Inconsistency management and prioritized syntax-based entailment

Salem Benferhat, Claudette Cayrol, Didier Dubois, Jerome Lang, Henri Prade
Institut de Recherche en Informatique de Toulouse (I.R.I.T.), University Paul Sabatier - C.N.R.S.
118 route de Narbonne, 31062 TOULOUSE Cedex (FRANCE)
E-mail: {benferha, testemal, dubois, lang, prade}@irit.irit.fr

Abstract

The idea of ordering plays a basic role in commonsense reasoning for addressing three inter-related tasks: inconsistency handling, belief revision and plausible inference. We study the behavior of non-monotonic inferences induced by various methods for priority-based handling of inconsistent sets of classical formulas. One of them is based on a lexicographic ordering of maximal consistent subsets, and refines Brewka's preferred sub-theories. This new approach leads to a non-monotonic inference which satisfies the "rationality" property while solving the problem of blocking of property inheritance. It differs from and improves previous equivalent approaches such as Gardenfors and Makinson's expectation-based inference, Pearl's System Z and possibilistic logic.

1 Introduction

It is noticeable, although very natural, that the notion of ordering (between logical formulas, between models, between subsets of formulas) has emerged from studies in nonmonotonic reasoning and belief revision as playing a crucial role. In both cases, the existence of such orderings is a direct consequence of a set of axioms which plausible inference, or revision processes, must obey. Makinson and Gardenfors [23] have pointed out that some nonmonotonic inference systems can be expressed in terms of the appropriate revision of a related set of propositional sentences, both inference and revision being guided by the ordering. Namely, given a set of sentences S and a revision procedure, given a formula a to be added to S, inferring β from a in the context defined by S can be achieved by checking whether β is a consequence of the result of revising S by a; the revision process takes care of the case when $S \cup \{\alpha\}$ is inconsistent. The ordering underlying these operations helps altogether coping with inconsistency, solving a revision problem, and guiding a nonmonotonic inference. In this paper we shall assume that a set of formulas is equipped with a complete preordering structure (or *priority ranking)* which, contrarily to Gardenfors' [13] view, is not related to the semantical entailment ordering between sentences. This kind of ordering has been considered by Brewka [4], Geffner [15], Nebel [24], Cayrol [5] in the recent past. Especially Nebel has used it to define syntax-based revision procedures, such that two semantically equivalent knowledge bases may, upon the arrival of some input inconsistent with them, result in non-semantically equivalent revisions. A first idea to revise an inconsistent knowledge base 8 is to select one of its maximal consistent subbases; another natural idea is to keep as many sentences as possible, i.e. consider a consistent subbase of S of maximum cardinality. The latter option helps reducing the number of revision candidates. The presence of an ordering on S leads to refine both approaches. Besides, one does not need to select a single preferred subbase when defining a non-trivial notion of inference from an inconsistent knowledge base. The task of this paper is precisely to study inferences of the form "S entails B if can be classically inferred in all the preferred consistent subbases of S". Here we shall focus on two meanings of "preferred": one, already considered by Brewka and Geffner, that combines priorities and maximal consistent subbases; another, which has not been studied in the literature, combines priorities and consistent subbases of maximum cardinality. Borrowing from Gardenfors the image that nonmonotonic reasoning and belief revision are two sides of the same coin, we pursue Nebel's work on syntax-based revision, by studying the other side of that particular coin; namely we study the properties of inference based on the two kinds of preferred subsets of formulas, and give algorithms for computing these inferences. Inconsistency-tolerant inferences are interesting since they can overcome some limitations of the kind of nonmonotonic inference that is the exact counterpart of Gardenfors' revision theory (see [14]). This kind of inference is also at work in possibilistic logic [10], [8] and in System Z [25]. This approach suffers from what we call the *"drowning effect":* all formulas which are not sufficiently entrenched are inhibited; this is an attenuated form of the fact that in classical logic anything follows from an inconsistent set of sentences; a particular case of this effect is the *property inheritance blocking* [25], [16].

2 Nonmonotonic Inference Relations Generated by a Flat Belief Base

Throughout this paper, $\mathcal{L}$ is a propositional language assumed finite for simplicity, since we deal with cardinalities. Formulas will be denoted by greek letters $\alpha, \beta, \ldots$; $\top$ (resp. $\bot$) denotes any tautology (resp. contradiction); Ω denotes the set of classical interpretations (denoted by $\omega, \omega', \ldots$) associated with $\mathcal{L}$; $\mathcal{M}(\alpha)$ denotes the set of models of α. A *belief base* S is a non-empty set of formulas $\{\varphi_1, \ldots, \varphi_n\}$ of $\mathcal{L}$. Subsets of S will be denoted by A, B,...; $|A|$ is the cardinality of A. Let $Cn(A)=\{\varphi \in \mathcal{L}, A \models \varphi\}$.

As discussed e.g. in [24], there are two ways of considering a set S of sentences that describes an epistemic state: the first one is to consider it logically ($\mathbf{S}$ is then called "belief set"), i.e. it can be represented equivalently by the set of its models and in this case it behaves as its closure under Cn. In this spirit, Gärdenfors and Makinson [14] generate nonmonotonic inference relations from sets of sentences which are consistent and closed under Cn; namely, β is inferred from α with respect to $\mathbf{S}$ iff β is derivable from α in all the "selected" maximal subsets of $\mathbf{S}$ consistent with α, where the selection function is to be defined. The second way is to consider the belief base "syntactically", i.e. each formula of $\mathbf{S}$ is a distinct piece of information; $\mathbf{S}$ is then more than the set of its models; for instance, $\{p,q\}$ is different from $\{p \wedge q\}$ since the former can be viewed as obtained from two experts, one asserting p and the other one q, while the latter has to be taken as a whole; in particular, in a revision context, p could be rejected or ignored independently from q in the former case, which is not true in the latter. Here, we assume neither that a belief base $\mathbf{S}$ is closed under Cn, nor that $\mathbf{S}$ is consistent. The inconsistency generally comes from the defeasible status of the pieces of information in $\mathbf{S}$; in particular, it may come from the combination of several independent (consistent) knowledge bases, as in [1]. The syntactic approach to nonmonotonicity or belief revision was used in different ways [11] [12] [24]. We first briefly deal with the non-prioritized case, for which inference can be based on maximal consistent subbases.

