ALDO FRANCO DRAGONI AND PAOLO GIORGINI

REVISING BELIEFS RECEIVED FROM MULTIPLE SOURCES

ALDO FRANCO DRAGONI AND PAOLO GIORGINI

REVISING BELIEFS RECEIVED FROM MULTIPLE SOURCES

ABSTRACT: Since the seminal, philosophical and influential works of Alchourrón, Gärdenfors and Makinson, ideas on "belief revision" have been progressively refined toward normative, effective and computable paradigms. Side by side to this "symbolic" line of research, there has been also a "numerical" approach to belief revision whose main contributes were the probabilistic and the evidence-based approaches. The opinion expressed in this paper is that, to be applied in a multi-sources environment, belief revision has to depart considerably from the original framework. In particular, it has to abandon the fundamental principle of "Priority to the Incoming Information" in preference to what we called the principle of "Recoverability". Furthermore the semantic approach should be blended with a syntactic treatment of consistency inspired by the Truth Maintenance Systems.

1 A BRIEF RETROSPECTIVE

1.1 The AGM paradigm

During the last decade, the logical framework layed down by Alchourrón, Gärdenfors and Makinson [1] deeply influenced the notion of belief revision. Conceiving both, the *knowledge space* K and incoming information p, as sentences of a propositional language L, they introduced three rational principles for the revision operator (*):

AGM1 *Consistency*: K^*p must be consistent
AGM2 *Minimal Change*: K^*p should alter as little as possible K
AGM3 *Priority to the Incoming Information*: p must belong to K^*p

From these principles they drew eight postulates:

K^*1. For each p and K, K^*p is still a cognitive state
K^*2. $p \in K^*p$
K^*3. $K^*p \subseteq K^+p$
K^*4. If $\neg p \notin K$ then $K^+p \subseteq K^*p$
K^*5. K^*p is inconsistent iff p is inconsistent
K^*6. If p and q are logically equivalent then $K^*p = K^*q$
K^*7. $K^*(p \wedge q) \subseteq (K^*p)^+q$
K^*8. If $\neg q \notin K^*p$ then $(K^*p)^+q \subseteq K^*(p \wedge q)$

These axioms describe the rational properties that revision should obey, but they do not suggest how to perform it. An operative definition came from the Levi's identity:

$$K^*p = Th((K^-\neg p) \cup \{p\})$$

(where "Th" represents the deductive closure operator) that suggests to delete $\neg p$ from K (reducing in some way K at a point that $\neg p$ is no longer derivable), add p and "calculate" the deductive closure. The deletion of $\neg p$ from K, $K^-\neg p$, is called *contraction*, and can be defined in terms of "Epistemic Entrenchment" [14], which is an ordering $\leq$, that envisages the logical dependencies of the formulae in K; it depends on K but it applies to all the formulae of L. $p \leq q$ means that p is less entrenched (more exposed to eventual changes) than q. $\leq$ satisfies the following postulates:

EE1. $\leq$ is transitive
EE2. For all $p, q \in L$, if $p \vdash q$ then $p \leq q$
EE3. For all $p, q \in L$, either $p \leq p \wedge q$ or $q \leq p \wedge q$
EE4. If K is consistent, then $p \notin K$ iff for all $q \in L$, $p \leq q$
EE5. If for all q of L, it holds $q \leq p$, then p is a tautology

Contraction could be defined from the Epistemic Entrenchment as follows: $q \in K^-p$ iff $q \in K$ and, either $p < q \vee p$, or p is a tautology.
There are three problems with such a kind of revision:

1. it deals with infinite sets of sentences

2. $\leq$ depends on K, so it is difficult to iterate the revision because the ordering defined on K^*p could be different from the one defined on K

3. the choice of a particular ordering $\leq$ satisfying the postulates EE1$\div$EE5 is arbitrary; as Gärdenfors wrote: "[the postulates] leave the main problem unsolved: what is a reasonable metric for comparing different epistemic states?".

Mary-Anne Williams [26] showed how the first two problems "can be solved in theoretically satisfying ways wholely within the AGM paradigm". She pointed out that, belief revision means "epistemic entrenchment revision": incoming information *transmutes* the old epistemic entrenchment into a new one which, in turn, yields a different contraction and, hence, a new revised knowledge space. Given a knowledge space K, let a *base* B be a finite set of sentences of L such that K is the deductive closure of B. She introduced a ranking function on a base B, called *partial entrenchment ranking* (p.e.r), that she proved can be regarded as a specification of a finite epistemic entrenchment on K. AGM revision for finite bases can be iterated by transmuting partial entrenchment rankings. The transmutation is provoked by an incoming, consistent and not-tautological sentence p with a degree of firmness i. A transmuted p.e.r. is one such that p has a degree i and the other beliefs are rearranged in order to respect epistemic entrenchment.

