Reasoning in inconsistent stratified knowledge bases

Salem BENFERHAT – Didier DUBOIS – Henri PRADE

Institut de Recherche en Informatique de Toulouse (I.R.I.T.), Université Paul Sabatier.
118 route de Narbonne – 31062 TOULOUSE – FRANCE
email: {Benferha, Dubois, Prade}@irit.fr

Abstract: This paper proposes a discussion of inconsistency-tolerant consequence relations in prioritized knowledge bases. These inference techniques extend methods for reasoning from inconsistent, non-stratified, knowledge bases to the case where priorities between formulas are available. Priorities between formulas are handled in the framework of possibility theory and allow for the use of pieces of information having various levels of confidence. A comparative analysis of several approaches is carried out, namely, the possibilistic inference and its extensions, three inference methods based on a selection of maximal consistent subsets of formulas, and two inference methods based on arguments.

Key words: inconsistency, argumentation, uncertainty, possibility theory.

1 - Introduction

The notion of priority is very important in the study of knowledge-based systems (Fagin et al., 1983). In practice, knowledge bases are sometimes inconsistent, due to the presence of rules having exceptions, or because the available knowledge comes from several, not necessarily agreeing sources. When priorities attached to pieces of knowledge are available, the task of coping with inconsistency is greatly simplified, since conflicts have a better chance to be solved. Gärdenfors (1988) has thus proved that any revision process of a belief set at the arrival of a new piece of propositional information, which satisfies natural requirements, is implicitly based on a priority ordering. Then the handling of priorities has been shown to be completely in agreement with possibilistic logic (Dubois and Prade, 1991; Benferhat et al., 1992). One way of tackling inconsistency is to revise the knowledge base and restore consistency. However, as pointed out in (Benferhat et al. 1995), in the case of multiple sources of information, it does not always make sense to revise an inconsistent knowledge base since it comes down to destroying part of the knowledge. In the context of merging several knowledge bases, the introduction of priorities between pieces of information can be explained by the two following situations:

- Each partial knowledge base issued from a source of information is "flat" (i.e., without any priority between items). But there is a total pre-ordering between the sources of information according to their reliability. In this case, merging different sources of information leads to a prioritized knowledge base, where the certainty level of each formula reflects the reliability of the source which supplies it. A particular case is when each piece of information is supported by a different source.
- All sources of information are equally reliable, but inside each partial knowledge base there exists a preference relation between pieces of information given by an expert, who rank-orders them according to their level of certainty. Here again, the combination of the different sources of information gives a prioritized knowledge base, provided that the scales of uncertainty used in each knowledge base are commensurate.

In (Benferhat et al., 1996), we have developed methods for inferring non-trivial conclusions from inconsistent *flat* sets of sentences in propositional logic, here called belief bases. These approaches fall in two categories: coherence-based methods use a selection of consistent subsets and apply classical inference to them; argument-based approaches exploit the belief base as it is and compute reasons to believe and to disbelieve propositions of interest. This paper extends inconsistency-tolerant inference methods from the flat case to the prioritized case and investigates other methods, specific to prioritized knowledge bases, for coping with inconsistency. Section 2 gives the background needed for reading this paper, and recalls the main results on the treatment of inconsistency for flat knowledge bases, developed in (Benferhat et al., 1996). Section 3 compares several coherence-based consequence relations that are inconsistency-tolerant, including several ones that come from the nonmonotonic logic literature. Two argument-based methods are presented in Section 4: an extension of the argumentation inference, originally proposed in (Benferhat et al., 1993b), and another approach which exploits a local treatment of inconsistency levels.

2 - Background

In this paper, we only consider a finite propositional language denoted by $\mathcal{L}$. We denote by $\vdash$ the classical consequence relation, Greek letters α, β,… represent formulas. Let Σ be a multiset of formulas (the same formula may be present several times), possibly inconsistent but not deductively closed. Cn(Σ) denotes the deductive closure of Σ. When the multiset Σ is not deductively closed, we call it a "belief base", following Nebel (1991). A formula in Σ is called a "belief". Each piece of information, or belief can be seen as supplied by some expert and is self-justifying. Beliefs are distinguished from plausible conclusions which are derived from beliefs.

Definition 1: A sub-base A of Σ is said to be *minimal inconsistent* (mI-subset, for short) iff: i) $A \vdash \bot$, and ii) $\forall \phi \in A$, $A - \{\phi\} \nvdash \bot$. A subbase A is said to be *maximal consistent* (MC-subset, for short) iff i) it is consistent and ii) either $A = \Sigma$ or adding any formula ϕ from $\Sigma - A$ to A entails the inconsistency of $A \cup \{\phi\}$.

Definition 2: A formula ϕ is said to be *free* iff it does not belong to any mI-sub-base of Σ. The subset of free formulas in Σ is called the free base of Σ, denoted by Free(Σ) (Benferhat et al., 1992).

