merging operators are / are not maximum likelihood estimators are pointed out. In Section 5 the focus is laid on the world swap noise, and in Section 6 on the atom swap noise. Section 7 gives some empirical results. Section 8 discusses some related work. Finally, Section 9 concludes the paper. For space reasons only the proofs of some of the main results are provided. A full-proof version of the paper is available on http://www.cril.fr/∼konieczny/ijcai20proofs.pdf.

2 Formal Preliminaries

We consider a propositional language $\mathcal{L}$ defined from a finite set $\mathcal{P}$ of m propositional variables and the usual connectives. An interpretation (or state of the world) ω is a total function from $\mathcal{P}$ to $\{0,1\}$. Ω is the set of all interpretations. An interpretation ω is a model of a formula $\phi \in \mathcal{L}$ if and only if it makes it true in the usual truth-functional way. Otherwise, ω is a counter-model of ϕ. $[\phi]$ denotes the set of models of formula ϕ, i.e., $[\phi] = \{\omega \in \Omega \mid \omega \models \phi\}$.

A *belief base* K is a finite set of propositional formulae, interpreted conjunctively (i.e., viewed as the conjunction of its elements). A *profile* E (of belief bases) is associated with a group of n agents that are involved in the merging process; formally, E is given by a vector of bases $E = \langle K_1, \dots, K_n\rangle$. $\bigwedge E$ denotes the conjunction of all elements of E, and $\sqcup$ denotes concatenation. Two profiles $E = \langle K_1, \dots, K_n\rangle$ and $E' = \langle K'_1, \dots, K'_n\rangle$ are said to be equivalent, noted $E \equiv E'$, iff there exist a permutation σ such that $\forall K_i$, we have $[K_i] = [K'_{\sigma(i)}]$. An *integrity constraint* μ is a formula restricting the possible results of the merging process. A *merging operator* $\triangle$ is a mapping which associates with a profile E and an integrity constraint μ a (*merged*) base $\triangle_\mu(E)$.

We expect belief merging operators to satisfy the following rationality postulates (the IC postulates):

Definition 1 *A merging operator $\triangle$ is an* IC merging operator *iff it satisfies the following properties:*

(IC0) $\triangle_\mu(E) \models \mu$

(IC1) *If μ is consistent, then $\triangle_\mu(E)$ is consistent*

(IC2) *If $\bigwedge E$ is consistent with μ, then $\triangle_\mu(E) \equiv \bigwedge E \wedge \mu$*

(IC3) *If $E_1 \equiv E_2$ and $\mu_1 \equiv \mu_2$, then $\triangle_{\mu_1}(E_1) \equiv \triangle_{\mu_2}(E_2)$*

(IC4) *If $K_1 \models \mu$ and $K_2 \models \mu$, then $\triangle_\mu(\langle K_1, K_2\rangle) \wedge K_1$ is consistent if and only if $\triangle_\mu(\langle K_1, K_2\rangle) \wedge K_2$ is consistent*

(IC5) $\triangle_\mu(E_1) \wedge \triangle_\mu(E_2) \models \triangle_\mu(E_1 \sqcup E_2)$

(IC6) *If $\triangle_\mu(E_1) \wedge \triangle_\mu(E_2)$ is consistent, then $\triangle_\mu(E_1 \sqcup E_2) \models \triangle_\mu(E_1) \wedge \triangle_\mu(E_2)$*

(IC7) $\triangle_{\mu_1}(E) \wedge \mu_2 \models \triangle_{\mu_1 \wedge \mu_2}(E)$

(IC8) *If $\triangle_{\mu_1}(E) \wedge \mu_2$ is consistent, then $\triangle_{\mu_1 \wedge \mu_2}(E) \models \triangle_{\mu_1}(E) \wedge \mu_2$*

See [Konieczny and Pino Pérez, 2002a] for explanations and justifications of these properties.

Let us give some examples of IC merging operators from the family of distance-based merging operators [Konieczny *et al.*, 2004]. We first need a notion of (pseudo-)distance:

Definition 2 *A* (pseudo-)distance *between interpretations is a mapping $d : \Omega \times \Omega \to \mathbb{R}^+$ such that for any ω_1, $\omega_2 \in \Omega$:*

- $d(\omega_1, \omega_2) = d(\omega_2, \omega_1)$
- $d(\omega_1, \omega_2) = 0$ *iff* $\omega_1 = \omega_2$

Every distance between interpretations can be easily lifted to a "distance" between worlds and formulae, by stating that $d(\omega, K) = min_{\omega' \in [K]} d(\omega, \omega')$.

Usual distances considered in merging [Konieczny and Pino Pérez, 2002a] are the Hamming distance d_H and the drastic distance d_D. $d_H(\omega_1, \omega_2)$ is the number of propositional letters on which the two interpretations differ. d_D is such that $d_D(\omega_1, \omega_2) = 0$ iff $\omega_1 = \omega_2$, and $= 1$ otherwise.

