Belief Revision through the Belief Function Formalism in a Multi-Agent Environment

Aldo Franco Dragoni and Paolo Giorgini

Istituto di Informatica, Univ.ersità di Ancona, v. Brecce Bianche, 60131, Ancona,

Abstract. To perform "belief revision" in a multi-agent environment, we followed two main guidelines.

1. replacing the "priority to the incoming information principle" with the "recoverability principle": an anytime received piece of information must belong to the current cognitive state whenever it is possible
2. dealing not just with pieces of information but with couples <source, information>, since the reliability of the source affects the credibility of the information and vice-versa.

The "belief-function" formalism provides a simple and intuitive way to transfer the sources' reliability to the information's credibility, but, along with its computational complexity, it suffers from the fact that the degrees of reliability must be established "a-priori". We give a unitary perspective of this composite subject, suggesting ways to overcome the mentioned limitations.

Content Areas: belief revision, multi-agent systems, reasoning under uncertainty

1. INTRODUCTION

Most of the symbolic and numerical models for belief revision developed so far obey the following three rationality principles [1,2,3,11]:

- *Consistency*: revision must yield a consistent knowledge space
- *Minimal Change*: revision should alter as less as possible the knowledge space
- *Priority to the Incoming Information*: the revised knowledge space must embody the information that caused the revision

However, the last principle is questionable. While it is acceptable when updating the representation of an evolving world, it is not generally justified when revising the representation of a static situation. In this case, the chronological sequence of the informative acts has nothing to do with their credibility or importance; we think in particular of multi-agent domains, in which the information come from different sources [4,5,6]. Furthermore, accepting the incoming information (hoping that it is not inconsistent) and remaining consistent, imply throwing away part of the previously held knowledge, but this change should not be irrevocable. Let K^*p be the cognitive state revised after the newcoming information p. 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 could be $r=\neg q$.

To make practical and useful belief revision in a multi-agent environment, we substitute the priority to the incoming information principle with the following one:

- *Recoverability*: an anytime received piece of information must belong to the current knowledge space if it is consistent with it.

In our view, a theoretical account for belief revision should consider two knowledge bases:

1. the *knowledge background* KB, which is the set of all the pieces of knowledge available to the reasoning agent; 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 maximally consistent, it can contain incredible pieces of knowledge.

The incoming information p, with its weight of evidence, should be confronted not just with B, but with the overal KB. A first advantage in doing so is that the weights of the sentences in $KB \cup \{p\}$ are reviewed on a broader and less prejudicial basis. Let KB be inconsistent. If we'd revise only B by p we could not recover information from KB; partial-meet base revision (Nebel 1994) 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 $B' \subseteq B''$. On the basis of the new ordering induced by p, the system will choose another base $B' \subseteq KB \cup \{p\}$, that may not contain p. B' could also be syntactically equal to B (meaning that p has been rejected), but it will have a different credibility distribution than before. B' may not contain p even in the case that $B' \neq B$, since the new ordering may induce the choice of a base which is different from the preceding one, but still contains sentences incompatible with p; since B' is a maximal consistent subset of KB, it can rescue sentences from KB, even if the incoming information has been rejected.

Our way to belief revision consists of four steps:

1. detection of all the minimally inconsistent subsets of $KB \cup \{p\}$ (*nogoods*)
2. generation of all the maximally consistent subsets of $KB \cup \{p\}$ (*goods*)
3. revision of the credibility weights of the sentences in $KB \cup \{p\}$
4. choice of a preferred maximally consistent subset of $KB \cup \{p\}$ as the new revised base B'.

As already pointed out, when p is consistent with B, not necessarily $B'=B \cup \{p\}$ (expansion), since the new credibility distribution produced at step 3 may yield a totally different choice at step 4. So, the rejection of the priority to the incoming information principle implies that the fourth and the fifth AGM postulates, respectively, "when p is consistent with B, $B^*p=B \cup \{p\}$" and "B^*p is inconsistent iff p is inconsistent", hold no more (if p is inconsistent it will be part of none of the goods produced at step 2, so it will never be part of a base).