Definition 1: $A \subseteq \mathbf{S}$ is *inclusion-maximal α-consistent* if and only if $A \cup \{\alpha\}$ is consistent and there is no strict superset B of A with $B \subseteq \mathbf{S}$ such that $B \cup \{\alpha\}$ is consistent.

A natural way of selecting among these subbases A is to choose those of maximal cardinality, which leads to:

Definition 2: $A \subseteq \mathbf{S}$ is *cardinality-maximal α-consistent* if and only if $A \cup \{\alpha\}$ is consistent and for any $B \subseteq \mathbf{S}$ such that $B \cup \{\alpha\}$ is consistent, then $|B| \leq |A|$.

The obvious way of generating a nonmonotonic inference relation from $\mathbf{S}$ is then:

Definition 3: $\alpha \vdash_{\mathbf{S}} \beta$ (resp. $\alpha \vdash^{c}_{\mathbf{S}} \beta$) if and only if for every inclusion-maximal (resp. cardinality-maximal) α-consistent subset A of $\mathbf{S}$, we have $A \cup \{\alpha\} \vdash \beta$.

If $\alpha = \top$, we recover in the case of inclusion maximality Brewka's definition of β-strong provability from $\mathbf{S}$. Since any cardinality-maximal α-consistent subset of $\mathbf{S}$ is inclusion-maximal α-consistent, we have:

Proposition 1: $\alpha \vdash_{\mathbf{S}} \beta \Rightarrow \alpha \vdash^{c}_{\mathbf{S}} \beta$.

Thus, $\vdash^{c}_{\mathbf{S}}$ is more adventurous than $\vdash_{\mathbf{S}}$. Beyond complexity considerations, the choice of cardinality for discriminating among subbases is justified by adopting the point of view (well-known in model-based diagnosis) that $\mathbf{S}$ is a set of independently generated pieces of information which can "fail" to be true with a small probability ε; thus, the greater the number of rejected formulas, the smaller the probability that all of them fail (actually, ε^n if n formulas are rejected). This justifies cardinality-based preferences. However we do not claim that inferences based on cardinality are always the best way of defining an inference relation since it may be too adventurous in some cases. The choice of a selection mode may be application-dependent (e.g. the use of cardinality in diagnosis).

Both $\vdash_{\mathbf{S}}$ and $\vdash^{c}_{\mathbf{S}}$ fall in the class of consequence relations based on preferential models, in the sense of Shoham [29].

Definition 4: $\omega \leq_{\mathbf{S}} \omega'$ iff $\{\varphi_i \in \mathbf{S}, \omega \models \varphi_i\} \subseteq \{\varphi_i \in \mathbf{S}, \omega' \models \varphi_i\}$; $\omega \leq^{c}_{\mathbf{S}} \omega'$ iff $|\{\varphi_i \in \mathbf{S}, \omega \models \varphi_i\}| \leq |\{\varphi_i \in \mathbf{S}, \omega' \models \varphi_i\}|$.

It is clear that $\leq_{\mathbf{S}}$ is a partial pre-ordering, while $\leq^{c}_{\mathbf{S}}$ is a complete pre-ordering. In the following, we take the convention that given a preordering on Ω, *maximal* elements are preferred ones.

Proposition 2: ω is maximal w.r.t. $\leq_{\mathbf{S}}$ (resp. $\leq^{c}_{\mathbf{S}}$) iff $\{\varphi \in \mathbf{S}, \omega \models \varphi\}$ is an inclusion-maximal (resp. cardinality-maximal) consistent subset of $\mathbf{S}$.

This leads to the following result (see [5] for (i))

Proposition 3:

i) $\alpha \vdash_{\mathbf{S}} \beta$ iff $\forall \omega$ maximal w.r.t. $\leq_{\mathbf{S}}$ in $\mathcal{M}(\alpha)$, $\omega \models \beta$

ii) $\alpha \vdash^{c}_{\mathbf{S}} \beta$ iff $\forall \omega$ maximal w.r.t. $\leq^{c}_{\mathbf{S}}$ in $\mathcal{M}(\alpha)$, $\omega \models \beta$

See Sec. 4 for properties of $\vdash_{\mathbf{S}}$ and $\vdash^{c}_{\mathbf{S}}$ in a general setting. A shortcoming of reasoning with maximal consistent subsets is their usually large number (see [24] for complexity consideration). To reduce their number, one may select them w.r.t. cardinality, or attach priorities to the formulas (see Sec. 3 & 4) as proposed in [4] [15] [24] [1].