One then needs a notion of aggregation function:

Definition 3 *An* aggregation function *is a mapping*[1] f *from* $(\mathbb{R}^+)^p$ *to* $\mathbb{R}^+$*, which satisfies:*

- *if* $x_i \geq x'_i$, *then* $f(x_1, \dots, x_i, \dots, x_p) \geq f(x_1, \dots, x'_i, \dots, x_p)$
- *if* $\forall i \in \{1, \dots, p\}$ $x_i = 0$ *then* $f(x_1, \dots, x_p) = 0$
- $f(x) = x$
- *If σ is a permutation over $\{1, \dots, p\}$, then* $f(x_1, \dots, x_p) = f(x_{\sigma(1)}, \dots, x_{\sigma(p)})$

Many aggregation functions have been pointed out so far, and considered in the belief merging literature. They include sum (Σ), $leximin$ (alias $Gmin$), $leximax$ (alias $Gmax$), and Σ^k (the sum of the k^{th} powers – where k is a fixed integer). Formal definitions can be found, e.g., in [Konieczny and Pino Pérez, 2002a; Everaere *et al.*, 2010b].

Definition 4 *Let d and f be respectively a pseudo-distance between interpretations and an aggregation function. The* distance-based merging operator $\triangle^{d,f}$ *is defined by* $[\triangle^{d,f}_\mu(E)] = \min([\mu], \leq_E)$*, where the total pre-order $\leq_E$ on Ω is defined in the following way (with $E = \langle K_1, \dots, K_n\rangle$):*

- $\omega \leq_E \omega'$ *iff* $d(\omega, E) \leq d(\omega', E)$,
- $d(\omega, E) = f(d(\omega, K_1), \dots, d(\omega, K_n))$.

For many standard aggregation functions, whatever the chosen pseudo-distance d between interpretations, the corresponding distance-based operator $\triangle^{d,f}$ is an IC merging operator [Konieczny and Pino Pérez, 2002a; Everaere *et al.*, 2010b]:

Proposition 1 *Let d be any pseudo-distance between interpretations and let f be an aggregation function among Σ, $leximin$, $leximax$, or Σ^k. $\triangle^{d,f}$ is an IC merging operator.*

More generally, properties on the aggregation function ensuring that the corresponding distance-based operator is an IC merging operator can be found in [Konieczny *et al.*, 2004].

3 Noise and Maximum Likelihood Estimators

In this work, we assume the existence of a true state of the world, that is represented by an interpretation $\omega^\star \in \Omega$. Whenever some integrity constraints μ are available, we assume that $\omega^\star$ is a model of μ, meaning that the integrity constraints are always supposed to hold in the true state of the world. Stated otherwise, integrity constraints are pieces of *knowledge*, i.e., of true beliefs.

[1] Strictly speaking, it is a family of mappings, where a mapping is defined for each positive integer p.

Let $\mathcal{I}$ be a set of n agents. We suppose that each agent has a noisy perception of the true state of the world $\omega^\star$: for every $i \in \mathcal{I}$, the base K_i representing the beliefs of agent i corresponds to the perception of $\omega^\star$ by i. The given (observed) profile $E = \langle K_1, \dots, K_n \rangle$ thus reflects the noisy perception of the true state of the world by the group of agents.

Because of the noise, there is no guarantee that $[K_i] = \{\omega^\star\}$ for at least one agent i, or even that there exists any *logical* connection between E and $\omega^\star$ (for instance, it can be the case that $\omega^\star$ is inconsistent with each K_i in E). However, depending on $\omega^\star$, the observation of a profile E is more or less plausible given the noise. Formally, a noise model makes precise the connection between $\omega^\star$ and the profiles using probabilities.

Given the worlds $\omega \in \Omega$ and the agents $i \in \mathcal{I}$, one considers $2^m \times n$ Boolean random variables of the form (ω, i). Each random variable (ω, i) is thus associated with two possible outcomes, the one noted (ω, i), where the random variable (ω, i) takes the truth value true, and the one noted $\overline{(\omega, i)}$, where the random variable (ω, i) takes the truth value false. The fact that (ω, i) takes the truth value true (resp. false) means that the world ω is not discarded (resp. discarded) by the observation achieved by agent i, i.e., ω is a model of K_i (resp. not a model of K_i). O denotes the Cartesian product of the set of outcomes for all the Boolean random variables (ω, i). Every element of O represents (from a semantical point of view) a profile E since it indicates for every world $\omega \in \Omega$ and every agent $i \in \mathcal{I}$ whether or not ω is a model of K_i. Similarly, every projection of an element of O over the outcomes associated with all the random variables of the form (ω, i) for a fixed $i \in \mathcal{I}$ represents (again, up to logical equivalence) a belief base K_i. Given those notations, a noise model can be formally defined as follows:

Definition 5 A noise model P *is a mapping associating with every consistent formula $\mu \in \mathcal{L}$ and every interpretation $\omega \in [\mu]$ a joint probability distribution, noted $P_{\mu,\omega}$, of the $2^m \times n$ Boolean random variables (ω, i), i.e., $P_{\mu,\omega}$ is a probability distribution over O, s.t. the following conditions are satisfied:*

- $\forall \omega' \in \Omega$, *if* $\omega' \not\models \mu$, *then* $\forall i \in \mathcal{I}$, $P_{\mu,\omega}((\omega', i)) = 0$,
- *If* $\mu' \equiv \mu$, *then* $P_{\mu,\omega} = P_{\mu',\omega}$.