She also introduced a particular transmutation, called *adjustment*, that inroduces minimal changes such that p is accepted with degree i.

Partial entrenchment rankings and transmutations respect Dalal's "principle of irrelevance of the syntax" [3], according to whom syntactically different but logically equivalent formulae represent the same knowledge space. The partisans of *syntax-dependent* belief revision consider knowledge spaces made up of a limited number of sentences. They claim that asserting facts is more important than deriving others from them. Nebel's *epistemic relevance ordering* [20] stratifies a base B into n priority classes $B_1, ..., B_n$. Epistemic relevance does not respect the logical contents of the sentences as epistemic and partial entrenchment do. A justification seems to rely on the logical paradoxes of the material implication: a rule $q \rightarrow p$ should not necessarily be considered more important than p just because $p \vdash q \rightarrow p$. Let $B{\downarrow}p$ denote the set of the subsets of B that fail to imply p. Nebel defines $B{\Downarrow}p$ as the subset of $B{\downarrow}p$ made of the elements that *contain as many sentences of the highest priority as possible*. The corresponding revision is defined as:

$$B \oplus p = Th\left(\left(\bigcap_{B' \in (B \Downarrow \neg p)} Th(B') \right) \cup \{p\} \right)$$

where $Th(B')$ denotes the deductive closure of B'. There are two problems with this revision:

- it does not satisfy all the AGM postulates
- it is still computationally hard.

We could adopt various criteria to sort and select the elements of $B{\downarrow}p$. Let $B' = B_1^{'} \cup ... \cup B_n^{'}$ and $B'' = B_1^{''} \cup ... \cup B_n^{''}$ two consistent subsets of B where $B_i^{'} = B^{'} \cap B_i$ and $B_i^{''} = B^{''} \cap B_i$. Benferhat et al. [13] suggest (implicitly) three ways to translate the epistemic relevance into a preference relation $\leq$ on $B \downarrow p$.

- *best-out ordering.* $B^{''} \leq B^{'}$ iff the most credible of the sentences in $B \setminus B^{''}$ is less credible than the most credible of the sentences in $B \setminus B^{'}$.
- *inclusion-based ordering.* $B^{''} \leq B^{'}$ iff there exists a stratum i such that $B_i^{'} \supset B_i^{''}$ and for any $j < i$, $B_j^{'} = B_j^{''}$. This preordering is strict but partial; its maximal consistent elements are also maximal for the best-out ordering.
- *lexicographic ordering.* $B^{''} \leq B^{'}$ iff there exists a stratum i such that $|B_i^{'}| > |B_i^{''}|$ and for any $j < i$ $|B_j^{'}| = |B_j^{''}|$, and $B^{''} = B^{'}$ iff for any j, $|B_j^{'}| = |B_j^{''}|$.

$B{\Downarrow}p$ contains the elements of $B{\downarrow}p$ maximal w.r.t. inclusion-based ordering.

1.2 The Truth Maintenance Systems

Even if the AGM paradigm cannot be set aside when speaking of belief revision, and even if Gärdenfors deeply contributed to the development of the field [14; 15], we must recognize that a parallel conception of belief revision started, even before the AGM framework, from researches on the so called Truth Maintenance Systems [4] (see, for instance, Martins and Shapiro's modelization [18]). Expecially De Kleer's ATMS [17], overcoming some limitations of Doyle's TMS, was immediately regarded as a powerful reasoning tool to revise beliefs (and to perform the dual cognitive operation of *diagnosis*). Crucial to the ATMS architecture is the notion of *assumption* which designates a decision to believe something without any commitment as to what is assumed. An *assumed* datum is the problem-solver output that has been assumed to hold after the assumption utilized to derive it. Assumptions are connected to assumed data via justifications and from the foundation to which every datum's support can be ultimately traced. The same assumption may justify multiple data or one datum be justified by multiple assumptions. Taking the notion of logical consequence as justification, the ATMS appears as a mechanism to perform syntax-dependent belief revision for finite bases.