3 Prioritized belief bases

3.1 Introduction

Most of the time, when revising an inconsistent knowledge base, it is not reasonable to give all formulas the same importance (in a belief revision terminology, the formulas in a belief base are not all equal regarding to rejection); for instance, if some formulas are more certain than others we wish to reject the least certain first; formulas may also be ordered according to their arrival time in the belief base and then the oldest ones might be preferably rejected in order to restore consistency. More generally, we assume that formulas of $\mathbf{S}$ are ranked by a complete pre-ordering $\leq$, or equivalently, that $\mathbf{S}$ consists in a collection $(\mathbf{S}_1, \ldots, \mathbf{S}_n)$ of belief bases, where $\mathbf{S}_1$ contains the formulas of highest priority, and $\mathbf{S}_n$ those of lowest priority. We shall denote $<$ the associated strict relation and $\approx$ the associated equivalence. $\mathbf{S}$ will be called a *prioritized belief base*, and $\leq$ a *priority relation*. The idea of selecting a preferred subbase induced by the priorities is not a new one since it goes back at least to [27]. This idea has been also put forward in the framework of nonmonotonic reasoning with prioritized circumscription [22], hierarchic autoepistemic logic [20] and more specifically Brewka's preferred subtheories [4] which extend Poole [26]'s logical approach (corresponding to the case of only two priority levels). Preferred subtheories have also been further studied (and in different ways) in [15], [5], in possibilistic logic by [8], in knowledge base combination by [1] and extensively in the context of belief revision, in [24]. Boutilier [3] gives a characterization of Brewka's preferred subtheories in terms of conditional logics. Roos

[28] studies a particular case of Brewka's work. The selection of preferred subbases of $\mathbf{S}$ presupposes that the preordering $\le$ be extended to compare sets of formulas. We argue that there might be not a single way to do it and we give several possible choices. This problem has been considered in specific cases in [4], [16], [8] and from a more general point of view in [6].

3.2 From priorities on formulas to preferences between subbases

The general idea is to account for priorities in order to further select some maximal consistent subbases. In this spirit we now give some possible definitions for extending the priority relation $\le$ on $\mathbf{S}$ into a preference relation $\ll$ on $2^{\mathbf{S}}$. Let $A=A_1\cup\ldots\cup A_n$, $B=B_1\cup\ldots\cup B_n$ be two consistent subsets of $\mathbf{S}$ (where $A_i=A\cap\mathbf{S}_i$ and $B_i=B\cap\mathbf{S}_i$).

Definitions 5:

- *"best out" ordering*: let A be a consistent subset of $\mathbf{S}$ and $a(A)$ the highest priority of a formula of $\mathbf{S}$ which is not in A, i.e. $a(A)=\mathrm{Min}\{i,\exists\varphi\in\mathbf{S}_i\backslash A_i\}$ with the convention $\mathrm{Min}\emptyset=n+1$; see [8]. The best-out ordering relation writes: $A \ll^{bo}_{\mathbf{S}} B$ iff $a(A)\le a(B)$. This ordering depends only on the most prioritary layer where there is at least one formula from $\mathbf{S}$ missing in the considered subbase. $\ll^{bo}_{\mathbf{S}}$ is a complete preordering. Note that this ordering is very rough since it does not take into account formulas of lower priority (of level $\ge a(A)$); consequently, the induced preference is not very selective, since some of the preferred subsets may not be maximal for set-inclusion. Indeed, let k be the maximal index such that $\mathbf{S}_1\cup\ldots\mathbf{S}_k$ is consistent. Then $A=\mathbf{S}_1\cup\ldots\mathbf{S}_k$ is a $\ll^{bo}_{\mathbf{S}}$-maximal element as well as any consistent subbase of $\mathbf{S}$ containing A.
- *inclusion-based ordering*: this is the strict ordering defined by $A \ll^{inc}_{\mathbf{S}} B$ iff $\exists$ i such that $A_i\subset B_i$ and for any $j<i$, $A_j=B_j$ ($\subset$ denotes strict inclusion). An equivalent definition of this ordering has been proposed by Geffner [15]:

$$A \ll^{inc}_{\mathbf{S}} B \text{ iff } \forall\alpha\in A\backslash B, \exists\beta\in B\backslash A \text{ such that } \alpha<\beta.$$

To see it, it is enough to consider $\beta\in B_i\backslash A_i$ in the above definition. This equivalence does not hold anymore if the priorities define only a partial ordering [6]. The $\ll^{inc}_{\mathbf{S}}$-maximal consistent elements are of the form $A=A_1\cup\ldots\cup A_n$ such that i) A_1 is a maximal consistent subbase of $\mathbf{S}_1$; ii) $A_1\cup\ldots\cup A_i$ is a maximal consistent subbase of $\mathbf{S}_1\cup\ldots\cup\mathbf{S}_i$ for i=1,n.

Clearly, when only one layer is present in $\mathbf{S}$, we recover the maximal consistent subbases. The $\ll^{inc}_{\mathbf{S}}$-maximal consistent elements are called preferred subtheories by Brewka and strongly maximal consistent subbases in [8] with another characterization: if $\mathrm{Inc}(A)=\max\{k, A_1\cup\ldots\cup A_{k-1}$ is consistent$\}$, A is a $\ll^{inc}_{\mathbf{S}}$-maximal consistent element iff $\forall\alpha\in\mathbf{S}_i\backslash A_i$, $\mathrm{Inc}(A\cup\{\alpha\})=i$, for i=1,n. Note that $\ll^{inc}_{\mathbf{S}}$-maximal consistent elements are also maximal for $\ll^{bo}_{\mathbf{S}}$ (the converse is false). Lastly, $\ll^{inc}_{\mathbf{S}}$ is a strict partial ordering, which means that there may be several incomparable preferred subtheories. The dual ordering, exchanging $\forall$ and $\exists$ in Geffner's definition, has been studied in the scope of abduction in [6].
- *lexicographic ordering*: this is the complete preordering whose strict part is $A \ll^{lex}_{\mathbf{S}} B$ iff $\exists$i such that $|A_i|<|B_i|$ and for any $j<i$, $|A_j|=|B_j|$ and the equivalence part is $A \approx^{lex}_{\mathbf{S}} B$ iff for any i, $|A_i|=|B_i|$.