Steps 1 and 2 deal with consistency and act on the symbolic part of the information. Operations are in ATMS style; to find out minimal inconsistencies and maximal consistencies, we adopt (and adapt) a well-known set-covering algorithm [7]; it will be presented in the next section.

Steps 3 and 4 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. The belief-function formalism that we present in section three, is able to perform both of them. However, an advantage of separating the two steps consists in flexibility; 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 the cardinality of the alternative goods. Step 3 is that of revising the credibility order of the sentences in KB. We see that the propositional belief-function formalism could do the job. Step 4 translates such ordering on the *sentences* in KB into an ordering on the *goods* of KB. The best classified good can then be selected as the preferred revised knowledge base. If the ordering on KB is not strict, then there can be multiple preferred goods. In this case we could take their intersection as revised knowledge base [8]; however, the intersection is not maximally consistent and this means that all the conflicting pieces of knowledge with the same credibility will be rejected. So it would be a better solution trying to make strict the ordering on the sentences of KB. But what about the ways to perform such translation? One of the expected requisite of this method for belief revision is its ability to recover previously discarded pieces of information. This rescue may happen even if the incoming information do not alter the ordering of the sentences in KB (but simply inserts itself somewhere in that classification), provided that it is inconsistent with the current knowledge base. However, as we have said, we increase the flexibility of the method by allowing that the incoming information yields a new ordering on the sentences in KB. In this case, the newcoming information *p* can rescue a previously rejected piece of knowledge *s* even if *p* is compatible with the current knowledge base B, simply by determining some upsetting between the credibility of the sentences in B and the credibility of *s*. It is almost unreasonable that a numerical revision method decreases the credibility of some sentences in B upon the arrival of an information *p* that is not in conflict with them; what may happen, however, is that *p* supports some other sentences which are not in B but in KB, increasing their credibility w.r.t. those that are in B. Another question is: the translation should consider only the *implicit* ordering of the sentences in KB (classification without the numerical weights) or could take advantage of the *explicit* ordering (numerical weights). The first approach seems closer to the human behavior (which normally refrains from numerical calculus). The second one seems more informative (it takes in count not only relative positions but also detachings) and, hence, more rational. Among the methods of the first kind described in [8], set-inclusion seems the most reasonable one since it corresponds to the intuitive idea of eliminating always the least credible one among conflicting pieces of knowledge. Among the methods of the second kind, ordering the goods according to their average credibility seems reasonable and easy to calculate. The main difference with the set-inclusion method is that now the preferred good may not contain the preferred sentence. However, it has its own drawbacks. In the short run (with goods made of few sentences) there is a strange and unintuitive dependence of the effect of the incoming information on the cardinality of the goods in which it enters. In the long run, the credibility of the goods will become very close to each others so that the differences turn out to be rather insignificant.

2. AN EFFICIENT ALGORITHM TO CALCULATE *GOODS* AND *NOGOODS*

Following de Kleer [9], we define *nogood* a minimally inconsistent subset of KB. Dually, we define *good* a maximally consistent subset of KB. An efficient algorithm to calculate goods and nogoods is the one adopted in [7] to calculate model-based diagnoses. Let F be a collection of sets. An *hitting-set* for F is a set $H \subset \bigcup_{S \in F} S$ such that $H \cap S \neq \emptyset$ for each $S \in F$. An hitting-set is *minimal* if none of its proper subsets is an hitting-set for F. Given F, the set-covering algorithm given in [7] yields all the minimal hitting-sets of F.

Given a finite and inconsistent knowledge background KB, the collection G of the goods and the collection NG of the nogoods are dual. If we remove from KB exactly one element for each nogood in NG, what remains is a good. Hence, G can be found by calculating all the minimal hitting-sets for NG, and keeping the complement of each of them w.r.t. KB.