1.3 Numerical approaches

Let us come to the third problem with the AGM approach: "what is a reasonable metric for comparing epistemic states?" . We believe that, practically, such computable and reasonable metric can be better provided by numerical approaches, for their natural ability to deal with fine grained values. In particular, when the epistemic states should be ranked according to some degree of certainty (or credibility), then such a metric might be provided by classic uncertainty management techniques. In fact, the knowledge space of a complex knowledge-based system does not suffer only from inconsistency; it can also be affected by uncertainty. Numerical distributions of credibility over the sentences of L or over the set Ω of the models of L, play the same role that "epistemic entrenchment", p.e.r. and epistemic relevance play in the symbolic frameworks. Generally, numerical approaches do not respect logical dependencies among the sentences. Logics of uncertainty often represent a cognitive state K and the incoming information p in terms of their sets of models (also said "possible worlds"), respectively, $[K]$ and $[p]$. A cognitive state is represented not simply by $[K]$, but by an assignment function $d(\omega) : \Omega \rightarrow [0, 1]$ such that $d(\omega^{'}) = 0$ for each $\omega^{'} \notin [K]$. The arrival of p generally means that the real world belongs to $[p]$. This event changes d into a new assignment (new prioritization) d'. Imposing the priority to the incoming information (AGM3) means assigning $d'(\omega) = 0$ to each $\omega \notin [p]$. Minimizing this change (AGM2) means minimizing some kind

of distance between d and d'.

In the probabilistic approach [22] a cognitive state is characterized by a *probability measure* P on 2^{Ω}, whose fundamental property is *additivity*: $\forall A, B \subseteq \Omega, A \cap B \neq \oslash \Rightarrow P(A \cup B) = P(A) + P(B)$. $P(\Omega) = 1$, so if $\bar{A} = \Omega - A$ then $P(A) + P(\bar{A}) = 1$. We might also consider the *probability distribution* $pr(\omega)$ that assignes a probability degree to each world in Ω, where $P(A) = \sum_{\omega \in A} pr(\omega)$. $pr(\omega) = 0$ means that ω is not a possible world. $pr(\omega) = 1$ means that ω is surely the real world. An incoming information p changes the probability measure of any sentence q of L through the very famous Bayes' Conditioning Rule:

$$P(q|p) = \frac{P([q] \cap [p])}{P([p])} = \frac{P([p]|[q]) \cdot P([q])}{P([p])}$$

which can also be expressed in terms of probability distribution:

$$pr(\omega|p) = \begin{cases} \dfrac{pr(\omega)}{P([p])} & \text{if } \omega \in [p] \\ 0 & \text{otherwise} \end{cases}$$

This modification is defined only for $P([p]) > 0$, hence it is not applicable when p is judged impossible by the previously determined probability measure P.
Bayesian conditioning obeys the principle of priority to incoming information (AGM3); it increases the probability of the not-impossible worlds belonging to $[p]$ to the prejudice of those external to $[p]$ which become all impossible.

In the probabilistic framework the probability of a sentence p is simply the probability measure $P([p])$. Thus, probability measures order the sentences of L, but, unfortunately, they do not generate epistemic entrenchments. In effect, probability measures satisfy EE1 since they are, obviously, transitive (if $P([p]) \leq P([q])$ and $P([q]) \leq P([r])$ then $P([p]) \leq P([r])$. EE2 too is verified since $p \vdash q$ means $[p] \subseteq [q]$ hence $P([p]) \leq P([q])$ (it is always easier to retract p than q). Even EE4 is verified; in fact, $p \notin K$ means $P([p]) = 0$, and if K is consistent then there are sentences q such that $P([q]) = 0$, hence $P([p]) = 0$ iff $\forall q \in L\ (P([p]) \leq P([q]))$. Finally, EE5 is verified since if $\forall q \in L\ P([q]) \leq P([p])$ then $[p] = \Omega$ which means that p is a tautology. Unfortunately, EE3 is generally unsatisfied since $[p \wedge q] \subseteq [p]$ and $[p \wedge q] \subseteq [q]$ so that $P([p \wedge q]) \leq P([p])$ and $P([p \wedge q]) \leq P([q])$; normally it is easier to retract a conjunction than any of its conjuncts.