The complete preordering $\ll^{lex}_{\mathbf{S}}$ generalizes the cardinality-based selection of subbases of Section 2 to the prioritized case and refines the inclusion-based ordering. Any lexicographically maximal consistent subbase of $\mathbf{S}$ is strongly maximal consistent but the converse is false (as seen in Sec. 2; see an example in Sec. 4). This ordering has been proposed in [8] with a different formulation: namely, consider the increasingly ordered multiset $\mathrm{List}(\mathbf{S}\backslash A)$ of layer levels of formulas of $\mathbf{S}$ not present in A, then $A \ll^{lex}_{\mathbf{S}} B$ iff $\mathrm{List}(\mathbf{S}\backslash A)\le_L \mathrm{List}(\mathbf{S}\backslash B)$ where $\le_L$ is the lexicographic ordering. Similarly to Brewka's constructive definition for preferred subtheories, we can give an equivalent characterization for $\ll^{lex}_{\mathbf{S}}$-maximal consistent bases: they are of the form $A=A_1\cup\ldots\cup A_n$ such that i) A_1 is a cardinality-maximal consistent subbase of S_1; ii) $A_1\cup\ldots\cup A_i$ is a cardinality-maximal consistent subbase of $\mathbf{S}_1\cup\ldots\cup\mathbf{S}_i$ for i=1,n.

Definition 6: A subbase A of $\mathbf{S}$ is said to be $\ll^{bo}_{\mathbf{S}}$ (resp. $\ll^{inc}_{\mathbf{S}}$, $\ll^{lex}_{\mathbf{S}}$)-maximal. α-consistent iff $A\cup\{\alpha\}$ is consistent and A is maximal w.r.t. $\ll^{bo}_{\mathbf{S}}$ (resp. $\ll^{inc}_{\mathbf{S}}$, $\ll^{lex}_{\mathbf{S}}$) among the α-consistent subbases of $\mathbf{S}$.

3.3 Semantics

Now, these (pre-)orderings on $2^{\mathbf{S}}$ induce (pre-)orderings on Ω, which provide us with clearer semantics.

Definition 7: let $\omega, \omega' \in \Omega$, $[\omega]$ be the set of formulas satisfied by ω.

i) $\omega \preccurlyeq^{bo}_{\mathbf{S}} \omega'$ iff $a(\mathbf{S}\cap[\omega]) \le a(\mathbf{S}\cap[\omega'])$.

ii) $\omega \prec^{inc}_{\mathbf{S}} \omega'$ iff $\exists$i such that $\forall j<i$, $\mathbf{S}_j\cap[\omega]=\mathbf{S}_j\cap[\omega']$, and $\mathbf{S}_i\cap[\omega]\subset\mathbf{S}_i\cap[\omega']$; $\omega \approx^{inc}_{\mathbf{S}} \omega'$ iff $\forall$ i, $\mathbf{S}_i\cap[\omega] = \mathbf{S}_i\cap[\omega']$; $\omega \preccurlyeq^{inc}_{\mathbf{S}} \omega'$ iff $\omega \prec^{inc}_{\mathbf{S}} \omega'$ or $\omega \approx^{inc}_{\mathbf{S}} \omega'$.

iii) $\omega \prec^{lex}_{\mathbf{S}} \omega'$ iff $\exists$i such that $\forall j<i$, $|\mathbf{S}_j\cap[\omega]|=|\mathbf{S}_j\cap[\omega']|$ and $|\mathbf{S}_i\cap[\omega]|<|\mathbf{S}_i\cap[\omega']|$. $\omega \preccurlyeq^{lex}_{\mathbf{S}} \omega'$ means "not $\omega' \prec^{lex}_{\mathbf{S}} \omega$".

Proposition 4:

i) $\preccurlyeq^{bo}_{\mathbf{S}}$ and $\preccurlyeq^{lex}_{\mathbf{S}}$ are complete pre-orderings; $\preccurlyeq^{inc}_{\mathbf{S}}$ is a partial pre-ordering.

ii) $\omega \prec^{bo}_{\mathbf{S}} \omega' \Rightarrow \omega \prec^{inc}_{\mathbf{S}} \omega' \Rightarrow \omega \prec^{lex}_{\mathbf{S}} \omega'$.

4 Nonmonotonic Inference Relations Generated by a Prioritized Belief Base

We now generalize the inference relations defined in Sec. 2 to the prioritized case.

4.1 Inference relations generated by best-out preorderings: the drowning effect

Definition 8: let $\mathbb{S}=(\mathbb{S}_1,\ldots,\mathbb{S}_n)$ be a belief base and α be a formula. The addition of α to $\mathbb{S}$, denoted by $Add(\alpha,\mathbb{S})$, is the belief base obtained by adding to $\mathbb{S}$ a new level, at the highest priority, containing only α. More formally, $Add(\alpha,\mathbb{S})=(\mathbb{S}_0=\{\alpha\},\mathbb{S}_1,\ldots,\mathbb{S}_n)$.

Definition 9: $\alpha \vdash_{\mathbb{S}}^{bo} \beta$ iff for any $\leq_{\mathbb{S}}^{bo}$-maximal α-consistent subbase A of $\mathbb{S}$, $A\cup\{\alpha\}\vdash\beta$.