Since each $P_{\mu,\omega}$ is a joint probability distribution, every (non-empty) subset of the $2^m \times n$ Boolean random variables of the form (ω, i) corresponds to a marginal probability distribution.

The first condition in Definition 5 states that any interpretation ω' that violates the given integrity constraint μ will never be reported by any agent as a model of her belief base (stated otherwise, every agent involved in the merging process is supposed to be aware of μ). Thus, the (marginal) probability of such an (impossible) event (ω', i) for the random variable (ω', i) is 0. The second condition in the definition expresses a form of syntax-independence for the integrity constraint.

A profile E is said to be generated by the noise model P if E is a random variate (outcome) from P. In this work, we observe a single profile E generated by P, and we want to determine whether a merging operator allows to identify the worlds that best explain E, i.e., to find (if it exists) a merging operator that is a maximum likelihood operator for P:

Definition 6 *Let P be a noise model. A merging operator Δ^P is a* maximum likelihood estimator (MLE) for P *iff for any consistent formula $\mu \in \mathcal{L}$ and for any profile E generated by P, we have*[2] $[\Delta^P_\mu(E)] = argmax_\omega P_{\mu,\omega}(E)$.

4 IC Postulates, IC Operators and MLEs

We first show that if a merging operator Δ is a MLE for some noise P (whatever this noise), then Δ must comply with some IC postulates.

Proposition 2 *Let P be any noise model. If Δ is a MLE for P, then Δ satisfies* **(IC0)***,* **(IC1)***,* **(IC3)***,* **(IC7)** *and* **(IC8)**.

In general, we can expect noise models to satisfy some additional constraints, like the following independence condition, which is very natural (and is quite standard for voting methods):

Definition 7 *A noise model P satisfies the* independence condition *iff for any for any consistent formula $\mu \in \mathcal{L}$, any $\omega \in [\mu]$, for each $i \in \mathcal{I}$, the set of 2^m random variables $\{(\omega, i) \mid \omega \in \Omega\}$ considered in $P_{\mu,\omega}$ is independent from all the remaining random variables.*

When P satisfies the independence condition, whatever the true state of the world, the beliefs of any agent are independent from the beliefs of the other agents.

The independence condition rules out some merging operators as MLEs:

Proposition 3 *Let P be any noise model satisfying the independence condition. If Δ is a MLE for P, then Δ satisfies* **(IC5)** *and* **(IC6)**.

To sum up, a merging operator that is a MLE for a noise model satisfying the independence condition must satisfy almost all IC postulates. The exceptions are **(IC2)** (that requires the merged base to be the conjunction of all the bases with the integrity constraints whenever this conjunction is consistent), and **(IC4)** (that requires an equal treatment of the two bases in any profile containing two bases).

In the literature (see for example [Conitzer and Sandholm, 2005]), another assumption about noise is often considered, namely the uniformity condition:

Definition 8 *A noise model P satisfies the* uniformity condition *iff for any for any consistent formula $\mu \in \mathcal{L}$, any $\omega \in [\mu]$, $\omega' \in \Omega$, any $i, j \in \mathcal{I}$, we have $P_{\mu,\omega}((\omega', i)) = P_{\mu,\omega}((\omega', j))$.*

Intuitively, the uniformity condition states that the noise model is the same one for all the agents.

Interestingly, many merging operators (especially IC ones) can be interpreted as MLEs for a noise model satisfying the independence condition and the uniformity condition:

Proposition 4 *Let d be any pseudo-distance on interpretations. Let k be any natural number. There exists a noise model P satisfying the independence condition and the uniformity condition such that the merging operators Δ^{d,Σ^k} and $\Delta^{d,leximax}$ are MLEs for P.*

[2] $argmax_\omega P_{\mu,\omega}(E) = \{\omega \in \Omega | \forall \omega' \in \Omega, P_{\mu,\omega}(E) \geq P_{\mu,\omega'}(E)\}$.

Proof: Let $\omega^\star \models \mu$ be the true state of the world and $E = \langle K_1, \ldots, K_n \rangle$ a given profile. Let us define the noise P such that the probability that any world ω is a model of the base associated with any agent j, namely $P_{\mu,\omega^\star}((\omega, j))$, is equal to $\alpha \times \frac{1}{2^{(d(\omega,K_j)+1)^k}}$, where α is a positive real number that does not depend on j (the uniformity assumption ensures that it exists) and k is a fixed integer.