Also the belief function framework [24; 25] assignes a probability P to the

subsets of Ω, with the constraints $P(\oslash) = 0$ and $\sum_{A \subseteq \Omega} P(A) = 1$. If $P(A) > 0$ then A is said to be a *focal element*. The *belief function* on the subsets of Ω is defined as

$$Bel(A) = \sum_{X \subseteq A} P(X)$$

$Bel(A)$ measures the persuasion that the real world is inside A; maybe that there is no evidence that directly support A but it cannot be excluded because there is evidence that supports some of its subsets. This function is not additive: $Bel(A) + Bel(\bar{A}) \leq 1$. The knowledge is:

- *certain* and *precise* if there exists a $\omega \in \Omega$ such that $P(\{\omega\}) = 1$
- *certain* and *imprecise* as if there exists a $A \subset \Omega$ such that $P(A) = 1$ but A is not singleton
- *consistent* if all the focal elements are nested
- *inconsistent* if all the focal elements are disjoint
- *void* if $P(\Omega) = 1$ and for all $A \subseteq \Omega$, $P(A) = 0$.

This framework deals also with uncertain inputs. They are treated as new probability assignments on 2^{Ω}. The change consists of merging the two evidences (the prior P_1 and the new P_2) through the Dempster's Rule of Combination:

$$P(A) = \frac{\sum_{X_1 \cap X_2 = A} P_1(X_1) \cdot P_2(X_2)}{\sum_{X_1 \cap X_2 = \oslash} P_1(X_1) \cdot P_2(X_2)}$$

for all $A \subset \Omega$. This rule, easily extensible to combine n probability assignments, reinforces concordant evidence and weakens conflicting ones. It can be applied only if evidences are independent and referred to the same Ω. Because of the commutativity of the product, the rule is independent from the sequence $P_1 ... P_n$ so *it violates the principle of priority to the incoming information*! From a knowledge engineering point of view, the worst problem with the Dempster's Rule of Combination is its computational complexity. One should generate a frame of $2^{|\Omega|}$ elements to calculate it! However, much work has been spent in reducing the complexity of that rule. Such methods range from "efficient implementations" [16] to "qualitative approaches" [21] through "approximate techniques" with statistical methods as the Montecarlo sampling algorithm [27; 19].

2 REQUIREMENTS FOR A BELIEF REVISION FRAMEWORK IN A MULTI SOURCE ENVIRONMENT

We think that to revise beliefs in a Multi-Agent scenario, where many sources give information about a same static situation, the framework should satisfy some requisites.

- Ability to reject incoming information

A belief revision system for multi source environments should drop the rationality principle of "priority to the incoming information" which is not acceptable since there is no strict correlation between the chronology of the informative acts and the credibility of their contents [12]; it seems more reasonable to treat all the available pieces of information as they had been collected at the same time.

- Ability to recover previously discarded beliefs

Cognitive agents should be able to recover previously discarded pieces of knowledge after that new evidence redeems them. The point is that this should be done not only when the new information directly "supports" a previously rejected belief, but also when the incoming information indirectly supports it, by disclaiming the beliefs that contradicted it, causing its ostracism. More formally, for each cognitive state K, and sentences p and q such that $K \vdash p$ and $K^*q \nvdash p$, there can always be another piece of information r such that $(K^*q)^*r \vdash p$, even if $r \nvdash p$. An obvious case should be $r = \neg q$. We elsewhere called this rule *principle of recoverability*: "any previously held piece of knowledge must belong to the current knowledge space if consistent with it" [12; 6; 8]. The rationale for this principle is that, if someone gave us a piece of information (sometime in the past) and currently there is no reason to reject it, then we should accept it! This is stronger than the traditional "coherence" spirit of belief revision, since the piece of knowledge to accept is not a *generic* sentence of the language but a *generated* piece of information; somewhere there is an utilitarian intelligent information source that guarantees for it. Of course, this principle does not hold for updating, where changes may be irrevocable. This feature could also be subtitled: "revocable treatment of consistency" . We remember of Minsky's lection: "I do not believe that consistency is necessary or even desirable in a developing intelligent system ... What is important is how one handles paradoxes or conflicts ... *Enforcing consistency produces limitations.* As we will see in a moment, we overcome this problem by defining a single global, never forgetting, eventually inconsistent *Knowledge Background*, upon which act multiple specific, competitive, ever changing, *consistent* cognitive states.

- Ability to *combine* contradictory and concomitant evidences

The notion of *beliefs integration* should blend that of revision [9]. Every incoming information changes the cognitive state. Rejecting the incoming information does not mean leaving beliefs unchanged since, in general, incoming information alters the distribution of the weights. Surely the last incoming information decreased the credibility of the beliefs with whom it got in contradiction, even in the case that it has been rejected. The same when receiving a piece of information which we were already aware of; it is not the case that nothing happened (as AGM K^*4 states) since we are now, in general, more sure about that belief. More generally, there is no reason to limit the changes introduced by the new information to an insertion into a pre-established relative order with consequent rearrangement of the ranking to accomplish the logical relations between beliefs (as Williams' transmutation does). If it is true that new incoming information affects the old one, it is likewise true that the latter affects the former. In fact, an autonomous agent (where "autonomous" means that his cognitive state is not *determined* by other agents) judges the credibility of new information on the basis of its previous cognitive state. "Revising beliefs" should simple mean "dealing with a new broader set of pieces of information".