Similarly to what we noticed in 3.2, the set of maximal α-consistent elements for $\leq_{\mathbb{S}}^{bo}$ has a least element, namely $\mathbb{S}_1\cup\ldots\mathbb{S}_k$ is α-consistent and $\mathbb{S}_1\cup..\mathbb{S}_{k+1}$ is α-inconsistent. Hence $\alpha\vdash_{\mathbb{S}}^{bo}\beta \Leftrightarrow \{\alpha\}\cup\mathbb{S}_1\ldots\cup\mathbb{S}_k\vdash\beta$, and no sentence of lower priority (i.e. level>k) is involved in the inference. This is the drowning effect. The semantic counterpart is

Proposition 5: $\alpha \vdash_{\mathbb{S}}^{bo} \beta \Leftrightarrow \forall\omega\in\mathcal{M}(\alpha)$, such that $a([\omega]\cap\mathbb{S})$ is maximal, $\omega\in\mathcal{M}(\beta)$.

In the terminology of possibility theory, $a([\omega]\cap\mathbb{S})$ is a kind of degree of impossibility of interpretation ω, and $\alpha\vdash_{\mathbb{S}}^{bo}\beta$ corresponds to the non-trivial deduction of possibilistic logic [10], also called expectation-based non-monotonic inference in [14]. Namely let us define for any formula ϕ, $\ell(\phi)=\min\{k,\mathbb{S}_1\cup\ldots\mathbb{S}_k\vdash\phi\}$. Then ℓ is closely related to a necessity measure, i.e. $\ell(\phi\wedge\psi)=\max(\ell(\phi),\ell(\psi))$ in a scale of integers running opposite to the unit interval.

It is easy to check that the ordering defined by $\ell(\phi)\leq\ell(\psi)$ is an expectation ordering [14] and a qualitative necessity measure [9], and that it is induced from $\mathbb{S}$ by the most compact ranking of interpretations (following Pearl [25]) and by the principle of minimum specificity of possibilistic logic. The non-monotonic inference is usually given by

$$\alpha\vdash_{\mathbb{S}}^{bo}\beta\Leftrightarrow\alpha\vdash\beta \text{ or } \ell(\alpha\rightarrow\beta)<\ell(\alpha\rightarrow\neg\beta)$$

a definition given under various guises by Gärdenfors and Makinson, Pearl, and Dubois and Prade. The relationship between the drowning effect laid bare by definition 9 and the blocking of property inheritance in Pearl's System Z can be viewed in the following example [2].

Example 1: Let p, b, f, w respectively mean penguin, bird, fly, winged and consider the set of defaults $\{p\rightarrow\neg f, b\rightarrow f, p\rightarrow b, b\rightarrow w\}$. Using the default ordering procedure of System Z, the two defaults $p\rightarrow\neg f$ and $p\rightarrow b$ are granted higher priority because they correspond to a more specific reference class. In [2] it has been shown that it corresponds to the following layered classical belief base $\mathbb{S}=\mathbb{S}_1\cup\mathbb{S}_2$ where $\mathbb{S}_1=\{\neg p\vee\neg f,\neg p\vee b\}$ and $\mathbb{S}_2=\{\neg b\vee f,\neg b\vee w\}$ and the inference made possible by System Z is exactly $\vdash_{\mathbb{S}}^{bo}$ (expressed in the terminology of possibilistic logic in the above mentioned paper). It is easy to check that $b\vdash_{\mathbb{S}}^{bo} f$, $b\vdash_{\mathbb{S}}^{bo} w$, $p\vdash_{\mathbb{S}}^{bo}\neg f$ but because of the drowning effect, $p\vdash_{\mathbb{S}}^{bo} w$ does not hold due to the presence of a conflict in $\mathbb{S}_1\cup\mathbb{S}_2$ ($Inc(\mathbb{S})=2$). Hence penguins cannot inherit from birds the property of having wings. The inference relationships studied in the following will cope with this problem.

4.2 Inference relations generated by inclusion-based and lexicographic preorderings

Since the preference relations defined in Section 3 generalize the criteria of Section 2, based on inclusion and cardinality, the following definitions naturally extend definition 3.

Definition 10: $\alpha\vdash_{\mathbb{S}}^{inc}\beta$ iff for any $\ll_{\mathbb{S}}^{inc}$-maximal α-consistent subbase A of $\mathbb{S}$, $A\cup\{\alpha\}\vdash\beta$. This definition is also in [24] [5]. $\alpha\vdash_{\mathbb{S}}^{lex}\beta$ iff for any $\ll_{\mathbb{S}}^{lex}$-maximal α-consistent subbase A of $\mathbb{S}$, $A\cup\{\alpha\}\vdash\beta$.

Example 2: let $\mathbb{S}=(S_1=\{a,\neg a\}, S_2=\{\neg a\vee b, a\vee c\}, S_3=\{\neg b\}, S_4=\{\neg c\})$. Then the two subbases $A=\{a,\neg a\vee b,a\vee c,\neg c\}$ and $B=\{\neg a,\neg a\vee b,a\vee c,\neg b\}$ are both $\ll_{\mathbb{S}}^{inc}$-maximal consistent, but only B is $\ll_{\mathbb{S}}^{lex}$-maximal α-consistent. Thus, $\vdash_{\mathbb{S}}^{lex}\neg a$ holds but not $\vdash_{\mathbb{S}}^{inc}\neg a$.

Example 3: let us consider again the belief base of Section 4.1. Now, it is easy to check that $p\vdash_{\mathbb{S}}^{inc} w$ and $p\vdash_{\mathbb{S}}^{lex} w$; thus, $\vdash_{\mathbb{S}}^{inc}$ and $\vdash_{\mathbb{S}}^{lex}$ escape the drowning problem.

