Resource-bounded inference from inconsistent belief bases

Pierre Marquis and **Nadège Porquet**
CRIL / Université d'Artois
rue de l'Université - S.P. 16
F-62307 Lens Cedex, France
{marquis, porquet}@cril.univ-artois.fr

Abstract

A family of resource-bounded paraconsistent inference relations is introduced. These relations are based on $S-3$ entailment, an inference relation logically weaker than classical entailment and parametrized by a set S of variables. Their properties are investigated, especially from the computational complexity point of view. Among the strong features of our framework is the fact that tractability is ensured each time $|S|$ is bounded and that binary connectives behave in a classical manner. Moreover, our family is large enough to include both $S-3$ inference, the standard inference relations based on the selection of consistent subbases and some additional forms of paraconsistent reasoning as specific cases.

1 Introduction

Few would dispute the fact that classical entailment is not suited to common-sense inference. One of its main drawbacks is its inability to deal with inconsistency. Indeed, any formula is a logical consequence of a contradiction. Such a trivialization of inference is often referred to as *ex falso quodlibet sequitur* and classical entailment is said to be explosive (i.e., not paraconsistent).

Many approaches have been proposed so far to address this issue. Some of them, like belief revision, aim at preventing inconsistencies from being generated; other approaches aim at removing inconsistencies once they appeared. Complementary to them, many approaches deal with inconsistencies. Among them are various paraconsistent logics, argumentative logics, belief merging and the so-called coherence-based approach[1] to inconsistency handling. Unfortunately, existing paraconsistent inference relations are typically intractable (see [Cayrol *et al.*, 1998] [Nebel, 1998]).

In this paper, we show that trivialization in presence of inconsistency and intractability can be handled within a uniform framework. A family of paraconsistent resource-bounded inference relations is presented. Our starting point is the notion of $S-3$ inference [Schaerf and Cadoli, 1995] and the coherence-based approach to inconsistency handling [Pinkas and Loui, 1992] [Benferhat *et al.*, 1993]. In $S-3$ logic, every propositional symbol is interpreted in a weak way, allowing it to be both true and false at the same time, except for the variables of a given set S that are interpreted classically. Tractability is ensured by limiting the size of S. Interestingly, $S-3$ entailment proves sufficient to handle some inconsistencies, those that are not reachable when classical inference is limited to the variables of S. In order to handle the remaining inconsistencies while avoiding trivialization, we introduce a family of inference relations which are to $S-3$ entailment what the inference relations of the standard coherence-based approach are to classical entailment. Specifically, consistency is restored by removing variables (from S) instead of removing some explicit beliefs from the belief base. The computational complexity of these inference relations is identified, and some other properties are investigated. Among the strong features of our framework is the fact that tractability is ensured each time a computational bound $|S|$ is set. Moreover, binary connectives behave in a classical manner while it is not the case for many paraconsistent logics. Finally, our framework is general enough to encompass both $S-3$ logic [Schaerf and Cadoli, 1995], the standard coherence-based approach to inconsistency handling [Pinkas and Loui, 1992] [Benferhat *et al.*, 1993] and some systems for paraconsistent reasoning given in [Priest, 1991] [Besnard and Schaub, 1998] as specific cases.

2 Formal preliminaries

In the following, we consider a propositional language $PROP_{PS}$ inductively defined from a finite set PS of propositional symbols, the boolean constants $true$ and $false$ and the connectives $\neg, \wedge, \vee$. For every formula ϕ from $PROP_{PS}$, $Var(\phi)$ denotes the symbols of PS occurring in ϕ. As usual, every finite set of formulas is considered as the conjunctive formula whose conjuncts are the elements of the set.

Every formula ϕ from $PROP_{PS}$ is interpreted in a classical way, unless stated otherwise. In order to avoid any ambiguity, we sometimes refer to 2-interpretations, 2-models, 2-consistent and 2-inconsistent formulas instead of the usual notions of (respectively) interpretations, models, consistent formulas and inconsistent formulas.