- Ability to deal with couples <source,information> rather than with information alone

The way the credibility ordering is generated and revised must reflect the fact that beliefs come from different sources of information, since the reliability and the number of independent informants affect the credibility of the information and vice-versa [5].

- Ability to maintain and compare multiple candidate *cognitive states*

This ability is part of humans intelligence which does not limit its action to comparing single pieces of information but goes on trying to reconstruct alternative cognitive scenarios as far as it is possible.

- Sensibility of the syntax

Despite Dalal's mentioned principle, syntax plays an important role in everyday life. The way we pack (and unpack) pieces of information reflects the way we organize thinking and judge credibility, importance, relevance and even truthfulness. A testimony of the form $\alpha \wedge \delta \wedge ... \wedge \zeta \wedge \neg\alpha$ from a defendant A in a trial has the same semantic truth value than the testimony $\beta \wedge \neg\beta$ from defendant B, but we remember many cases in which B has been condemned while A has been absolved, being regarded his/her testimony "partially true", contrasting with the B's one regarded as "absolutely contradictory". A set of sentences seems not to be logically equivalent to their conjunction and we could change a cognitive state by simply clustering the same beliefs in a different way.

3 OUR BELIEF REVISION FRAMEWORK FOR A MULTI SOURCE ENVIRONMENT

From this discussion it should be evident that we cannot rely on the AGM framework to build a method for revising belief in a multi source environment. We better repose on the ATMS abstract conception. Our sentence-based approach for belief revision [6] conceives two knowledge repositories:

1. the *knowledge background* KB, which is the set of all the propositional sentences available to the reasoning agent (as assumptions); since it can be inconsistent, it cannot be used as a whole to support reasoning and decision processes

2. the *knowledge base* $B \subseteq KB$, which is the maximally consistent, currently preferred piece of knowledge that should be used for reasoning and decision supporting; since it is determined by imposing the maximal consistency, it may contain pieces of knowledge with low degrees of credibility.

The reason why we pick *maximally* consistent subsets of KB is that this is the perfect implementation of the "Principle of Recoverability": everything were known in the past should be believed again now (or still believed) if nothing prevents that. Computationally, our way to belief revision consists of four steps (figure 1) [8; 9]:

S1. Detection of the minimally inconsistent subsets of $KB \cup \{p\}$ (*nogoods*)
S2. Generation of the maximally consistent subsets of $KB \cup \{p\}$ (*goods*)
S3. Revision of the credibility weights of the sentences in $KB \cup \{p\}$
S4. Choice of a preferred *good* as the new revised base B'

The incoming information p, with its weight of evidence, is evaluated not just within the current base B, but within the overall knowledge background KB. Doing so, the degrees of credibility of the sentences in $KB \cup \{p\}$ are reviewed on a broader and less prejudicial basis (S3). As already explained, the main advantage is that we can rescue sentences from KB by virtue of the maximal consistency of B'. If we'd revise only B by p, we could not recover information from KB. For instance, Nebel's revision would select some $B' \in B \Downarrow \neg p$, but it will be always possible to find out some $B'' \in KB \Downarrow \neg p$ such that $B' \subseteq B''$.

S4 might choose a new base B' syntactically equal to the previous B (meaning that p has been rejected) but, in general, B' will have a different credibility distribution than B. p might be rejected even if S4 chooses a base B' different from B, but that still containing sentences incompatible with p.