We now recall two inconsistency-tolerant inferences in flat belief bases useful in the next sections:

Definition 3: A formula ϕ is said to be a *free consequence* of Σ, denoted by $\Sigma \vdash_{\text{Free}} \phi$, iff ϕ is logically entailed from Free(Σ), namely:

$$\Sigma \vdash_{\text{Free}} \phi \qquad \text{iff} \qquad \text{Free}(\Sigma) \vdash \phi.$$

Definition 4: A formula ϕ is said to be an *existential consequence* of Σ, denoted by $\Sigma \vdash_{\exists} \phi$, iff there exists at least one maximal consistent sub-base which entails ϕ. A formula ϕ is said to be a *universal consequence* (or *MC-consequence*) of Σ, denoted by $\Sigma \vdash_{MC} \phi$, iff ϕ can be classically inferred from all maximal consistent subsets of Σ.

When the belief base is flat, i.e., made of equally reliable propositional beliefs, the MC-consequence, a proposal made by Rescher and Manor (1970) (see also Rescher (1976)), is often used. However this consequence relation is also very conservative hence rather unproductive while the existential one is too permissive and leads to pairs of mutually exclusive conclusions. A mild inference approach, called *argumentation*, is proposed in (Benferhat et al., 1993b, 1996), that is more productive than the universal consequence but does not lead to conclusions which are pairwise contradictory. In this approach a conclusion can be inferred from an inconsistent belief base if the latter contains a consistent sub-base that entails this conclusion, but no consistent sub-base that entails its negation. In this paper, we extend them and propose new ones so as to deal with prioritized belief bases.

3 - Inconsistency-Tolerant Inference Relations: Coherence Approaches

Prioritized belief bases are layered, that is, of the form $\Sigma = S_1 \cup \ldots \cup S_n$, such that beliefs in S_i have the same level of priority or certainty and are more reliable than the ones in S_j where $j > i$. Each S_i is a multiset. Moreover copies of the same formula may appear in several layers, and are considered as distinct. The notation $\phi \in \Sigma$ means that there is a copy of ϕ in Σ. We use the notation $\phi \in_i \Sigma$ for $\phi \in S_i$, the ith layer of Σ. Throughout this paper, subbases are denoted by capital letters A, B,… and they are also represented in a stratified way, namely $A = A_1 \cup \ldots \cup A_n$ where $\forall$ j = 1,n, $A_j \subseteq S_j$ (A_j may be empty). In this section, we study several modes of inference based on the selection of consistent prioritized subbases.

3.1 - The Possibilistic Consequence Relation and Its Extensions

Issued from Zadeh (1978)'s possibility theory, possibilistic logic is a weighted logic where each classical logic formula is associated with a weight understood as a lower bound of a necessity measure; this weight accounts for the level of certainty (or the priority) of the information represented by the logical formula. A necessity measure N is a mapping from the language to the unit interval [0,1], such that $N(\phi \wedge \psi) = \min(N(\phi),N(\psi))$. A possibility measure Π is defined, by duality as $\Pi(\phi) = 1 - N(\neg\phi)$. The stratification of Σ is modelled in possibilistic logic by attaching a weight to each formula with the convention that $(\phi\ a_i)$ encodes $\phi \in_i \Sigma$, $\forall i$ and $1 \geq a_1 > \ldots > a_n > 0$. In this section, we only recall how possibilistic logic addresses the problem of inconsistency. See (Dubois et al., 1994b) for a complete exposition of possibilistic logic. A possibilistic belief base is thus a multiset of weighted sentences $\{(\psi\ a_j), \psi \in S_j, j = 1,n\}$, and all sentences in a layer S_i receive the same weight. The treatment of inconsistency is based on an entailment dedicated to layered belief bases, named the *π-entailment* and defined in the following way:

Definition 5: Let $\Sigma = S_1 \cup \ldots \cup S_n$ be a layered belief base. A formula ϕ is said to be an *i-consequence* of Σ, denoted by $\Sigma \vdash_i \phi$, iff: i) $S_1 \cup \ldots \cup S_i$ is consistent, ii) $S_1 \cup \ldots \cup S_i \vdash \phi$, and (iii) $\forall\, j < i$, $S_1 \cup \ldots \cup S_j \nvdash \phi$.

The definition of $\vdash_i$ is identical to the one proposed in possibilistic logic, namely $\vdash_i$ comes down to deduce $(\phi\ a_i)$ from the weighted formulae $\{(\psi\ a_j), \psi \in S_j, j = 1,i\}$ using an extended modus ponens rule whereby a conclusion inherits the smallest weight of the premises (Rescher, 1976), in agreement with the necessity measure-based semantics (Dubois et al., 1994b). From now on, we simply say that ϕ is *a π-consequence* of Σ if ϕ is an i-consequence of Σ for some i, and denote it by $\Sigma \vdash_\pi \phi$. It is clear that in the presence of inconsistency the π-entailment and the classical entailment have not the same behaviour. Indeed in classical logic if Σ is inconsistent then any formula can be deduced from Σ. In a stratified base, the situation is better since it is possible to use a unique consistent subbase of Σ, denoted by $\pi(\Sigma)$, induced by the levels of priority and defined in this way: $\pi(\Sigma) = S_1 \cup \ldots \cup S_i$ such that $S_1 \cup \ldots \cup S_i$ is consistent and $S_1 \cup \ldots \cup S_{i+1}$ is inconsistent. If Σ is consistent then $\pi(\Sigma) = \Sigma$. The level a_{i+1} is called the *inconsistency degree* $Inc(\Sigma)$ of Σ. The remaining subbase $\Sigma - \pi(\Sigma)$ is simply inhibited. It is not hard to check that:

$$\Sigma \vdash_\pi \phi \qquad \text{iff} \qquad \pi(\Sigma) \vdash \phi.$$

However, this way of dealing with inconsistency is not entirely satisfactory, since it suffers from an important drawback named "drowning problem" in (Benferhat et al., 1993a), as we can see in the following example: let $\Sigma = \{\{\neg\alpha \vee \neg\beta\}, \{\alpha\}, \{\beta\}, \{\delta\}\}$. This stratified base is inconsistent, and only the subset $\pi(\Sigma) = \{\{\neg\alpha \vee \neg\beta\}, \{\alpha\}\}$ is kept, and therefore δ cannot be deduced despite the fact that $\delta \in Free(\Sigma)$ is not involved in the inconsistency. One way of solving the drowning problem is to recover the inhibited free formulas, and to enlarge the selected consistent subset of Σ from $\pi(\Sigma)$ to $\pi(\Sigma) \cup Free(\Sigma)$. Then the drowning problem is solved, if the consequences of $\pi(\Sigma) \cup Free(\Sigma)$ are derived. However, $Free(\Sigma)$ is computed without taking the stratification into account. A natural extension of the free base to the stratified case is the *dominant subbase* defined as follows:

Definition 6: The dominant subbase of Σ is $\Sigma^* = \text{Free}(S_1) \cup \text{Free}(S_1 \cup S_2) \cup \ldots \cup \text{Free}(S_1 \cup \ldots \cup S_n)$.

Proposition 1: The dominant subbase Σ^* of a belief base Σ is consistent. Moreover $\Sigma^* \supseteq \pi(\Sigma) \cup \text{Free}(\Sigma)$.[1]

Note that for a given $i \geq 1$, if a formula ϕ does not belong to $\text{Free}(S_1 \cup \ldots \cup S_i)$ then it will not belong to $\text{Free}(S_1 \cup \ldots \cup S_k)$ for $k \geq i$. And in general, there is no inclusion relation between $\text{Free}(S_1 \cup \ldots \cup S_i)$ and $\text{Free}(S_1 \cup \ldots \cup S_k)$ where $k > i$. Σ^* is the set of non-defeated beliefs in Σ, in the sense that the dominant subbase of a belief base Σ is made of all formulas not contained in the least priority layer of any mI-subset of Σ, namely:

Proposition 2: For any mI-subset A let $\underline{A} = A - \pi(A)$ be the least priority layer of A. Then $\Sigma^* = \{\phi, \nexists \text{mI-subset } A \text{ of } \Sigma, \phi \in \underline{A}\}$.

This result shows that Σ^* is the largest subset of Σ containing undisputed beliefs. Clearly the *non-defeated* inference from Σ, defined as $\Sigma \vdash_{ND} \phi$ iff $\Sigma^* \vdash \phi$ is more productive than the possibilistic inference. Σ^* differs from the consistent subbase $\ell(\Sigma)$ obtained by dropping layers S_i when they are inconsistent with the previous ones. Namely,

$$\ell(\Sigma) = \ell(S_1 \cup \ldots \cup S_n).$$

$$i > 1: \quad \ell(S_1 \cup \ldots \cup S_i) = \begin{cases} \ell(S_1 \cup \ldots \cup S_{i-1}) \cup S_i & \text{if } \ell(S_1 \cup \ldots \cup S_{i-1}) \cup S_i \text{ is consistent} \\ \ell(S_1 \cup \ldots \cup S_{i-1}) & \text{otherwise} \end{cases}$$

$$\ell(S_1) = \begin{cases} S_1 & \text{if } S_1 \text{ is consistent} \\ \emptyset & \text{otherwise} \end{cases}$$

This idea, suggested by Dubois and Prade (1991) and more systematically studied by Nebel (1994), leads to define $\Sigma \vdash_{\ell} \phi$ iff $\ell(\Sigma) \vdash \phi$. The inference relation $\vdash_{\ell}$ does not solve the drowning problem, since if formulas in S_i cause an inconsistency, the whole subset S_i is inhibited, including free formulas therein. It is clear that $\vdash_{\ell}$ is more productive than the π-consequence. But however $\vdash_{\ell}$ remains incomparable with the ND-consequence.