For the sake of simplicity, we assume that every formula ϕ of $PROP_{PS}$ is in Negation Normal Form (NNF), i.e., only

[1] In this approach, beliefs are represented by a belief base B, i.e., a finite set of propositional formulas; some consistent subbases of B that are preferred w.r.t. a given selection policy are first selected, then inference amounts to classical inference from some of them [Pinkas and Loui, 1992] [Benferhat *et al.*, 1993].

variables are in the scope of any occurrence of $\neg$ in ϕ. For instance, $\neg a \vee (b \wedge \neg c)$ is in NNF while the (classically equivalent) formula $\neg(a \wedge \neg(b \wedge \neg c))$ is not. Nevertheless, this assumption could be given up easily (cf. Section 4.4).

We assume the reader familiar with the complexity classes P, $\Delta_2^p[\mathcal{O}(log\ n)]$, Δ_2^p, and Π_2^p of the polynomial hierarchy (see [Papadimitriou, 1994] for details).

3 $S-3$ inference

Our family of inference relations is based on $S-3$ logic [Schaerf and Cadoli, 1995], a multivalued logic which can be viewed as a generalization of Levesque's logic of limited inference [Levesque, 1984]. In $S-3$ logic, every propositional variable is interpreted in a weak way, allowing it to be both true and false at the same time, except for the variables of a given set S.

Definition 3.1 *Let S be a subset of PS.*

- *An $S-3$-interpretation over PS is a mapping I from the set L_{PS} of all literals over PS to $BOOL = \{0,1\}$ s.t. for every literal l of L_{PS}, we never have $I(l) = I(\neg l) = 0$, and we always have $I(l) = 1 - I(\neg l)$ whenever the variable of l belongs to S. $true$ is always interpreted as 1 while $false$ is always interpreted as 0.*
- *An $S-3$-interpretation I is an $S-3$-model of a NNF formula Σ iff the formula obtained by replacing in Σ every occurrence of $\neg l$ by $false$ (resp. $true$) when $I(\neg l) = 0$ (resp. 1), then every occurrence of l by $false$ (resp. $true$) when $I(l) = 0$ (resp. 1) evaluates classically to 1.*
- *Let B be a belief base (i.e., a conjunctively-interpreted finite set of formulas from $PROP_{PS}$). Let γ be a formula from $PROP_{PS}$. γ is an $S-3$ consequence of B, noted $B \models_S^3 \gamma$, iff every $S-3$-model of B is an $(S-3$-$)$ model of γ.*

A formula is said to be $S-3$ consistent when it has an $S-3$-model. A 3-interpretation (resp. 3-model) is just an $\emptyset-3$-interpretation (resp. $\emptyset-3$-model) and a formula is 3-consistent when it has a 3-model. Every formula from $PROP_{PS}$ which does not contain any occurrence of $false$ is 3-consistent (the 3-interpretation I s.t. $I(l) = 1$ for every literal l of L_{PS} is a 3-model of it).

As shown in [Schaerf and Cadoli, 1995], for every $S \subseteq PS$, the corresponding inference relation $\models_S^3$ is an approximation of $\models$ by below: if $B \models_S^3 \gamma$ then $B \models \gamma$.

Interestingly, the limited power of $S-3$ inference enables *some* inconsistencies to be handled:

Example 3.1 *Let $B = \{(\neg a \vee b), a, (\neg c \vee b), (\neg b \vee d), (c \wedge \neg d), e, (\neg e \vee f)\}$. B is (2-)inconsistent.*

- *Let $S = \{c, d, e\}$. We have $B \models_S^3 a \wedge c \wedge \neg d \wedge b \wedge \neg b \wedge e \wedge f$ but $B \not\models_S^3 d$.*
- *Let $S = \{b, c, d, e\}$. We have $B \models_S^3 false$. Here, $\models_S^3$ is explosive: any formula is an $S-3$ consequence of B.*

Tractability of $S-3$ entailment is ensured by limiting the size of S. To be more specific, the time complexity of $S-3$ inference is in $\mathcal{O}(2^{|S|}.|B|.|\gamma|)$ when B is a CNF formula[2] and γ is a clause [Schaerf and Cadoli, 1995]. Accordingly, the complexity of $S-3$ inference depends essentially on the number of variables in S.

4 Resource-bounded paraconsistent inference

4.1 Dealing with inconsistency

Adhering to $S-3$ entailment is a way to prevent some inconsistencies from being harmful, namely those which are derivable only if classical reasoning over some variables outside S is performed. In the previous example with $S = \{c, d, e\}$, while we have both $B \models_S^3 b$ and $B \models_S^3 \neg b$, we do **not** have $B \models_S^3 d$.