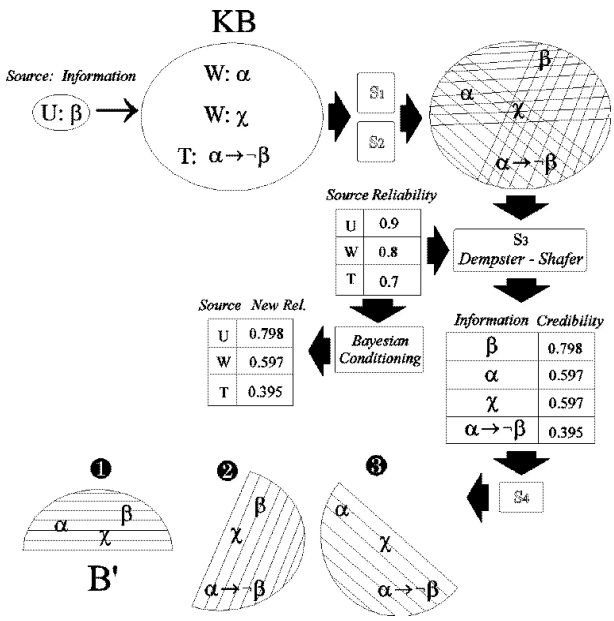

Figure 1. Dempster's and Bayes rules at work

When p is consistent with B, not necessarily $B'=B\cup\{p\}$, since S3 may yield a totally different choice at S4. Previously rejected pieces of knowledge $R\subset KB$ can be rescued simply by determining some upsetting between the credibility of a set $S\subset B$ and the credibility of R, this may happen if p *supports* R against S. The rejection of the priority to the incoming information principle implies that K^*4 and and K^*5 hold no longer (if p is inconsistent it will be part of none of the goods produced at S2, so it will never be part of a base).

S1 and S2 deal with consistency and derivation, and act on the symbolic part of the information. Operations are in ATMS style; to find out nogoods and goods, we adopt (and adapt) the most efficient set-covering algorithm that we are aware of [23]. Notwithstanding this, even in the propositional case, determining all the minimal inconsistencies can be very hard. However, such condition can be relaxed (the consequence is that some of the goods are not really consistent) and in practical applications dealing with commonsense knowledge (for instance [11]), such minimal inconsistencies could be provided interactively by the external user.

S3 and S4 deal with uncertainty and work with the numerical weight of the information. Both contribute to the choice of the revised knowledge

space so their reasonableness should be evaluated as a couple. Numerical formalisms are able to perform both of them since the credibility of a single sentence p is determined in the same way as the credibility of a set of sentences B by the weights attached to $[p]$ and $[B]$, respectively. Flexibility is an advantage in separating the two steps; for instance, depending on the characteristics of the knowledge domain under consideration and the kind of task and/or decision that should be taken on the basis of the revision outcome, the selection function could consider also one (or a combination) of the methods described in [2].

Probabilistic methods with uncertain inputs seem inadequate for the strong dependence that they impose on the credibility of a sentence and that of its negation. The "drowning problem" is a limit of the possibilistic approach (see [2]). We see that the belief-function formalism, in the special guise in which Shafer and Srivastava apply it to auditing [25], could work well because it treats all the pieces of information as they had been provided at the same time.

The method has the following I/O:

INPUT: list of pairs <source,piece of information>
list of pairs <source,reliability>
OUTPUT: list of pairs <piece of information,credibility>
list of pairs <source,reliability>

Let S=$\{s_1, ..., s_n\}$ be the set of the sources, and let kb_i be the subset of KB received from s_i. Each source s_i is associated with a *reliability* $R(s_i)$, that is regarded as the *probability* that the source is faithful. The main idea with this multi-source version of the belief function framework is that a reliable source cannot give false information, while an unreliable source can give correct information; the hypothesis that s_i is reliable is compatible only with the models of kb_i, while the hypothesis that s_i is unreliable is compatible with the overall Ω. Each source s_i is an evidence for KB and generates the following *bpa* $m_i(\cdot)$ on 2^Ω:

$$m_i(X) = \begin{cases} R(s_i) & \text{if } X = [kb_i] \\ 1 - R(s_i) & \text{if } X = \Omega \\ 0 & \text{otherwise} \end{cases}$$

All these *bpas* will be then combined through the Dempster Rule of Combination. From the combined *bpa* $m(\cdot)$, the credibility of a sentence p of L is given, as usual, by:

$$Bel(p) = \sum_{X \subseteq [p]} m(X)$$

From this mechanism we obtained an easy way to calculate the new reliabilities of the sources. Let Φ be an element of 2^S. If the sources are independent, the reliability of Φ is

$$R(\Phi) = \prod_{S \in \Phi} R(s) \cdot \prod_{S \in \Phi} (1 - R(s))$$

It holds that

$$\sum_{\Phi \in 2^S} R(\Phi) = 1$$

It maybe that some source fall in contradiction, so that some elements of 2^S are impossible. The remaining elements are subjected to bayesian conditioning so that their reliabilities sum up again to 1. The revised reliability $R^*(s)$ of a source s is the sum of the new reliabilities of the surviving elements of 2^S that contain s. If a source has been involved in some contradictions, then $R^*(s) \leq R(s)$, otherwise $R^*(s) = R(s)$.