3.2 - Inferences Based on Maximal Consistent Subsets

In the previous section, inconsistency-tolerant inference relies on a revision process that selects a suitable consistent prioritized subset of Σ, and infers from it. But this selection is

[1] Proofs of this paper and illustrative examples can be found in a longer paper entitled "A comparative study of some syntactic approaches to the handling of inconsistent knowledge bases: Part 2: The prioritized case", available as a Technical Report.

generally not a MC-subbase of Σ. The presence of priorities makes it possible to operate a selection among MC-subbases. A MC-subbase is *unacceptable* if it excludes high importance formulas, and it is acceptable otherwise. Rescher (1976) has proposed to rule out certain MC-subbases which are considered as unacceptable. More formally, for each maximal consistent subbase A, define a rank, denoted by BOMC(A) and read Best Out of the MC-subbase A, in the following way:

$$\mathrm{BOMC}(A) = \min\{i \,/\, \phi \in_i \Sigma \text{ and } \phi \notin A_i\}.$$

Then a sub-base A is said to be unacceptable iff there exists a maximal consistent subbase B≠A such that: BOMC(B) > BOMC(A), and A is said to be acceptable otherwise. A similar ranking on consistent (not necessarily maximal) subbases has been defined in (Benferhat et al., 1993b).

Brewka (1989) has introduced so-called "*preferred subbases*" which further restricts the class of acceptable maximal consistent subbases. Preferred subbases have also been introduced by (Dubois et al., 1992) in the setting of possibilistic logic under the name of *strongly maximal consistent subbases* (SMC-subbases for short). In the following the terminology "SMC-subbase" is adopted so as to emphasize the fact that the selection is made among maximal consistent subbases of Σ. A SMC-subbase A is constructed by starting with a MC-subbase of S_1, then we add to A as many formulas of S_2 as possible (while preserving consistency), and so on. Formally, a consistent subbase A is a *SMC-subbase* of Σ iff it is of the form $A = A_1 \cup \ldots \cup A_n$, where $\forall$ j = 1,n, $A_1 \cup \ldots \cup A_j$ is a MC-subbase of $S_1 \cup \ldots \cup S_j$. Let SMC(Σ) be the set of SMC-subbases of Σ, then:

Definition 7: A formula ϕ is said to be an *SMC-consequence* of Σ, denoted by $\Sigma \vdash_{SMC} \phi$, iff it is entailed from each element of SMC(Σ), namely: $\Sigma \vdash_{SMC} \phi$ iff $\forall$ A $\in$ SMC(Σ), $A \vdash \phi$.

As it turns out, the set of dominating formulas of Σ is included in the intersection of all SMC subbases:

Proposition 3: $\Sigma^* \subseteq \cap\{A \mid A \in \mathrm{SMC}(\Sigma)\}$ (the intersection is performed level by level). The converse is false.

One of the main drawbacks of the SMC-consequence relation is the cardinality of SMC(Σ) which increases exponentially with the number of conflicts in the belief base. One may think of selecting a non-empty subset of SMC(Σ), denoted by Lex(Σ), and defined as:

$$A = A_1 \cup \ldots \cup A_n \in \mathrm{Lex}(\Sigma) \text{ iff } \forall B = B_1 \cup \ldots \cup B_n \in \mathrm{SMC}(\Sigma),$$
$$\nexists i, \text{ s.t. } |B_i| > |A_i| \text{ and } \forall j<i, |B_j| = |A_j|$$

where |A| is the cardinality of A. A is said to be a lexicographically maximal consistent subset of Σ. The definition of Lex(Σ) is proposed in another form in (Dubois et al., 1992). Once Lex(Σ) is computed, we define a Lex-consequence relation in the following way: namely a formula ϕ is a *Lex-consequence* of Σ, denoted by $\Sigma \vdash_{Lex} \phi$, if and only if it is entailed from each element of Lex(Σ). It is obvious that each SMC-consequence of Σ is also a Lex-consequence of Σ. The converse is false since the Lex-consequence generally uses a strict subset of SMC(Σ).

4 - Arguments in Prioritized Belief Bases

The idea of argumentation is pervasive in artificial intelligence as testified by the literature on reason-maintenance systems initiated by Doyle (1979) and pursued by De Kleer (1986). It goes back to Toulmin (1956), and more recent proposals by Poole (1985), Pollock (1987), and Simari and Loui (1992) suggested argumentation as a key tool for the modelling of exception-tolerant defeasible reasoning. Fox et al. (1992) propose a typology of arguments, and insist that reasoning under uncertainty should exhibit reasons for uncertainty, and not only degrees of uncertainty. Argumentation is a natural method for handling inconsistent belief bases that treat inconsistency in local way. Two notions of consequence based on arguments are described in the following. One extracts and weights consistent subparts of a belief base viewed as arguments, and the other evaluates beliefs both in terms of certainty and inconsistency so as to distinguish between safe and unsafe arguments.

4.1 - Argued Consequences

We suggest that a conclusion can be inferred from an inconsistent belief base if the latter contains an argument that supports this conclusion with some certainty, but there is no argument that supports its negation with a higher or equal level of certainty. More formally, the argumentation approach is summarized in the two following definitions:

Definition 8: A subbase A of Σ is said to be an *argument* of rank i for a formula ϕ, if it satisfies the following conditions:

(i) $A \nvdash \bot$ (consistency), (ii) $A \vdash_i \phi$ (relevance), and

(iii) $\forall \psi \in A, A - \{\psi\} \nvdash_i \phi$ (economy)

An argument A which supports ϕ to a rank i is a *best* argument for ϕ if and only if each argument of rank j which supports ϕ is such that $i \leq j$. In this definition, the higher the rank of the argument, the weaker it is.