If F is a collection of sets and $F' \subseteq F$ is the collection of all the minimal elements (w.r.t. set inclusion) of F, then F and F' have the same minimal hitting-sets. This simplifies our task since, in order to generate G, we do not need to generate precisely NG (i.e. the collection of all and only the *minimally* inconsistent subsets of KB) but just a collection $M \supseteq NG$ of inconsistent subsets of KB, which is easier.

The proofs of the following results are in [10]

Lemma. Given a collection F of sets, in any HS-Tree T relative to F generated and pruned with the rules 1-5 described in (Reiter 1987), every minimal element of F (w.r.t. set-inclusion) appears as the label of at least one node of T. ■

The collection $M \supseteq NG$ can be generated while generating the HS-Tree. After reducing in clauses all the sentences in KB, we start a refutation process on it. If we find the empty clause, we label the root of T with the set of clauses in KB that have been involved in the refutation. Such set is not necessarily minimal, i.e., may be that some of them were not indispensable to generate the empty clause. Each clause σ_m in this node labels an arc toward a successor node. To calculate the label of this successor node we start a tentative refutation on $KB \backslash \sigma_m$. If $KB \backslash \sigma_m$ is consistent, then the node is labelled with √, otherwise it is labelled with the proper subset of $KB \backslash \sigma_m$ from which the empty clause has been derived. The process iterates till there are no more successor nodes.

Theorem. Given a collection KB of clauses, in any HS-Tree T generated as described before, every nogood of KB appears as the label of at least one node of T. ■

Summarizing, given KB, the following algorithm calculates NG and G. *Inc*(C) is a function on set of clauses that returns √ if C is consistent, returns a set $D \subseteq C$ if D is inconsistent.

1. NG := ∅, G := ∅
2. When a node *n* needs a new label, give it the label *Inc*(KB\H(*n*)) and add it to NG
3. When the rule 5 (in [7]) applies, eliminate S from NG
4. For each H(*n*) such that *n* is labelled with √, put KB\H(*n*) in G
5. Return NG and G

3 THE BELIEF-FUNCTION FORMALISM IN A MULTI-AGENT ENVIRONMENT

The belief-function formalism is able to assign a degree of credibility both to single sentences and to goods (i.e., set of sentences). However, as we'll see, the way it evaluates the credibility of a good is not quite satisfactory. The belief-function formalism, as Shafer and Srivastava apply it to auditing [6], is a process that takes in input a list of couples <source, information> and a degree of reliability for each source (in the range [0,1]), and gives in output a degree of credibility (still in [0,1]) for any subset of the received pieces of information (hence, for any single information received too). In practice, it transforms an array of reliability values (probability degrees that the given sources are reliable) into an array of credibility values (one for each subset of the received sentences).

An agent exchanging information in a multi-agent domain is often challenged by the problem of establishing which of the contrasting pieces of information coming from different sources is the most credible one, and who, among its colleagues, is the most reliable. The belief-function formalism offers an elegant and reasonable way to solve the first question, provided that the reliability degrees of the sources are known. However, a main problem with such formalism is its computational complexity. In the next section we will give some heuristics and hints to reduce this complexity, making the system able to manage a reasonably large set of sentences. Following Shafer and Srivastava (to which the reader should refer for a less synthetic description and a "philosophical" discussion), we present the formalism starting from the simplest case.