Due to the independence condition on P, the probability of ω to be a model of every base of E is thus equal to

$$\Pi_{j=1}^{n} \frac{\alpha}{2^{(d(\omega,K_j)+1)^k}} = \frac{\alpha^n}{2^{\Sigma_{j=1}^{n}((d(\omega,K_j)+1)^k)}}.$$

Then the problem of determining the maximum likelihood estimates of $\omega^\star$ consists in exhibiting the interpretations $\omega \models \mu$ such that $\frac{\alpha^n}{2^{\Sigma_{j=1}^{n}((d(\omega,K_j)+1)^k)}}$ is maximal. Obviously, $\frac{\alpha^n}{2^{\Sigma_{j=1}^{n}((d(\omega,K_j)+1)^k)}}$ is maximal if and only if $\Sigma_{j=1}^{n}((d(\omega, K_j) + 1)^k)$ is minimal, and $\Sigma_{j=1}^{n}((d(\omega, K_j) + 1)^k)$ is minimal if and only if the weighted sum of Hamming distances of ω to every base from the profile $E = \langle K_1, \ldots, K_n \rangle$ is minimal. This shows that every Δ^{d,Σ^k} is a MLE for P.

Finally, it is well-known [Konieczny and Pino Pérez, 2002b] that for a sufficiently large k whatever the pseudo-distance d over interpretations and the profile $E = \langle K_1, \ldots, K_n \rangle$, two worlds ω and ω' are such that $\sum_{i=1}^{n} d(\omega, K_i)^k \leq \sum_{i=1}^{n} d(\omega', K_i)^k$ if and only if the *leximax* aggregation of the values $d(\omega, K_1), \ldots, d(\omega, K_n)$ is less than or equal to the *leximax* aggregation of the values $d(\omega', K_1), \ldots, d(\omega', K_n)$ w.r.t. the lexicographic ordering. Hence an interpretation ω minimizing the *leximax* aggregation of the values $d(\omega, K_1), \ldots, d(\omega, K_n)$ w.r.t. the lexicographic ordering maximizes the probability $\Pi_{j=1}^{n} \frac{\alpha}{2^{(d(\omega,K_j)+1)^k}}$ when k is sufficiently large. Accordingly, the merging operator $\Delta^{d,leximax}$ is also a MLE for P. □

Proposition 4 shows that many distance-based merging operators Δ (namely those for which the aggregation function is one of the Σ^k or *leximax*) are rationalizable in the sense that there exists a noise model P such that Δ is a MLE for P.

The fact that a merging operator can be rationalized as a MLE (i.e., showing that there exists a noise model P such that the operator is a MLE for P) is a valuable property of the operator when one is interested in making the best possible guess about $\omega^\star$. Indeed, as already discussed in [Everaere *et al.*, 2010a], merging a profile of belief bases is a useful operation when one wants to synthesize the information given in the profile, i.e., to characterize a belief base which best represents the beliefs of the input profile (the synthesis view of belief merging), or when one wants to best approximate the true state of the world from the information give in the profile (the epistemic view of belief merging). In the second case, rationalizable merging operators must clearly be preferred to non-rationalizable ones. The rationalization property can thus be viewed as an interesting criterion for helping to choose a merging operator that is suited to the epistemic purpose.

In some situations, more information about the noise model P are available, and an interesting issue is then to determine whether a merging operator that is a MLE for the specific noise model P exists. This is a research question that is somewhat converse to the one discussed in the previous paragraphs: instead of starting with a merging operator Δ and looking for the existence of a noise model P such that Δ is a MLE for P, one starts with a noise model P and looks for a merging operator Δ that is a MLE for P.We address this issue in the rest of the paper. To be more precise, two sensible noise models are considered and merging operators that are suited to them are identified.

5 The World Swap Noise

A plausible noise model is obtained by assuming that some interpretations ω of Ω are misperceived (and thus misclassified) by the agents: for an agent i, when $\omega = \omega^\star$ (resp. $\omega \neq \omega^\star$), it is considered as a counter-model (resp. a model) of K_i. This noise model is parameterized by a *(world) swap probability* p, which represents the chances of an interpretation to be considered as a model of the beliefs of an agent while it is not the true state of the world. We make the (reasonable) assumption that this probability of error is less than $\frac{1}{2}$, otherwise the noise would be too important. More formally, the *world swap noise model* P^{wsn} is a mapping associating with every consistent formula $\mu \in \mathcal{L}$, every interpretation $\omega \in [\mu]$, the joint probability distribution, noted $P^{wsn}_{\mu,\omega}$, of the $2^m \times n$ Boolean random variables (ω', i) defined as the product of the distributions of the random variables (ω', i) when ω' varies in Ω and i varies in $\mathcal{I}$:

Definition 9 *Let $p \in [0, \frac{1}{2})$, the (world) swap probability is:*

$$P^{wsn}_{\mu,\omega}((\omega', i)) = \begin{cases} 0 & \textit{if } \omega' \not\models \mu \\ p & \textit{if } \omega' \models \mu \textit{ and } \omega' \neq \omega, \\ 1-p & \textit{otherwise.} \end{cases}$$

Stated otherwise, one considers here that every Boolean random variable (ω, i) is associated with a Bernoulli distribution, so that the joint distribution $P^{wsn}_{\mu,\omega}$ for the $2^m \times n$ random variables is a binomial distribution with parameters n and p.