The next result, generalizing Prop. 3, states that $\vdash_{\mathbb{S}}^{inc}$ and $\vdash_{\mathbb{S}}^{lex}$ inferences can be defined equivalently by means of preferential models à la Shoham (see [5] for $\vdash_{\mathbb{S}}^{inc}$):

Proposition 6: let $\alpha\neq\perp$; ω is maximal in $\mathcal{M}(\alpha)$ w.r.t. $\leq_{\mathbb{S}}^{inc}$ (resp. $\leq_{\mathbb{S}}^{lex}$) iff $[\omega]\cap\mathbb{S}$ is maximal w.r.t. $\ll_{\mathbb{S}}^{inc}$ (resp. $\ll_{\mathbb{S}}^{lex}$) among the consistent subsets of $Add(\alpha,\mathbb{S})$.

Proposition 7: $\alpha\vdash_{\mathbb{S}}^{inc}\beta$ (resp. $\alpha\vdash_{\mathbb{S}}^{lex}\beta$) iff for any $\omega\leq_{\mathbb{S}}^{inc}$-maximal (resp. $\leq_{\mathbb{S}}^{lex}$-maximal) in $\mathcal{M}(\alpha)$, we have $\omega\in\mathcal{M}(\beta)$.

Sketch of proof: ($\Rightarrow$) follows from Proposition 4 and the definitions of $\alpha\vdash_{\mathbb{S}}^{inc}\beta$, $\alpha\vdash_{\mathbb{S}}^{lex}\beta$.

($\Leftarrow$) let A be a $\ll_{\mathbb{S}}^{inc}$- maximal α-consistent subbase of $\mathbb{S}$. Let ω be a model of $A\cup\{\alpha\}$; then $[\omega]\cap\mathbb{S}\supseteq A$. As the relation $\ll_{\mathbb{S}}^{inc}$ refines set inclusion, A $\ll_{\mathbb{S}}^{inc}$-maximal implies that $[\omega]\cap\mathbb{S}$ is $\ll_{\mathbb{S}}^{inc}$- maximal and due to Prop. 6, ω is $\leq_{\mathbb{S}}^{inc}$- maximal in $\mathcal{M}(\alpha)$. Similar proof for $\ll_{\mathbb{S}}^{lex}$.

We also give other equivalent formulations:

Proposition 8: If $\alpha \neq \bot$, $\alpha \vdash^{inc}_{\mathcal{S}} \beta$ iff $\vdash^{inc}_{Add(\alpha,\mathcal{S})} \beta$. The same holds for lex-ordering instead of inc-ordering.

Sketch of proof for $\ll^{inc}_{\mathcal{S}}$: let $A \subseteq \mathcal{S}$; A is $\ll^{inc}_{\mathcal{S}}$-maximal α-consistent iff $A \cup \{\alpha\}$ is $\ll^{inc}_{Add(\alpha,\mathcal{S})}$-maximal consistent subbase of $Add(\alpha, \mathcal{S})$. Indeed in $Add(\alpha,\mathcal{S})$ the first level is a singleton $\{\alpha\}$, so set-inclusion verifications are equivalent on the first level A_0. Similar proof for $\ll^{lex}_{\mathcal{S}}$.

Proposition 9:
$\alpha \vdash^{bo}_{\mathcal{S}} \beta \Rightarrow \alpha \vdash^{inc}_{\mathcal{S}} \beta \Rightarrow \alpha \vdash^{lex}_{\mathcal{S}} \beta$ (as shown with Examples 2 and 3, the converses are false).

4.3 Properties

Proposition 10: $\vdash^{inc}_{\mathcal{S}}$-inferences satisfy Supraclassicality, Left logical equivalence, right weakening, And, Weak conditionalization, Consistency preservation, Cumulativity, Or (System P in [21]).

Proposition 11: $\vdash^{inc}_{\mathcal{S}}$-inferences do not satisfy Weak rational monotony [14] (nor a fortiori, Rational monotony).

Proof: a counterexample for Weak rational monotony is $\mathcal{S}$= $(\{\neg a \vee \neg b, \neg a \vee c\}, \{a,b\}, \{\neg c\})$; then $\vdash^{inc}_{\mathcal{S}} \neg c$, $\vdash^{inc}_{\mathcal{S}} a \vee \neg c$ and nevertheless $c \not\vdash^{inc}_{S} a \vee \neg c$.

Proposition 12: the class of $\vdash^{inc}_{\mathcal{S}}$ inference relations strictly includes the class of Gärdenfors and Makinson [14]'s comparative relations and is strictly included in the class of Kraus et al. [21]'s preferential relations.

Sketch of proof: for the first inclusion, we need to show that for any expectation ordering E there is an $\mathcal{S}$ such that $\vdash_E = \vdash^{inc}_{\mathcal{S}}$); n being the number of levels induced by E, this can be done by considering the belief set $\mathcal{S}$ obtained by putting every formula at its E-level. Then $\alpha \vdash^{inc}_{\mathcal{S}} \beta$ is equivalent to $\mathcal{S}^*_\alpha \vdash \beta$ where the revision $\mathcal{S}^*_\alpha$ is a partial meet in the sense of Gärdenfors. The strictness of the inclusion is a corollary of Proposition 11. The second inclusion comes from the fact that preferential models are strictly more general than ours since the partial ordering is defined on so-called states; each state is labelled by a single interpretation, where an interpretation may label distinct states.

The properties of lex-inferences stem from this result:

Proposition 13: any $\vdash^{lex}_{\mathcal{S}}$-inference is a rational inference relation, i.e. satisfies all properties in proposition 10 plus Rational monotony.