Definition 9: A formula ϕ is said to be an *argued consequence* of Σ, denoted by $\Sigma \vdash_{\mathscr{A}} \phi$, iff:

(i) there exists an argument of rank i for ϕ in Σ , and

(ii) arguments for $\neg\phi$ (if any) are of rank $j > i$.

The next proposition shows that $\vdash_{\mathscr{A}}$ is a faithful extension of the π-entailment.

Proposition 4: If $\Sigma \vdash_{\pi} \phi$ then $\Sigma \vdash_{\mathscr{A}} \phi$. The converse is false.

However, the argumentation inference is neither comparable to the SMC-consequence nor to the Lex-consequence.

4.2 - Safe Arguments

In the definition of the argumentation inference, there is no constraint on beliefs used to build arguments in favor of plausible conclusions of the belief base. Namely, one argument may for instance contain pieces of information which are directly involved in the inconsistency of the belief base. Levels of priority or of certainty attached to formulas have only been used to distinguish between strong and less strong arguments in favour of a proposition or of its contrary. However it is possible to go one step further in the use of the certainty or priority levels. The idea is to attach to any formula in the stratified belief base Σ two weights *a* and *b* reflecting respectively the extent to which there is some certainty that the formula is true and to what extent there is some certainty that the formula is false. More formally, to each $\phi \in \Sigma$, the triple (ϕ, a, b) is computed, where

i) The degree of certainty *a* is the (greatest) weight of a possibilistic logic proof of ϕ from a consistent subbase: $a = a_i$ if the best argument for ϕ is of rank i

ii) The degree of doubt *b* is the (greatest) weight of a possibilistic logic proof of $\neg\phi$ from a consistent subbase: $b = a_j$ if the best argument for $\neg\phi$ is of rank j

When $b = 0$, then ϕ is free since it means that no argument for its negation exists so that ϕ is not involved in the inconsistency of Σ. When $\min(a,b) \neq 0$, ϕ is said to be *paraconsistent* or *attacked* by some arguments. When $b \geq a$ then ϕ is said to be *defeated* by some argument.

Let $\Sigma_p = \{(\phi, a, b), \phi \in \Sigma\}$. Conclusions from the set Σ_p of biweighted formulas should be derived by propagation of the weights *a* and *b*. (Dubois et al., 1994a) already proposes an extension of the possibilistic resolution principle which allows for the propagation of paraconsistency degrees from premises to conclusions. In the following, the propagation of paraconsistency weights to a conclusion ψ is carried out directly from the set of arguments supporting it. To see if a formula ψ is a plausible consequence of Σ, first check if there is an argument in favour of ψ in Σ. It is clear that if there is no argument in favour of ψ in Σ then ψ cannot be a consequence of Σ. Assume that we have an argument A of ψ where all elements of A are free formulas, then ψ can safely be considered as a consequence of Σ. The situation differs if some elements of A are not free, and here we must be more careful. Indeed, let ϕ be a formula in A (the argument for ψ) and consider $(\phi, a, b) \in \Sigma_p$. If $b > a$ then there exists an argument in Σ which supports $\neg\phi$ more strongly than those supporting ϕ. The argument A is thus not reliable and cannot safely justify conclusion ψ. The idea is then to associate two degrees to arguments in Σ. St(A), called the *strength* of A, and Br(A), called the degree of *brittleness* of A. These two measures are computed in the following way:

$$\mathrm{St}(A) = \min\{a_i \mid (\phi_i\ a_i\ b_i) \in \Sigma_p \text{ and } \phi_i \in A\},$$
$$\mathrm{Br}(A) = \max\{b_i \mid (\phi_i\ a_i\ b_i) \in \Sigma_p \text{ and } \phi_i \in A\}.$$

In general, we can have several arguments which support ψ and we denote the set of all arguments with their pair of weights thus obtained for ψ by: $\mathrm{Label}(\psi) = \{(A_i, \mathrm{St}(A_i), \mathrm{Br}(A_i)) \mid A_i$ is an argument for $\psi\}$ From Label(ψ), we may think of two criteria to select the best argument for ψ: a certainty degree obtained from $\mathrm{St}(A_i)$ and a degree of doubt obtained from $\mathrm{Br}(A_i)$ and induced by the arguments A_i. It is clear that the best argument for ψ is the one which allows to deduce ψ with the highest certainty degree and the lowest degree of doubt. But in general, such an argument does not always exist. Then, we may: first either minimize the brittleness degree among the arguments, or first maximize the certainty degree. These two possibilities lead to two definitions of inconsistency-tolerant consequence relations. A third inference is to consider that ψ is a plausible consequence of Σ as soon as it possesses an argument A_i such that $\mathrm{St}(A_i) > \mathrm{Br}(A_i)$. More formally:

Definition 10:

i) Let Label$*(\psi)$ be the subset of Label(ψ) obtained by choosing the arguments with the lowest brittleness degree. Let A be an argument of Label$*(\psi)$ such that St(A) is maximal. Then ψ is said to be *BS-consequence* of Σ (BS: short for "first brittleness then strength"), denoted by $\Sigma \vdash_{BS} \psi$, iff St(A) > Br(A).

ii) Let Label* (ψ) be the subset of Label(ψ) obtained by choosing the arguments with the highest strength. Let A be an argument in Label$^*(\psi)$ such that A has the lowest brittleness degree. Then ψ is said to be a *SB-consequence* of Σ (SB: short for "first strength then brittleness"), denoted by $\Sigma \vdash_{SB} \psi$, iff St(A) > Br(A)

iii) ψ is said to be a *safely supported- consequence* of Σ (SS-Consequence for short), denoted $\Sigma \vdash_{SS} \psi$ iff there exists a triple $(A_i, \mathrm{St}(A_i), \mathrm{Br}(A_i))$ in Label(ψ) such that $\mathrm{St}(A_i) > \mathrm{Br}(A_i)$.

In order to assess the merits of $\vdash_{BS}$ and $\vdash_{SB}$ consider the following example where the knowledge base is:

$$\Sigma = \{(\phi\ a), (\neg\phi\ b), (\neg\phi \vee \psi\ c), (\chi\ d), (\neg\chi \vee \psi\ e), (\neg\chi\ f)\}.$$

Then: $\Sigma_p = \{(\phi\ a\ b), (\neg\phi\ b\ a), (\neg\phi \vee \psi\ \max(b,c)\ 0), (\chi\ d\ f), (\neg\chi \vee \psi\ \max(e,f)\ 0), (\neg\chi\ f\ d)\}$.

Notice that only the formulas $\neg\phi \vee \psi$ and $\neg\chi \vee \psi$ are free in Σ, all the others are paraconsistent. We are interested in knowing if ψ can be deduced from Σ. We have:

$$\mathrm{Label}(\psi) = \{(A, \min(a,c), b), (B, \min(d,e)\ f))\}$$

obtained using the following arguments respectively: $A = \{(\phi\ a), (\neg\phi \vee \psi\ c)\}$, and $B = \{(\chi\ d), (\neg\chi \vee \psi\ e)\}$. Then:

– assume that a > c > b > f > d > e. Then minimizing the degree of doubt of the conclusion ψ will lead to select B as the best argument for ψ, and since min(d,e) < f then ψ will not be BS-inferred. This result is somewhat debatable since using argument A, ψ is inferred from the two most certain formulas in Σ. In contrast, A is the strongest argument for ψ and hence ψ is a SB-consequence of Σ (since min(a, max(b,c)) > b). Notice from the example, that the BS-consequence relation does not recover all the possibilistic consequences of the knowledge base (e.g., here ψ is a possibilistic consequence of the knowledge base). In contrast, SB-consequences do recover all the possibilistic consequences of the knowledge base.

– assume that a > b > c > d > e > f then we obtain ψ as a plausible consequence of Σ if we first minimize the brittleness degree in Label(ψ). In contrast, if we first maximize the strength degree, A is selected since A is the best argument of ψ in this case, but ψ is no longer inferred since b > min(a, max(c,b)). This result is somewhat debatable especially if the certainty degree of $\neg\chi$ is very low.

The above example suggests that $\vdash_{BS}$ is not so interesting insofar as the possibilistic inference since some possibilistic consequences are not derived by this inference method. On the contrary, if $\Sigma \vdash_{\pi} \psi$ then $\Sigma \vdash_{SB} \psi$. Notice that in the two above cases of the example, ψ is a safely supported consequence of Σ. Indeed St(A) > Br(A) in the first case and St(B) > Br(B) in the second case. Actually it is obvious that generally if $\Sigma \vdash_{BS} \psi$ (resp. $\Sigma \vdash_{SB} \psi$) then $\Sigma \vdash_{SS} \psi$. Moreover, as suggested by the previous examples, we may have (A, a, b) $\in$ Label(ψ) with a > b (in this case ψ is a safely supported consequence), while ψ is neither a SB-consequence nor a BS-consequence of the knowledge base. Indeed, it is sufficient to imagine a situation where Label(ψ) = {(A, a + ε, 1), (B, a, b), (C, ε', b – ε'')} with $1 > a > b > b - \varepsilon'' \geq \varepsilon'$. The following result indicates that any set of safely supported consequences is consistent:

Proposition 5: Let $K = \{\psi \mid \Sigma \vdash_{SS} \psi\}$. Then the set of formulas K is consistent.

This result contrasts with the case of argumentation-based inference, where a set of three or more argued consequences can be inconsistent (Benferhat et al., 1993b). The following proposition positions the consequence relation $\vdash_{SS}$ with respect to the ones proposed in the previous sections.