We can easily check that the world swap noise model P^{wsn} satisfies the independence and uniformity conditions. It turns out that this noise is linked to a pseudo-distance between formulae. The distance of two formulae is defined here as the number of interpretations one have to add or to retrieve to one of the two formulae to make it equivalent to the other formula. We call it the *swap distance* d_{swap} between formulae.

Definition 10 (swap distance) *Let ϕ, ϕ' be two formulae from $\mathcal{L}$. $d_{swap}(\phi, \phi') = |\{\omega \in \Omega \mid \omega \models \phi \oplus \phi'\}|$.*[3]

Example 1 *Consider the formulae $\phi = a \wedge b$, $\phi' = a \vee b$, $\phi'' = \neg a$ where $\mathcal{P} = \{a, b\}$. We have: $d_{swap}(\phi, \phi') = 2$, $d_{swap}(\phi, \phi'') = 3$, and $d_{swap}(\phi', \phi'') = 3$. According to d_{swap}, $a \wedge b$ is thus closer to $a \vee b$ than to $\neg a$.*

Proposition 5 *d_{swap} is a pseudo-distance between formulae.*

[3] $\oplus$ is the exclusive disjunction connective.

When the world swap noise model P^{wsn} is considered, for any model ω of the integrity constraint μ and any base K of E, the largest $d_{swap}(K, \omega)$, the less $P^{wsn}_{\mu,\omega}(K)$. As a consequence:

Proposition 6 *Let ω be any model of the integrity constraint μ, ϕ be any formula having ω as its single model, and let K_i, K_j be two belief bases of E. If $d_{swap}(K_i, \phi) \geq d_{swap}(K_j, \phi)$, then $P^{wsn}_{\mu,\omega}(K_i) \leq P^{wsn}_{\mu,\omega}(K_j)$.*

On this ground, we can show that the drastic merging operator $\Delta^{d_D,\Sigma}$ is a MLE for this noise:

Proposition 7 $\Delta^{d_D,\Sigma}$ *is a MLE for* P^{wsn}.

6 The Atom Swap Noise

Another interesting noise is the atom swap noise P^v, obtained by considering that the observed values of variables from $\mathcal{P}$ may differ from their values in $\omega^\star$. In such a case, the probability that an interpretation belongs to a base decreases with the Hamming distance between this interpretation and $\omega^\star$. Such a noise model is suited for instance to scenarios where the observed values of the variables come from noisy sensors.

P^v is parameterized by a probability $p < 1$, that indicates the probability that the truth value of any observed variable $v \in \mathcal{P}$ differs from the actual one (i.e., its value in $\omega^\star$). Formally the *atom noise model* P^v is a mapping associating with every consistent formula $\mu \in \mathcal{L}$, every interpretation $\omega \in [\mu]$, a joint probability distribution, noted $P^v_{\mu,\omega}$, of the $2^m \times n$ Boolean random variables (ω', i) defined as the product of the distributions of the random variables (ω', i) when ω' varies in Ω and i varies in $\mathcal{I}$:

Definition 11 *Let $p \in (0, 1)$. The (atom) swap probability is*

$$P^v_{\mu,\omega}((\omega', i)) = \begin{cases} 0 & \text{if } \omega' \not\models \mu \\ p^{d_H(\omega,\omega')+1} & \text{otherwise.} \end{cases}$$

So, with this noise model, the true state of the world $\omega^\star$ appears in any belief base with a probability p, and any other world has a smaller probability, which exact value depends on its distance to $\omega^\star$.

Clearly enough, P^v satisfies the independence condition (by definition) and the uniformity condition (since $P^v_{\mu,\omega}((\omega', i))$ does not depend on i).

Let us illustrate the behavior of this noise model on an example, with three variables ($m = 3$) and a profile for $n = 5$ agents.

Example 2 *Suppose that the true state of the world $\omega^\star$ is $\omega^\star = 001$, with a swap noise P^v such that $p = 30\%$. For 5 agents, a possible outcome given that noise is the profile $E = \langle K_1, K_2, K_3, K_4, K_5 \rangle$ with $[K_1] = \{011\}$, $[K_2] = \{110, 001, 111, 101\}$, $[K_3] = \{000, 001, 111\}$, $[K_4] = \{000, 110, 001\}$, $[K_5] = \{001, 111\}$. In this case, the resulting merged base with $\Delta^{d_H,\Sigma}$ from this profile is $[\Delta^{d_H,\Sigma}(E)] = \{001\}$.*

Proposition 8 $\Delta^{d_H,\Sigma}$ *is a MLE for* P^v.