Sketch of proof: from proposition 7, the fact that $\ll^{lex}_{\mathcal{S}}$ is complete, and Gärdenfors and Makinson's characterization of nonmonotonic inference by ranked models.

The last result states the other direction, i.e. that the class of lex-inferences contains the class of comparative inference relations.

Proposition 14: for any comparative inference relation $\vdash$ there is a belief base $\mathcal{S}$ such that $\vdash$ is equal to $\vdash^{lex}_{\mathcal{S}}$.

Corollary 15: the class of lex-inferences is equal to the class of comparative inference relations.

As a conclusion, what we gain with inclusion-based inferences is that we avoid the drowning effect, but at the same time we lose Rational monotony; now, with lex-inferences, we still avoid the drowning effect and we furthermore recover Rational monotony.

5 Computing Nonmonotonic Inference Relations from a Belief Base

In this section we briefly describe a procedure for computing inclusion-based and lex-inferences. The procedure presupposes the existence of an ATMS computing incrementally the set of *minimal candidates* [7], which are defined as minimal sets of assumptions of a belief base whose deletion restores the consistency (i.e., a minimal candidate is the complement of a maximal consistent subbase). Note that, due to Proposition 8, it is sufficient to have a procedure deciding if $\vdash^{inc}_{\mathcal{S}} \beta$ (resp. $\vdash^{lex}_{\mathcal{S}} \beta$).

Decision procedure for $\vdash^{inc}_{\mathcal{S}} \beta$

$\mathcal{SB} \leftarrow \{\emptyset\}$ {set of $\ll^{inc}_{\mathcal{S}}$-preferred subbases of $\mathcal{S}_1 \cup \ldots \mathcal{S}_k$}
$k \leftarrow 0$; {working level}; Stop $\leftarrow$ False
Repeat
1 $k \leftarrow k+1$; New-$\mathcal{SB} \leftarrow \emptyset$; Answer $\leftarrow$ Yes
2 **For each** B **in** $\mathcal{SB}$ **do**
3 $\Sigma_k \leftarrow$ set of justifications obtained from $\mathcal{S}_k$ [1]
4 Compute the set $\mathcal{C}_k$ of minimal candidates of $B \cup \Sigma_k$ using an ATMS[2];
5 **For each** C **in** $\mathcal{C}_k$ **do**
6 **If** $B \cup (\mathcal{S}_k \setminus C) \models \neg\beta$ **then** Stop $\leftarrow$ True; Answer $\leftarrow$ No[3]
else if $B \cup (\mathcal{S}_k \setminus C) \not\models \beta$ **then**
If $k = n$ **then** Stop $\leftarrow$ True; Answer $\leftarrow$ No [4]
else New-$\mathcal{SB} \leftarrow$New-$\mathcal{SB} \cup \{B \cup (\mathcal{S}_k \setminus C)\}$;[5]
If New-$\mathcal{SB} = \emptyset$ **then** Stop $\leftarrow$ True
$\mathcal{SB} \leftarrow$New-$\mathcal{SB}$
until Stop

Deciding whether $\vdash^{lex}_{\mathcal{S}} \beta$ holds is very similar, except at (4) where only candidates of minimal cardinality are computed, and at (6) where the test is not done any more, since it is

---

[1] This means that an assumption A_j is associated to each formula φ_j of $\mathcal{S}_k$, in the following way: let $\{C_{i,1},\ldots,C_{i,p}\}$ a set of clauses equivalent to φ_j; then we replace φ_j by the justifications $A_j \rightarrow C_{i,1},\ldots,A_j \rightarrow C_{i,p}$.

[2] This means that we obtain the maximal consistent sub-bases of $B \cup \mathcal{S}_k$ such that all rejected formulas are formulas of $\mathcal{S}_k$.

[3] If there is a $\ll^{inc}_{\mathcal{S}}$-preferred consistent sub-base of $\mathcal{S}_1 \cup \ldots \mathcal{S}_k$ deriving $\neg\beta$ then there will be a $\ll^{inc}_{\mathcal{S}}$-preferred consistent sub-base of $\mathcal{S}$ deriving $\neg\beta$ and it is not worth going on.

[4] If we are at the last level, finding a preferred sub-base not deriving β is sufficient to conclude "no".

[5] Note that in the case where $B \cup (\mathcal{S}_k \setminus C) \models \beta$, we do nothing, even if $k < n$: indeed, in this case, it is guaranteed that this sub-base will lead only to preferred sub-bases of $\mathcal{S}$ deriving β, so it is not necessary to go on computing them.

not guaranteed that $B \cup (\Sigma_k \setminus C)$ will lead to a lex-preferred subbase of Σ.

The procedures are based on the constructive characterization of preferred maximal consistent subbases and the following results: i) for any k and for each $\ll^{inc}_{\Sigma}$-preferred subbase B of $\Sigma_1 \cup \ldots \Sigma_k$ there is a $\ll^{inc}_{\Sigma}$-preferred subbase A of Σ such that $A \cap (\Sigma_1 \cup \ldots \Sigma_k) = B$ (this is not true for lex-preferred subbases); ii) for any $\ll^{inc}_{\Sigma}$ (resp. $\ll^{lex}_{\Sigma}$)-preferred subbase A of Σ and for any k there is a $\ll^{inc}_{\Sigma}$ (resp. $\ll^{lex}_{\Sigma}$)-preferred subbase B of $\Sigma_1 \cup \ldots \cup \Sigma_k$ such that $A \cap (\Sigma_1 \cup \ldots \cup \Sigma_k) = B$.