Suppose that X gives the information p. The single couple $<X,p>$ generates two "frames": $S=\{X^r$: X is reliable, X^u: X is unreliable$\}$ and $I=\{p^t$: p is true, p^f: p is false$\}$. Let C be the *compatibility relation* over the elements of the two frames, sCi meaning that s is compatible with i. A fundamental assertion is that *a reliable source cannot give a false information, while an unreliable source can (occasionally) give a correct information*, so we have X^rCp^t, X^uCp^t and X^uCp^f, but not X^rCp^f. The reliability of X is regarded as a "probability" of X's being reliable, so $P(X^r)+P(X^u) = 1$. In general, from the probability distribution P on the sources' frame S, a function Bel will assign a degree of credibility to every subset B of the information's frame I as follows:

$Bel(B)=\Sigma\{P(\{s\}) \mid \text{if } s\in S,\ i\in I \text{ and } sCi \text{ then } i\in B\}$ (1

i.e., $Bel(B)$ is the sum of the probabilities of the elements of S such that all the respective compatible elements of I belong to B. The extreme points are $Bel(\varnothing)=0$ and $Bel(I)=1$. $Bel(B)$ is not a probability measure since it is not additive: $Bel(B)=\alpha \not\Rightarrow Bel(\bar{B})=1-\alpha$. In the simplest case of a single couple $<X,p>$, the frame I has four subsets: $\varnothing$, $\{p^t\}$, $\{p^f\}$ and I. $Bel(\varnothing)=P(\varnothing)=0$, $Bel(\{p^t\})=P(\{X^r\})=0.9$, $Bel(\{p^f\})=P(\varnothing)=0$ (since X^uCp^t), $Bel(I)=P(S)=1$. From a reliability α for X we obtain the credibility α for p^t, however we do not obtain the credibility $1-\alpha$ for p^f, but 0, which means that we have no evidence against p.

Suppose that n sources, $X_1,\ldots,X_n$, give the same piece of information p. Since the temporal sequence of the testimonies is not significant, and if the n testimonies can be considered independent, we can adopt the Dempster Rule to combine the evidence. The n couples $<X_1,p>,\ldots,<X_n,p>$ yield a frame S which contains 2^n elements, the cartesian product of the n frames $\{X_1^r,X_1^u\},\ldots,\{X_n^r,X_n^u\}$. Since the testimonies are independent, the probability of the elements of S is given by the product of the probabilities of their elements. The notion of compatibility is extended from each source compatibility C_i to the *global* compatibility *SCI* defined as follows:

$\{X_1^{\delta_1},\ldots,X_n^{\delta_n}\}Ci \Leftrightarrow X_1^{\delta_1}C_1 i,\ldots,X_n^{\delta_n}C_n i$ (2

where $i\in\{p^t, p^f\}$ and either δ_i=r or δ_i=u. This global compatibility is then used in the formula (1) to calculate the credibility of the subsets of I. Let us consider now the more general case in which the n sources give n different pieces of information. The n couples $<X_1,p_1>,\ldots,<X_n,p_n>$ yield a frame I which contains also 2^n elements, i.e., the cartesian product of the n frames $\{p_1^t, p_1^f\},\ldots,\{p_n^t, p_n^f\}$. The formula (2) can be simply generalized:

$\{X_1^{\delta_1},\ldots,X_n^{\delta_n}\}C\{p_1^{\gamma_1},\ldots,p_n^{\gamma_n}\} \Leftrightarrow X_1^{\delta_1}C_1p_1^{\gamma_1},\ldots,X_n^{\delta_n}C_np_n^{\gamma_n}$ (3

where either δ_i=r or δ_i=u and either γ_i=t or γ_i=f. This global compatibility can again be used in the formula (1) to calculate the credibility of the subsets of I. From n pieces of information we obtain 2^{2^n} subsets of I!! Luckily, not all of them are interesting. If we want to assess only the credibility of the n sentences $p_1,\ldots,p_n$, we have to calculate only the credibility of the n subsets of I made of all the elements that contain, respectively, $p_1^t,\ldots,p_n^t$. Sometimes, we are also interested in the credibility of the falsity of $p_1,\ldots,p_n$; in this case, we have to calculate the credibility of the n subsets of I made of all the elements that contain, respectively, $p_1^f,\ldots,p_n^f$. In our belief revision framework, a very important role is played by the notion of "good", which is a maximally consistent subsets of $p_1,\ldots,p_n$; each good g corresponds to a particular element i of I: $p_i^t\in i \Leftrightarrow p_i\in g$ and $p_i^f\in i \Leftrightarrow p_i\notin g$. So, the belief-function formalism is able to attach a credibility to the goods too. However, as we'll see, this value of credibility is quite unsatisfactory.