Proof: Let w be any interpretation. Let $d_{H,\Sigma}(\mathsf{w}, E)$ denote $d_H(\mathsf{w}, K_1) + \ldots + d_H(\mathsf{w}, K_n)$, where $d_H(\mathsf{w}, K_i) = \min_{\omega \in K_i} d_H(\mathsf{w}, \omega)$. Let ω^{w}_i be a model of K_i such that $d^H(\mathsf{w}, \omega^{\mathsf{w}}_i) = d^H(\mathsf{w}, K_i)$.

We have $d_{H,\Sigma}(\mathsf{w}, E) = d_H(\mathsf{w}, \omega^{\mathsf{w}}_1) + \ldots + d_H(\mathsf{w}, \omega^{\mathsf{w}}_n)$.

Now, by definition of the atom swap noise, the probability that ω^{w}_i is a model of K_i given the interpretation w is $P^v_{\mu,\mathsf{w}}((\omega^{\mathsf{w}}_i, i)) = p^{d_H(\mathsf{w},\omega^{\mathsf{w}}_i)+1}$. In addition $P((\omega^{\mathsf{w}}_1, 1) \cap \ldots \cap (\omega^{\mathsf{w}}_n, n)) = p^{d_H(\mathsf{w},\omega^{\mathsf{w}}_1)+1} \times \ldots \times p^{d_H(\mathsf{w},\omega^{\mathsf{w}}_n)+1}$ since the atom swap noise satisfies independence. Furthermore $p^{d_H(\mathsf{w},\omega^{\mathsf{w}}_1)+1} \times \ldots \times p^{d_H(\mathsf{w},\omega^{\mathsf{w}}_n)+1} = p^n \times p^{d_H(\mathsf{w},\omega^{\mathsf{w}}_1)+\ldots+d_H(\mathsf{w},\omega^{\mathsf{w}}_n)} = p^n \times p^{d_{H,\Sigma}(\mathsf{w},E)}$.

We want to prove that $\Delta^{d_H,\Sigma}$ is a MLE for P^v, which means that $[\Delta^{d_H,\Sigma}_\mu(E)] = argmax_\omega P_{\mu,\omega}(E)$. This is a direct consequence of the following equivalences:

$$\mathsf{w} \in argmax_\omega P_{\mu,\omega}(E) \Leftrightarrow \forall \mathsf{w}' \models \mu,\ P_{\mu,\mathsf{w}}(E) \geq P_{\mu,\mathsf{w}'}(E)\}$$

$$\Leftrightarrow \forall \mathsf{w}' \models \mu,\ p^{d_H(\mathsf{w},\omega^{\mathsf{w}}_1)+\ldots+d_H(\mathsf{w},\omega^{v}_n)} \geq p^{d_H(\mathsf{w}',\omega^{\mathsf{w}'}_1)+\ldots+d_H(\mathsf{w}',\omega^{\mathsf{w}'}_n)}$$

$$\Leftrightarrow \forall \mathsf{w}' \models \mu,\ d_H(\mathsf{w}, \omega^{\mathsf{w}}_1) + \ldots + d_H(\mathsf{w}, \omega^{\mathsf{w}}_n) \leq d_H(\mathsf{w}', \omega^{\mathsf{w}'}_1) + \ldots + d_H(\mathsf{w}', \omega^{\mathsf{w}'}_n)$$

$$\Leftrightarrow \forall \mathsf{w}' \models \mu,\ d_{H,\Sigma}(\mathsf{w}, E) \leq d_{H,\Sigma}(\mathsf{w}', E)$$

$$\Leftrightarrow \mathsf{w} \models \Delta^{d_H,\Sigma}_\mu(E).$$ □

7 Experimental Results

In this section we study how the merging operators identified as maximum likelihood estimators for the world swap noise and for the atom swap noise are efficient noise cancellers in practice. This requires to evaluate the number of agents needed for getting a merged base equals to the true state of the world with high probability. To this end, we perform some experiments. We report hereafter the obtained results for the world swap noise and the atom swap noise.

In our experiments, the true state of the world $\omega^\star$ has been generated by considering every variable v from $\mathcal{P}$ in an iterative way, assigning v to 1 with probability $\frac{1}{2}$. We did not assume any integrity constraint to be fulfilled (i.e., μ is supposed to be a valid formula). Then, for several values of the parameter p, the noise model P^{wsn} (resp. P^v) has been applied to $\omega^\star$ in order to generate 1000 profiles E consisting of n bases for increasing values of n.

Each resulting profile has then been merged using $\Delta^{d_D,\Sigma}$ for the world swap noise and using $\Delta^{d_H,\Sigma}$ for the atom swap noise. A success has been obtained whenever $[\Delta^{d_D,\Sigma}(E)] = \{\omega^\star\}$ for the world swap noise and $[\Delta^{d_H,\Sigma}(E)] = \{\omega^\star\}$ for the atom swap noise. For each value of n, the success rate was defined as the percentage of profiles for which a success has been obtained.