However, restricting S to an arbitrary subset of the variables occurring in B is not sufficient to avoid trivialization in every situation (see the previous example). In order to deal with the remaining inconsistencies, we suggest to focus on some subsets S' of S, those for which the corresponding inference relations are not explosive. The approach is similar to the standard coherence-based approach to inconsistency handling, except that the inference relation is weakened by removing variables from S instead of removing explicit beliefs from B. Formally, we first need a notion playing a role similar to the notion of consistent subbase in the standard coherence-based approach.

Definition 4.1 *Let B be a belief base. Let $S \subseteq PS$ and $S_0 \subseteq S$ s.t. $B \not\models_{S_0}^3 false$. A* consistent subset *$S'$ of S w.r.t. B and S_0 is a subset of S containing S_0 and s.t. $B \not\models_{S'}^3 false$. $\mathcal{S}(B, S, S_0)$ denotes the set of all consistent subsets of S w.r.t. B and S_0.*

In this definition, S_0 is a given set of variables which *must* be interpreted classically. In our framework, S_0 plays a role similar to the one played by integrity constraints. Requiring the existence of such a set S_0 (possibly empty) imposes that B is 3-consistent (which is not a strong assumption).

In the following, we adhere to a skeptical approach (every preferred consistent subset is considered):

Definition 4.2 *Let B be a belief base and S_0 be a subset of PS s.t. $B \not\models_{S_0}^3 false$. Let $\mathcal{P}$ be a selection policy s.t. $\mathcal{S}_{\mathcal{P}}(B, S, S_0)$ is a subset of $\mathcal{S}(B, S, S_0)$ and let γ be a formula from $PROP_{PS}$. γ is a* consequence *of B w.r.t. $\mathcal{P}$ and S_0, noted $B \approx_S^{\mathcal{P}, S_0} \gamma$, iff $\forall S' \in \mathcal{S}_{\mathcal{P}}(B, S, S_0)$, $B \models_{S'}^3 \gamma$.*

Now, we can define selection policies for consistent subsets that are similar to the ones defined for consistent subbases in the standard coherence-based approach [Pinkas and Loui, 1992] [Benferhat *et al.*, 1993] [Benferhat *et al.*, 1995].

Definition 4.3 *Let B be a belief base. Let $S = (S_0, \ldots, S_j)$ be a stratification[3] of a given subset of PS where S_0 is s.t. $B \not\models_{S_0}^3 false$.*

- *The set $\mathcal{S}_{\mathcal{PO}}(B, S, S_0)$ of all preferred consistent subsets of S w.r.t. B and S_0 for the possibilistic policy is the singleton $\{\bigcup_{i=0}^{s-1} S_i\}$, where s is either the smallest index ($1 \leq s \leq j$) s.t. $\bigcup_{i=0}^{s} S_i$ is an inconsistent subset of S w.r.t. B and S_0 or $j+1$ if S is consistent w.r.t. B.*

[2] Combining both Theorem A.3 from [Schaerf and Cadoli, 1995], Lemma 3.1 and Theorem 4.1 from [Cadoli and Schaerf, 1996], this result can be easily generalized to the case where B is any NNF formula and γ is a CNF formula; in this situation, $S-3$ inference can still be decided in time linear in the size of the belief base.

[3] A stratification of S is a totally ordered partition of it.