The main problem with this formalism is its computational complexity. Contradictions help us because they reduce the size of Ω. However, to cancel an impossible world, we must first generate it. This imposes severe limits on the number of sentences acceptable in input (about 30) even for "off-line" applications.

Yet, from an input-output analysis of the mechanism we realized that two properties hold.

1. Sentences not involved in contradictions and received from a single source s, do not contribute to the mechanism, in the sense that the degrees of credibility of the other sentences do not depend on their presence. Their own credibility is $R^*(s)$.

2. Multiple contradictions involving sentences received exclusively from exactly the same sources are redundant; all the sentences from the same source receive the same degree of credibility, independently of the number and the cardinality of the contradictions.

Property (1) implies that we can temporarily leave out of the process those sentences received from a single source, and that are not involved in contradictions. In many cases, this dramatically reduces the size of Ω. Property (2) says that, what is important is that a set of sources was contradictory, not how many times nor about what or about how many sentences they did. This allows us to temporarily leave out of the process also some sentences involved in contradictions; this is significant in situations like that of two

sources systematically in contradiction only with each others.

S4 translates such ordering on the *sentences* in $KB \cup \{p\}$ into an ordering on the *goods* of $KB \cup \{p\}$. The best classified good is selected as the preferred revised knowledge base. If the ordering on $KB \cup \{p\}$ is not strict, then there can be multiple preferred goods. In this case we could take their intersection as revised knowledge base ([2]); however, the intersection is not maximally consistent and this means that all the conflicting pieces of knowledge with the same credibility will be rejected.

Another question is: S4 should consider only the *qualitative* ordering of the sentences in $KB \cup \{p\}$ (relative classification without the numerical weights) or could it take advantage of the *explicit* ordering (numerical weights). The first approach seems closer to the human cognitive behavior (which normally refrains from numerical calculus). The second one seems more informative (it takes into account not only relative positions but also the gaps between the items). In our model we do not use the "best-out" ordering for its "drowning effect" [2]. The lexicographic one could be justified in some particular application domains (e.g. diagnosys). The inclusion-based method seems the most reasonable since it eliminates always the least credible one among conflicting pieces of knowledge.

As an example of a numerical way to perform S4, ordering the goods according to their average credibility seems reasonable and easy to calculate. With this method the preferred good may not contain the most credible sentence.

In the belief function framework, a "good" g is an element of Ω, precisely the one in which all the sentences in g are considered "true" and all the sentences out of g are considered "false" . This implies that the belief-function formalism is able to attach directly a degree of credibility to g, bypassing S4 in our framework. Unfortunately, when a good contains only part of the information supplied by a source, the belief-function formalism puts at zero its degree of credibility. This is unreasonable and, unluckily, the event is all but infrequent, so that often the credibility of all the goods is null.

A final step in our revision mechanism is the selection of the derived sentences which are still derivable from B' since the assumptions on which they rely are all contained in B'. Theoretically, it simply consists in applying classical entailment on the preferred good to deduce plausible conclusion from it. We adopted an ATMS and we stored each sentence derived by the Theorem Prover with an *origin set* [18], i.e., a set of basic assumptions which are all *necessary* to derive it. Practically, this step (not represented in figure 1) consists in selecting from the derived sentences, all those whose

origin set is subset of the preferred good. We could relax the definition of origin set to that of a set of basic assumptions used to derive the sentence. This is easier to compute and does not have harmful consequences; the worst it can happen is that, being this relaxed origin set a superset of the real one, it is not certain that it will be a subset of the preferred good as the real one is, and so some derived logical consequences of the preferred good may be not recognized (at first).

Besides recoverability, this computational model for belief revision overcomes various limitations of other classic approaches, in particular:

- the revision can be iterated
- inconsistent incoming information does not yield inconsistent revised knowledge spaces
- the numerical revision is performed on a broader base (the overall KB)
- the revision is more flexible;
- the complete numerical ordering renders the revision as least drastic as possible
- the splitting between the symbolic treatment of the inconsistencies and the numerical revision of the credibility weights, provides a clear understanding of what is going on and lucid explanations for the choices.

We are currently applying this conception of belief revision in a distributed monitoring system [10] and in the police inquiry domain [7].

REFERENCES

[1] C.E. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: Partial meet contraction and revision functions. *The Journal of Simbolic Logic*, 50:510–530, 1985.