Figure 1 shows some of the results obtained for the world swap noise and the atom swap noise, for several values of p when $m = 10$. Though considering 10 propositional variables may look not that high, one must keep in mind that 2^m possible worlds exist since no integrity constraint is considered, so that the task is to identify a single world in a set of 1024 worlds. To interpret the results correctly, one must remind that the noise increases with p when the world swap

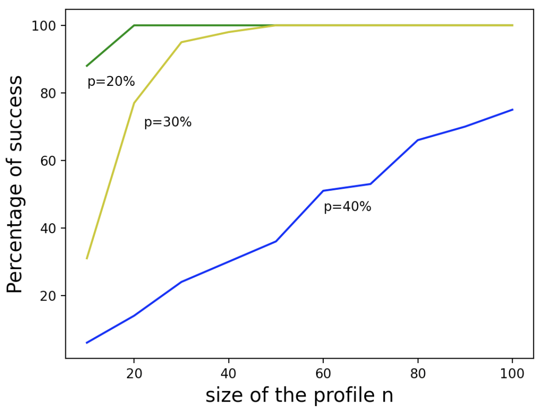

(a) World swap noise

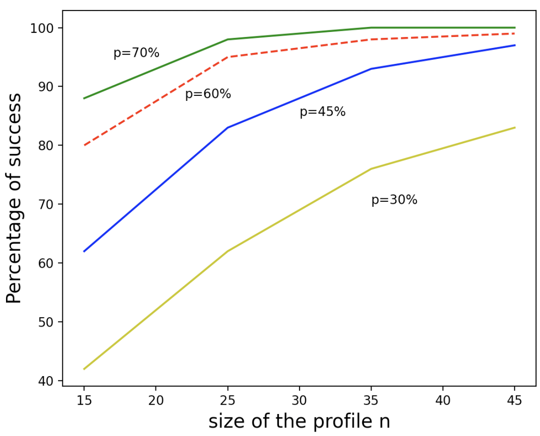

(b) Atom swap noise

Figure 1: Success rate for several values of p, with $m = 10$.

noise is considered, while it decreases with p when the atom swap noise is considered.

Our experimental results confirm the ability of the merging operators $\Delta^{d_D,\Sigma}$ and $\Delta^{d_H,\Sigma}$ to "cancel" their respective noises. It is interesting to note that when the noise level is low ($p = 30\%$ for the world swap noise and $p = 60\%$ for the atom swap noise), as few as 30 bases are enough for being nearly sure ($> 90\%$) to identify the true world. One can observe that for the world swap noise, the number of agents to be considered for getting with high probability a merged base having the true state of the world as its sole model remains quite small, even when the noise level is high. The empirical results also show that the noise cancellation task requires less effort (i.e., less agents) when dealing with the atom swap noise than when dealing with the world swap noise.

Finally, let us note that similar results have been obtained for the other values of m considered in our experiments (we let m vary from 2 to 15; the shapes of the resulting curves are close to those of the curves given in Figure 1, the number of agents required for getting a success increasing with m).

8 Related Work

This work is related to the rationalization of voting rules as maximum likelihood estimators, as discussed in [Conitzer and Sandholm, 2005; Young, 1995; Xia *et al.*, 2010; Conitzer, 2014; Pivato, 2012]. Such an interpretation of voting rules requires to suppose that some candidates are better than others, in an absolute/objective sense, and that the agent's preferences are noisy estimates of the "absolute" agents' ranking. The objective of voting is then to determine this best ranking from the agents' votes.

This work has also connections with the truth tracking problem [Bovens and Rabinowicz, 2006; Pigozzi and Hartmann, 2007], that has been studied for belief merging in [Everaere *et al.*, 2010a]. This problem consists in determining, when the agents are reliable enough, whether the true state of the world $\omega^\star$ can be identified in the limit (i.e., when the number of agents participating in the merging process tends to infinity) as the unique model of the corresponding merged base. In [Everaere *et al.*, 2010a], an extension of Condorcet's jury theorem to belief merging has been pointed out, giving a positive answer to the issue under some standard assumptions about the independence and the reliability of the agents.

The maximum likelihood estimator question differs from the truth tracking one in several aspects. In the truth tracking case, $\omega^\star$ is known and every agent is more or less reliable. In the maximum likelihood case, $\omega^\star$ is not known, agents have a noisy perception of it, and the corresponding noise model is more or less known. Furthermore, the observed profile is fixed (so the number of agents does not tend to infinity). So, while the truth tracking problem and the MLE one are related to the issue of providing an epistemic justification for belief merging operators, they offer complementary views.