Even if we are interested in a small number of subsets of I, we have still to face with the cardinality of I itself, which is 2^n. However, when there are absurdities or incompatibilities between the sentences $p_1,\ldots,p_n$, we can reduce the cardinality of I by eliminating those elements that contain impossible sets. For instance, if p_1 is an absurdum, then we delete from I all the elements that contain p_1^t (thus eliminating half the elements of I); if p_1 and p_2 are mutually exclusive and exaustive, then we delete from I all the elements that contain either $<p_1^t, p_2^t>$ or $<p_1^f, p_2^f>$ (still eliminating half the elements of I); if p_1 and p_2 are mutually exclusive but not exaustive, then we delete from I all the elements that contain $<p_1^t, p_2^t>$ (thus eliminating a quarter of the I's elements) and so on (there can be contradictions involving more than two sentences). Often, after these cancellations, it happens that some elements of S have no more compatible elements in I. For example, given the two couples $<X_1,p_1>$ and $<X_2,p_2>$ and the incompatibility $<p_1^t, p_2^t>$, the elements of S which are supersets of $\{X_1^r,X_2^r\}$ are compatible with no elements of I (X_1 and X_2 cannot both be reliable). Such elements will be cancealed from S, and the probability of the remaining elements will be subject to bayesian conditioning to restore the additivity property (the probability of the remaining elements in S have to sum up to 1). If $\{s_i,s_j,\ldots,s_k\}$ are the elements of S which are compatible with no elements in I, then the probability of any other element of S will be divided by $(1 - P(s_i)P(s_j)\ldots P(s_k))$. These new values of probability will then be used in the formula (1) to compute the degrees of credibility of the subsets of I we are interested in.

Example.

Input: <X, x>, $P(X^r)=0.9$

P	S	*C*	I	*Bel*
0.9	X^r		x^t	**0.9**
0.1	X^u		x^f	**0**

Input: <X, x>, <Y, x>, $P(X^r)=0.9$, $P(Y^r)=0.8$

P	X	*Cx*	I
0.9	X^r		x^t
0.1	X^u		x^f
1			

P	Y	*Cy*	I
0.8	Y^r		x^t
0.2	Y^u		x^f
1			

P	S	*C*	I	*Bel*
0.72	X^r,Y^r		x^t	**0.98**
0.18	X^r,Y^u		x^f	**0**
0.08	X^u,Y^r			
0.02	X^u,Y^u			
1				

Corroborating testimonies increase the credibility of the information ($Bel(x^t)=0.98$) while its negation is still without evidence ($Bel(x^f)=0$)
.

Input: <X, x>, <Y, y>, $P(X^r)=0.9$, $P(Y^r)=0.8$ Contradictions: $<x^t, y^t>$, $<x^f, y^f>$

P	X	*Cx*	x
0.9	X^r		x^t
0.1	X^u		x^f
1			

P	Y	*Cy*	y
0.8	Y^r		y^t
0.2	Y^u		y^f
1			

P	S	*C*	I	*Bel*
~~0.72~~	~~X^r,Y^r~~		~~x^t,y^t~~	
0.64	X^r,Y^u		x^t,y^f	**0.64**
0.29	X^u,Y^r		x^f,y^t	**0.29**
0.07	X^u,Y^u		~~x^f,y^f~~	
1				

The conflicts among the testimonies reduce the credibility of x^t ($Bel(x^t)=0.64$) but increase the credibility of x^f ($Bel(x^f)=0.29$) because now there is an evidence for it

3.1 Heuristics and Hints

Contradictions help us because they reduce the size of the information frame *I* so that, in order to evaluate the credibility of the sentences, we have to manage smaller subsets of *I*. However, to canceal the *I*'s elements which contain a contradiction, we must first generate them (generate and test for superset). This imposes severe limits to the number of sentences acceptable in input (about 30).