[2] S. Benferhat, C. Cayrol, D. Dubois D., J. Lang, and H. Prade. Inconsistency management and prioritized syntax-based entailment. In *Proc. of the 13th Inter. Joint Conf. on Artificial Intelligence*, pages 640–645, 1993.

[3] M. Dalal. Investigations into a theory of knowledge base revision. In *Proc of the 7th National Conference on Artificial Intelligence*, pages 475–479, 1988.

[4] J. Doyle. A truth maintenance system. *Artificial Intelligence*, 12(3):231–272, 1979.

[5] A.F. Dragoni. A model for belief revision in a multi-agent environment. In Werner E. and Demazeau Y., editors, *Decentralized A.I. 3*. North Holland Elsevier Science, 1992.

[6] A.F. Dragoni. Belief revision: from theory to practice. *The Knowledge Engineering Review*, 12(2), 1997.

[7] A.F. Dragoni, C. Ceresi, and V. Pasquali. A system to support complex detective inquiries. In *Proc. of the V Iberoamerican Conference on Computer Science and Law*, La Habana, 1996.

[8] A.F. Dragoni and P. Giorgini. Belief revision through the belief function formalism in a multi-agent environment. In Wooldridge M., Jennings N.R., and Muller J., editors, *Intelligent Agents III*, number 1193 in LNCS. Springer-Verlag, 1997.

[9] A.F. Dragoni and P. Giorgini. Distributed knowledge revision-integration. In *Proc. of the Sixth ACM International Conference on Information Technlogy and Management*. ACM Press, 1997.

[10] A.F. Dragoni and P. Giorgini. Sensor data validation for nuclear power plants through bayesian conditioning and dempster's rule of combination. *Computers and Artificial Intelligence*, 17(2-3):151–168, 1998.

[11] A.F. Dragoni and M. Di Manzo. Supporting complex inquiries. *International Journal of Intelligent Systems*, 10:959–986, 1995.

[12] A.F. Dragoni, F. Mascaretti, and P. Puliti. A generalized approach to consistency-based belief revision. In M. Gori and G. Soda, editors, *Proc. of the Conference of the Italian Association for Artificial Intelligence*, number 992 in LNAI. Springer Verlag, 1995.

[13] D. Dubois and H. Prade. Belief change and possibility theory. In Gärdenfors P., editor, *Belief Revision*. Cambridge University Press.

[14] P. Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. MIT Press., Cambridge Mass., 1988.

[15] P. Gärdenfors. *Belief Revision: an introduction*. Cambridge University Press., 1992.

[16] R. Kennes. Computational aspects of the möbius transform of a graph. *IEEE Transactions in Systems, Man and Cybernetics*, 22:201–223, 1992.

[17] J. De Kleer. An assumption based truth maintenance system. *Artificial Intelligence*, 28:127–162, 1986.

[18] J.P. Martins and S.C. Shapiro. A model for belief revision. *Artificial Intelligence*, 35:25–97, 1988.

[19] S. Moral and N. Wilson. Importance sampling monte-carlo algorithms for calculation of dempster-shafer belief. In *Proc. of IPMU'96*, Granada, 1996.

[20] B. Nebel. Base revision operations and schemes: Semantics, representation, and complexity. In A.G. Cohn, editor, *Proc. of the 11th European Conference on Artificial Intelligence*. John Wiley & Sons, 1994.

[21] S. Parson. Some qualitative approaches to applying the dempster-shafer theory. *Information and Decision Technologies*, 19:321–337, 1994.

[22] J. Pearl. *Probabilistic Reasoning for Intelligent Systems*. Morgan Kaufmann, 1988.

[23] R. Reiter. A theory of diagnosis from first principles. *Artificial Intelligence*, 53, 1987.

[24] G. Shafer. Belief functions. In G. Shafer and J. Pearl, editors, *Readings in Uncertain Reasoning*. Morgan Kaufmann, 1990.

[25] G. Shafer and R. Srivastava. The bayesian and belief-function formalisms a general perspective for auditing. In G. Shafer and J. Pearl, editors, *Readings in Uncertain Reasoning*. Morgan Kaufmann, 1990.

[26] M.A. Williams. Iterated theory base change: A computational model. In *Proc. of the 14th Inter. Joint Conf. on Artificial Intelligence*, pages 1541–1547, 1995.

[27] N. Wilson. A monte-carlo algorithm for dempster-shafer belief. In Smets P. D'Ambrosio B.D. and Bonissone P., editors, *Proc. of the Seventh Conference*, pages 414–417, 1991.