Proposition 6: The set of non-defeated consequences of Σ is the deductive closure of $Cn_{SS}(\Sigma) = \{\psi \mid \Sigma \vdash_{SS} \psi\}$.

Finally, we can show that ND-consequence and the argued consequence of Σ are incomparable, and that the consequence relation $\vdash_{1}$ is incomparable with the SS-consequence, ND-consequence, SMC-consequence relations and also with the Lex-consequence relations.

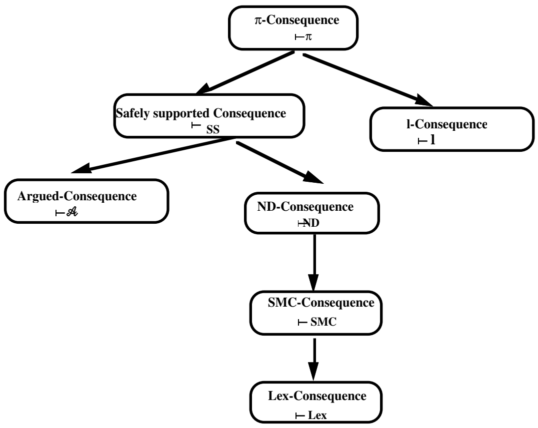

Figure 1: A comparative study of inference relations

Figure 1 summarized the links existing between the different consequence relations studied here, the edges mean the set inclusion relation between the set of results generated by each consequence relation. The top of the diagram thus corresponds to the most conservative inferences. Cayrol (1995) has shown that most existing argumentation systems come down to one of the consequence relations studied here.

5 - Conclusion

Argued and SS-consequence relations are appealing for reasoning from inconsistent belief bases storing multiple source information. Indeed, they do not reduce the belief base to one, or several, consistent subbases, but rather estimate pros and cons for a conclusion and evaluate their strength. They differ from the other consequence relations since the latter produce the classical consequences of one or several consistent subsets of the belief base. Moreover, a specific weight accounts for the level of local contradiction, attached to each formula handled by the SS-consequence relation. Only argued consequences look really faithful to the actual contents of the knowledge base, and do not give up information contained in it, as opposed to the approaches based on SMC and lexicographically preferred subbases. All approaches which are at least as productive as the SS-consequence partially avoid the drowning effect of standard possibilistic logic by salvaging sentences whose level of certainty is low but which are not involved in any contradiction set. However, approaches which are more adventurous than the SS-consequence relation can lead to undesirable conclusions. Indeed, take the following example $\Sigma = \{S_1 = \{\phi\}, S_2 = \{\neg\phi\}, S_3 = \{\neg\phi \vee \psi\}\}$, where ψ is a plausible consequence of Σ using $\vdash_{SMC}$, $\vdash_{Lex}$, $\vdash_{\mathscr{A}}$ while ψ is not a safely supported consequence of Σ. Indeed ψ is supported only by $\{\phi, \neg\phi \vee \psi\}$, an argument of rank 3, while $Br(\{\phi, \neg\phi \vee \psi\}) < 3$ since the paraconsistency of ϕ is stronger than the strength of the other component, $\neg\phi \vee \psi$, of the argument.

Another advantage of some methods is that they are amenable to efficient standard implementation methods based on classical resolution, especially the ND-inference which is the most productive one among those whose complexity is like the one of the possibilistic inference (once Σ^* is computed). The SMC and Lex-inferences look more computationally complex. All methods but one produce consistent sets of results from inconsistent belief bases. Argumentation inference avoids outright contradictory responses (such as ϕ and $\neg\phi$), although more than three deduced sentences can be globally inconsistent. But as pointed out earlier, the arguments supporting a set of more than two globally contradictory sentences are distinct, so that the reality of this contradiction is debatable, and only reflects the presence of different points of view. Anyway it seems that it is the price to pay in order to remain faithful to an inconsistent belief base. In a long version of this paper, we show how the argumentation inferences can be applied for merging several knowledge bases. Lastly, the idea of the approaches presented in this paper differs from the one used in "paraconsistent logics" (Da Costa, 1963). These logics even if they reject the "ex facto quolibet" principle (namely, $\phi \wedge \neg\phi \rightarrow \psi$ is not a theorem in these logics), are often monotone in the sense that they still infer contradictory set of conclusions from inconsistent belief base. This is not the case with most of the approaches developed in this paper since not only they reject the "ex facto quolibet" principle, but moreover they do not infer contradictory conclusions.

Acknowledgments: This work has been partially supported by the European ESPRIT Basic Research Action No. 6156 entitled "Defeasible Reasoning and Uncertainty Management Systems (DRUMS-II)".

References

Benferhat S., Dubois D., Prade H. (1992) Representing default rules in possibilistic logic. Proc. of the 3rd Inter. Conf. on Principles of Knowledge Representation and Reasoning (KR'92), Cambridge, MA, Oct. 26-29, 673-684.

Benferhat S., Cayrol C., Dubois D., Lang J., Prade H. (1993a) Inconsistency management and prioritized syntax-based entailment. Proc. of the 13th Inter. Joint Conf. on Artificial Intelligence (IJCAI'93), Chambéry, Aug. 28-Sept. 3, 640-645.