From an input-output analisys of the mechanism, we realized that two properties hold.

1. Sentences not involved in (minimal) contradictions and received from a single source, do not contribute to the mechanism; the degrees of credibility of the other sentences do not depend on their presence. Their degree of negative credibility is zero ($Bel(p^f)=0$; there is no evidence against them). Their degree of positive credibility ($Bel(p^t)$) is the minimum value among the degrees of positive credibility assumed by the other sentences (that were involved in contradictions) received from the same source and the degree of reliability of the source itself.
2. In assessing the degrees of positive credibility of the sentences, multiple (minimal) contradictions involving sentences received exclusively from exactly the same sources are redundant; all the sentences from the same source receive the same positive degree of credibility, independently of the number and the cardinality of the contradictions.

Property (1) implies that we can temporarly leave out of the process those sentences received from a single source that are not involved in contradictions. In many cases, this dramatically reduce the size of *I*. 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 temporarly 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 other.

The belief-function formalism is more powerful than what requested to fit in our belief revision framework. It provides not only a credibility value for each sentence, but also a negative credibility value, which is a parameter that represents the evidence against the sentence. While the notion of "nogood" (minimal set of sentences which cannot all be true) has a counterpart in the belief-function formalism, contradictions involving negated sentences (es. "*p*, *q* and *r* cannot all be false") have not counterpart in our framework. However, if we are interested only in the positive credibility values of the sentences, then only contradictions involving positive sentences (i.e., nogoods) are needed.

A "good" *g* corresponds to an element *i* of *I*, 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 the step 4 in our framework. From the formula (1), this degree of credibility is the sum of the degrees of reliability of the elements of *S* which are compatible only with *i*. But actually, it holds the following result.

Proposition. If $i \in I$ is a good, then there is only one element $s \in S$ (if any) which is compatible *only* with *i*: the one in which all the sources that gave information which are true in *i* are reliable in *s*, and all the sources that gave information which are false in *i* are unreliable in *s*.

Proof. If $s \in S$ is such that all the sources that gave information which are true in *i* are reliable and all the sources that gave information which are false in *i* are unreliable, then from the formula (3) it follows trivially that *sCi*. Furthermore, if *i* is a good, then it cannot exists another $i' \neq i$ such that *sCi'*. In fact, if there are sentences true in *i* and false in *i'*, then *s* is not compatible with *i'* since it is impossible that a reliable source gave false information. But, if there are sentences false in *i* and true in *i'*, then *i* would not be a good. Hence *s* is compatible only with *i*. On the other hand, suppose that there is another $s' \in S$ which is compatible only with *i*. If there is a source X which is reliable in *s* and unreliable in *s*', then *s*' is compatible not only with *i* but also with any *i*' in which at least one of the sentences from X is false. But if there is a source X which is unreliable in *s* and reliable in *s*', then *s*' is not compatible with *i* because the sentences from X are false in *i*. Hence there does not exist another $s' \in S$ compatible only with *i*. ■

Corollary. Let B be a set of sentences received exclusively from a same source X, and let $i \in I$ be a good. If there exist two partitions of B, B' and B", such that only the sentences of B' belong to *i*, then there does not exist a $s \in S$ which is compatible *only* with *i*.

Proof. From the previous proposition, X should be considered reliable and unreliable in the same $s \in S$. ■

If there are no elements of *S* which are compatible only with the good, its credibility is set at zero. This event is all but infrequent. So, the belief-function formalism is a bad estimator of the goods' credibility since it does not discriminate sufficiently among them. It is not infrequent the case that the credibility of all of them goes to zero. It attaches a non-null credibility only to goods supported by sources that have never been contradictory regarding any information provided. Besides, it is unreasonable to put at zero the credibility of a good only because it contains pieces of information received from a source that gave other pieces of information that we do not believe (are considered as false in the good).