9 Conclusion

In this work we carried out an evaluation of belief merging operators as maximum likelihood estimators. We have shown close connections between noise models and postulates for IC merging. Especially, we have proved that a merging operator that is a MLE for an independent and uniform noise model has to satisfy most IC merging postulates.We have also shown that all distance-base merging operators based on the Σ^k aggregation functions are MLEs for some noise. We have studied two specific, yet sensible noise models: the world swap noise and the atom swap noise. We have shown that some well-studied distance-based merging operators are MLEs for those noise models. Finally, we have reported results from experiments showing that finding the true state of the world is feasible in practice for each of the two specific noises. Based on those results, our work provides a new justification for IC merging operators, of epistemic nature.

This work calls for several perspectives for further research. One of them consists in identifying some additional pairs (P, Δ) where Δ is a MLE for the noise model P. In application scenarios where merging beliefs is required, whenever the noise model can be guessed, a catalog of such pairs would be a useful methodological tool for deciding which belief merging operator to choose.

Acknowledgements

This work has benefited from the support of the AI Chair BE4musIA of the French National Research Agency (ANR).

References

[Bovens and Rabinowicz, 2006] Luc Bovens and Wlodek Rabinowicz. Democratic answers to complex questions. *Synthese*, 150:131–153, 2006.

[Condorcet, 1785] Marquis de Condorcet. *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. Imprimerie Royale, Paris, 1785.

[Conitzer and Sandholm, 2005] Vincent Conitzer and Tuomas Sandholm. Common voting rules as maximum likelihood estimators. In *Proceedings of the Twenty-First Conference on Uncertainty in Artificial Intelligence (UAI'05)*, pages 145–152, 2005.

[Conitzer, 2014] Vincent Conitzer. The maximum likelihood approach to voting on social networks. In *International Symposium on Artificial Intelligence and Mathematics, ISAIM 2014, Fort Lauderdale, FL, USA, January 6-8, 2014*, 2014.

[Elkind and Slinko, 2016] Edith Elkind and Arkadii Slinko. *Rationalizations of Voting Rules*, chapter 8. Handbook of Computational Social Choice. Cambridge University Press, 2016.

[Everaere *et al.*, 2007] Patricia Everaere, Sébastien Konieczny, and Pierre Marquis. The strategy-proofness landscape of merging. *Journal of Artificial Intelligence Research*, 28:49–105, 2007.

[Everaere *et al.*, 2010a] P. Everaere, S. Konieczny, and P. Marquis. The epistemic view of belief merging: Can we track the truth? In *Proceedings of the Twenty-Ninth European Conference on Artificial Intelligence (ECAI'10)*, pages 621–626, 2010.

[Everaere *et al.*, 2010b] Patricia Everaere, Sébastien Konieczny, and Pierre Marquis. Disjunctive merging: Quota and Gmin merging operators. *Artificial Intelligence*, 174(12-13):824–849, 2010.

[Haret *et al.*, 2016] Adrian Haret, Andreas Pfandler, and Stefan Woltran. Beyond IC postulates: Classification criteria for merging operators. In *Proceedings of the Twenty Second European Conference on Artificial Intelligence (ECAI'16)*, pages 372–380, 2016.

[Konieczny and Pino Pérez, 2002a] Sébastien Konieczny and Ramon Pino Pérez. Merging information under constraints: a logical framework. *Journal of Logic and Computation*, 12(5):773–808, 2002.

[Konieczny and Pino Pérez, 2002b] Sébastien Konieczny and Ramon Pino Pérez. On the frontier between arbitration and majority. In *Proceedings of the Eighth International Conference on Principles of Knowledge Representation and Reasoning (KR'02)*, pages 109–118, 2002.

[Konieczny and Pino Pérez, 2011] Sébastien Konieczny and Ramon Pino Pérez. Logic based merging. *Journal of Philosophical Logic*, 40(2):239–270, 2011.

[Konieczny *et al.*, 2004] Sébastien Konieczny, Jérôme Lang, and Pierre Marquis. DA^2 merging operators. *Artificial Intelligence*, 157:49–79, 2004.

[Pigozzi and Hartmann, 2007] Gabriella Pigozzi and Stephan Hartmann. Judgment aggregation and the problem of truth-tracking. In *Proceedings of the Eleventh International Conference on Theoretical Aspects of Rationality and Knowledge (TARK'07)*, pages 248–252, 2007.

[Pivato, 2012] Marcus Pivato. Voting rules as statistical estimators. *Social Choice and Welfare*, 40:581–630, 2012.

[Xia *et al.*, 2010] Lirong Xia, Vincent Conitzer, and Jérôme Lang. Aggregating preferences in multi-issue domains by using maximum likelihood estimators. In *Proceedings of the Ninth International Conference on Autonomous Agents and Multiagent Systems (AAMAS'10)*, pages 399–408, 2010.

[Young, 1995] Hobart Peyton Young. Optimal voting rules. *Journal of Economic Perspectives*, 9(1):51–64, 1995.

[Young, 2003] Hobart Peyton Young. Condorcet's theory of voting. *The American Political Science Review*, 82(4):1231–1244, 2003.