Benferhat S., Dubois D., Prade H. (1993b) Argumentative inference in uncertain and inconsistent knowledge bases. Proc. of the 9th Conf. on Uncertainty in Artificial Intelligence (D. Heckerman, A. Mamdani, eds.), 411-419.

Benferhat S., Dubois D., Prade H. (1995) How to infer from inconsistent beliefs without revising?. Proc. of the 14th Inter. Joint Conf. on Artificial Intelligence (IJCAI'95), Montréal, Canada, Aug. 20-25, 1449-1455.

Benferhat S., Dubois D., Prade H. (1996) Some syntactic approaches to the handling of inconsistent knowledge bases: a comparative study — Part 1: The flat case. To appear in Studia Logica Journal.

Brewka G. (1989) Preferred subtheories: an extended logical framework for default reasoning. Proc. of the 11th Inter. Joint Conf. on Artificial Intelligence (IJCAI'89), Detroit, MI, Aug. 20-25, 1043-1048.

Cayrol C. (1995) On the relation between argumentation and non-monotonic coherence-based entailment. Proc. of the 14th Inter. Joint Conf. on Artificial Intelligence (IJCAI'95), Montréal, Canada, Aug. 20-25, 1443-1448.

Da Costa N.C.A. (1963) Calculus propositionnels pour les systèmes formels inconsistants. Compte Rendu Acad. des Sciences (Paris), 257, 3790-3792.

De Kleer J. (1986) An assumption-based TMS. Artificial Intelligence, 28, 127-162.

Doyle J. (1979) A truth maintenance system. Artificial Intelligence, 12, 231-272.

Dubois D., Lang J., Prade H. (1994a) Handling uncertainty, context, vague predicates, and partial inconsistency in possibilistic logic. In: Fuzzy Logic and Fuzzy Control (Proc. of the IJCAI'91 Workshops on Fuzzy Logic and Fuzzy Control, Aug. 1991) (D. Driankov, P.W. Eklund, A.L. Ralescu, eds.), Springer-Verlag, Lecture Notes in Artificial Intelligence, Vol. 833, 45-55.

Dubois D., Lang J., Prade H. (1992) Inconsistency in possibilistic knowledge bases — To live or not live with it. In: Fuzzy Logic for the Management of Uncertainty (L.A. Zadeh, J. Kacprzyk, eds.), Wiley, New York, 335-351.

Dubois D., Lang J., Prade H. (1994b) Possibilistic logic. In: Handbook of Logic in Artificial Intelligence and Logic Programming, Vol. 3 (D.M. Gabbay et al., eds.), Oxford Univ. Press, 439-513.

Dubois D., Prade H. (1991) Epistemic entrenchment and possibilistic logic. Artificial Intelligence, 50, 223-239.

Fagin R., Ullman J.D., Vardi M.Y. (1983) On the semantics of updates in database. Proc. of the 2nd ACM SIGACT-SIGMOD Symp. on the Principles of Databases Systems, Atlanta, 352-365.

Fox J., Krause P., Ambler S. (1992) Arguments, contradictions and practical reasoning. Proc. of the 10th Europ. Conf. on Artificial Intelligence (ECAI'92) (B. Neumann, ed.), Vienna, Austria, Aug. 3-7, 623-627.

Gärdenfors P. (1988) Knowledge in Flux — Modeling the Dynamic of Epistemic States. The MIT Press, Cambridge.

Nebel B. (1991) Belief revision and default reasoning: Syntax-based approaches. Proc. of the 2nd Inter. Conf. on Principles of Knowledge Representation and Reasoning (KR'91) (J. Allen, R. Fikes, E. Sandewall, eds.), Morgan Kaufmann, Cambridge, MA, April 22-25, 417-428.

Nebel B. (1994) Base revision operator and schemes: semantics representation and complexity. Proc. of the 11th Europ. Conf. on Artificial Intelligence (ECAI'94) (A.G. Cohn, ed.), Amsterdam, The Netherlands, Aug. 8-12, 1994, Wiley, New York, 341-345.

Pollock J.L. (1987) Defeasible reasoning. Cognitive Science, 11, 481-518.

Poole D. (1985) On the comparison of theories: preferring the most specific explanation. Proc. 9th Inter. Joint Conf. Artificial Intelligence (IJCAI'85), Los Angeles, CA, August, 144-147.

Rescher N. (1976), Plausible Reasoning: An Introduction to the Theory and Practice of Plausibilistic Inference. Van Gorcum, Amsterdam.

Rescher N., Manor R. (1970) On inference from inconsistent premises. Theory and Decision, 1, 179-219.

Simari G.R., Loui R.P. (1992) A mathematical treatment of defeasible reasoning and its implementation. Artificial Intelligence, 53, 125-157.

Toulmin S. (1956) The Uses of Argument. Cambridge Univ. Press, Cambridge, UK.

Zadeh L.A. (1978) Fuzzy sets as a basis for a theory of possibility. Fuzzy Sets and Systems, 1, 3-28.