Our solution is that of taking as credibility of the good, the average of the degrees of credibility, respectively positive and negative, of all the sentences inside and outside of the good. This estimator discriminates finely the goods and never goes to zero. The top good is the same obtained calculating directly the credibility of the goods (when there is at least one with a non-null credibility).

3.2 Changing the reliability of the sources

The main limit of this application of the belief-function formalism is that it takes the reliability of the sources as input values! These values will affect deeply the credibility of the sentences, so we need to be confident in them. A Salomonic solution to the problem of imposing such values is that of giving all the sources the same reliability. We are currently studing the possibility of taking advantage from the presented mechanism to reasonably change the reliability of the sources after any informative events, thus simulating some forms of learning (about the sources' reliability) from experience.

4 CONCLUSIONS

We have presented a method that tries to integrate symbolic and numerical formalisms for belief revision toward more practical and useful results. This method comes from researches in multi-agent domains, in which agents need to assess the credibility of information coming from different sources. It disconceives the principle of "priority to the incoming information" and follows the principle of "recoverability": an anytime received piece of information *p* must belong to the current knowledge space B whenever *p* is consistent with B. Pieces of knowledge may not only be abandoned ("non-monotonicity" of belief revision) but also rescued ("recoverability" of belief revision) after a newcoming information. To allow this, it is necessary to maintain two knowledge repositiories; the knowledge backgroung KB, which is the (inconsistent) collection of all the pieces of knowledge received/available, and the knowledge base B, which is a maximal consistent and preferred subset of KB. The selection of B among the many possible maximal consistent subsets of KB will be performed on the basis of a numerical credibility ordering on the sentences in KB. Besides recoverability, by doing so we overcome many limitations of other classic ways to belief revision, 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; for instance, the incoming information could be rejected even if it is consistent with the current knowledge base
- the complete numerical ordering renders the revision as least drastic as possible

Furthermore, 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. Numerical revision is performed in the belief-functon framework. We gave suggestions to limit the impact of the computational complexity of this formalism.

REFERENCES

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

[2] P. Gärdenfors, Belief Revision, Cambridge University Press, (1992)

[3] W. Nebel, Base Revision Operations and Schemes: Semantics, Representation, and Complexity, in Cohn A.G. (eds.), Proc. of the 11th European Conference on Artificial Intelligence, John Wiley & Sons, (1994)

[4] A.F. Dragoni and P. Puliti, Distributed Belief Revision versus Distributed Truth Maintenance, in Proc. 6th IEEE Conf. on Tools with A.I., IEEE Computer Press (1994).

[5] A.F. Dragoni and M. Di Manzo, Supporting Complex Inquiries, «International Journal of Intelligent Systems», vol 10, 12, (1995).

[6] Shafer G. and Srivastava R., The Bayesian and Belief-Function Formalisms a General Perpsective for Auditing, in G. Shafer and J. Pearl (eds.), Readings in Uncertain Reasoning, Morgan Kaufmann Publishers (1990.

[7] R. Reiter, A Theory of Diagnosis from First Principles, in «Artificial Intelligence», 53, (1987).

[8] Benferhat S., Cayrol C., Dubois D., Lang J. and Prade H. (1993), Inconsistency Management and Prioritized Syntax-Based Entailment, in Proc. of the 13th Inter. Joint Conf. on Artificial Intelligence, pp. 640-645.

[9] de Kleer J. (1986), An Assumption Based Truth Maintenance System, in «Artificial Intelligence», 28, pp. 127-162.

[10] Dragoni A.F. (1995), A consistency-Based Approach to Belief Revision, Tech. Rep, Istinfo 29/9/95.

[11] Dubois D. and Prade H. (1994), A Survey of Belief Revision and Update Rules in Various Uncertainty Models, in «International Journal of Intelligent Systems», 9, pp. 61-100.