6 Applications and Conclusion

In this paper we started by noticing the limitations of several approaches to non-monotonic reasoning such as expectation-based inferences, possibilistic logic and System Z, which cope with inconsistency at the expense of taking away too many pieces of knowledge (the drowning effect). We then proposed two ways of coping with this problem: the first one, inferring a conclusion iff it is deducible from all inclusion-based preferred subbases, avoids it but fails to satisfy Weak rational monotony. The second one, less cautious, inferring a conclusion iff it is deducible from all lex-preferred subbases, still avoids the drowning effect and enables us to recover Rational Monotony.

There are many situations where all this can be applied, of which we mention some examples. First in default reasoning, Geffner and Pearl [16] have already used inclusion-based prioritized inference to mend System Z. Here we have shown that the lexicographically-preferred subbases led to another solution recovering property inheritance. It would be interesting to compare this solution to the approach based on maximum entropy [17]. A second potential application is model-based diagnosis (in this spirit, see [19]) where S describes the functioning of the system to diagnose (the levels reflecting the certainty of the rules, and the reliability of the components); a corresponds to the observed situation; then each of the maximal consistent subbases correspond to a consistency-based diagnosis (where the absent formulas correspond to faulty components). In the non-gradual case, subbases of maximum cardinality correspond to a minimum number of faulty components. Lex-inferences are a generalisation of this principle, and are thus very natural in this context. Other potential applications are consistency maintenance in temporal data bases (where recent informations are preferred to older ones), prioritized constraint satisfaction problems (where overconstrained problems are solved by taking priorities into account), or to minimisation of surprizes in a logic of time and action.

Acknowledgements: This work has been partially supported by the European ESPRIT Basic Research Action n° 6156 entitled "Defeasible Reasoning and Uncertainty Management Systems (DRUMS-2)". Thanks also to Thomas Schiex for checking the algorithm.

References

1 C.Baral, S.Kraus, J.Minker, V.S.Subrahmanian Combining knowledge bases consisting in first-order theories. Computational Intelligence 8:45-71:1992.

2 S.Benferhat, D.Dubois, H.Prade Representing default rules in possibilistic logic. KR'92, 673-684.

3 C.Boutilier What is a default priority? Proc. 9th Canadian Conf. on AI, 140-147, 1992.

4 G.Brewka Preferred subtheories: an extended logical framework for default reasoning. IJCAr89, 1043-1048.

5 C.Cayrol Un modele logique general pour le raisonnement rdvisable. Revue d'Intell. Artif. 6:255-284:1992.

6 C.Cayrol, V.Royer, C.Saurel Management of preferences in assumption-based reasoning. IPMU'92, 221-224. Ext. vers. Tech. Rep. IRIT/92-13/R.

7 J.De Kleer Extending the ATMS. Artif. Intell. 28:163-196:1986.

8 D.Dubois, J.Lang, H.Prade Inconsistency in possibilistic knowledge bases. In Zadeh, Kacprzyk (eds.), Fuzzy Logic for the Management of Uncertainty, Wiley, 335-351, 1992.

9 D.Dubois, H.Prade Epistemic entrenchment and possibilistic logic. Artificial Intelligence 50:223-239:1990.

10 D.Dubois, H.Prade Possibilistic logic, preferential models, nonmonotonicity and related issues. IJCAr91, 419-424.

11 R.Fagin, J.D.Ullman, M.Y.Vardi. On the semantics of updates in databases. 2nd ACM SIGACT-SIGMOD Symp. on Principles of Databases Systems, 352-365.

12 M.Ginsberg, D.E.Smith Reasoning about action I: A possible world approach. Artif. Intell. 35:165-195:1988.

13 P.Gardcnfors Knowledge in Flux: Modeling the Dynamics of Epistemic States. MIT Press, 1988.

14 P.Gardcnfors, D.Makinson Nonmonotonic inferences based on expectations. Artificial Intelligence, 1993, to appear.

15 H.Geffner Default Reasoning: Causal and Conditional Theories. MIT Press, 1992.

16 H.Geffner, J.Pearl Conditional entailment; bridging two approaches to default reasonnig. Artif. Intell. 53:209-244:1992.

17 M.Goldszmidt, P.Morris, J.Pearl A maximum entropy approach to nonmonotonic reasoning. AAAI'90, 646-652.

18 B.Grosof Generalizing priori tarization. KR'91, 289-300.

19 U.Junker Prioritized defaults: implementation by TMS and application to diagnosis. IJCAT91, 310-315.

20 K.Konolige Hierarchic autoepistemic theories for nonmonotonic reasoning. AAAI'88, 439-443.

21 S.Kraus, D.Lehmann, M.Magidor Nonmonotonic reasoning, preferential models and cumulative logics. Artificial Intelligence 44:167-207:1990.

22 V.Lifschitz Circumscriptive theories: a logic-based framework for knowledge representation. AAAT87, 364-368.

23 D.Makinson, P.Gardenfors Relations between the logic of theory change and nonmonotonic logic. In A.Fuhrmann, M.Morreau (eds.). The Logic Of Theory Change, Springer-Verlag LNAI 465, 185-205, 1990.

24 B.Nebel Belief revision and default reasoning: syntax-based approaches. KR'91, 417-428.

25 J.Pearl System Z: a natural ordering of defaults with tractable applications to default reasoning. 3rd TARK, 121-135, 1990.

26 D.Poole A logical framework for default reasoning. Artificial Intelligence 36:1988.

27 N.Rescher Hypothetical reasoning. North-Holland, 1964.

28 N.Roos A logic for reasoning with inconsistent knowledge. Artificial Intelligence 57:69-103:1992.

29 Y.Shoham Reasoning About Change. MIT Press, 1988.