- *The set $\mathcal{S}_{\mathcal{LO}}(B, S, S_0)$ of all preferred consistent subsets of S w.r.t. B and S_0 for the linear order policy is the singleton $\{\bigcup_{i=0}^{j} S'_i\}$, where S'_i ($i \in 0, \ldots, j$) is defined by $S'_i = S_i$ if $S_i \cup \bigcup_{l=0}^{i-1} S'_l$ is a consistent subset of S w.r.t. B and S_0, $\emptyset$ otherwise.*
- *The set $\mathcal{S}_{\mathcal{IP}}(B, S, S_0)$ of all preferred consistent subsets of S w.r.t. B and S_0 for the inclusion-preference policy is $\{S' \mid S'$ is a consistent subset of S w.r.t. B and S_0 and $\forall S'' \neq S'$ s.t. $S'' \in \mathcal{S}(B, S, S_0)$, $\forall i \in 0, \ldots, j$ $((\forall l < i (S' \cap S_l = S'' \cap S_l)) \Rightarrow S' \cap S_i \not\subset S'' \cap S_i)\}$.*
- *The set $\mathcal{S}_{\mathcal{LE}}(B, S, S_0)$ of all preferred consistent subsets of S w.r.t. B and S_0 for the lexicographic policy is $\{S' \mid S'$ is a consistent subset of S w.r.t. B and S_0 and $\forall S'' \neq S'$ s.t. $S'' \in \mathcal{S}(B, S, S_0)$, $\forall i \in 0, \ldots, j$ $((\forall l < i(|S' \cap S_l| = |S'' \cap S_l|)) \Rightarrow |S' \cap S_i| \geq |S'' \cap S_i|)\}$.*

These definitions assume that a stratification of the given set S of variables is available. Such a stratification is the formal counterpart of the preferential information used to discriminate among consistent subsets.

Here is an example showing the variety of conclusions that can be drawn in our framework (if both a formula γ and its negation are consequences of B, the truth of γ given B must be considered as doubtful):

Example 4.1 *Let $B = \{(\neg a \vee b), a, (\neg c \vee b), (\neg b \vee d), (c \wedge \neg d), e, (\neg e \vee f)\}$. With $S = (\emptyset, \{a, b, c, d, e\})$ or $S = (\emptyset, \{a\}, \{b, c, d\}, \{e\})$, a, c, $\neg d$ and e are consequences of B (but their negations are not) whatever $\mathcal{P}$ among the four policies considered above. Moreover:*

- *With $S = (\emptyset, \{a, b, c, d, e\})$:*
  - *w.r.t. $\mathcal{PO}$ (or $\mathcal{LO}$), none of b, $\neg b$, f and $\neg f$ are consequences of B.*
  - *w.r.t. $\mathcal{IP}$, f is a consequence of B, but b and $\neg b$ are not.*
  - *w.r.t. $\mathcal{LE}$, b and f are consequences of B, but their negations are not.*
- *With $S = (\emptyset, \{a\}, \{b, c, d\}, \{e\})$:*
  - *w.r.t. $\mathcal{PO}$, b is a consequence of B, but f and $\neg f$ are not.*
  - *w.r.t. $\mathcal{LO}$, $\mathcal{IP}$ or $\mathcal{LE}$, b and f are consequences of B, but their negations are not.*

Clearly enough, in the case where $B \not\models_S^3 false$, all our inference relations $B \approx_S^{\mathcal{P},S_0}$ coincides with $B \models_S^3$. Accordingly, $S - 3$ inference can be recovered as a specific case in our framework in this situation.

The next proposition states that one of our basic objectives (avoiding trivialization) is always reached.

Proposition 4.1 *None of the relations $B \approx_S^{\mathcal{P},S_0}$ with $\mathcal{P}$ among $\mathcal{PO}$, $\mathcal{LO}$, $\mathcal{IP}$, $\mathcal{LE}$ is explosive.*

4.2 Instantiating our framework

There are many ways to define a stratified set S from a stratified belief base (SBB) B [4]. First, S_0 can be defined as the set of all variables from the certain beliefs (i.e., pieces of knowledge) of B when available, provided that such a set is small enough and satisfies $B \not\models_{S_0}^3 false$. Otherwise, S_0 is defined as $\emptyset$.

Defining S is known as a difficult problem in the general case [Schaerf and Cadoli, 1995]. Interestingly, when B is a SBB, it is possible to exploit the given preferential information so as to define S. A very natural way consists in considering in priority the variables of the most plausible beliefs until a preset bound b on the size of S is reached. Formally, we consider the following set of variables S^b:

Definition 4.4 *Let $B = (\Delta_1, \ldots, \Delta_k)$ be a SBB. Let b be a positive integer. Let S_0 be a subset of PS s.t. $B \not\models_{S_0}^3 false$ and $|S_0| \leq b$. The* inference level[5] *of B given b and S_0, noted i_B^b is the smallest i ($i \in 0..k-1$) s.t. $S_i^b = S_{i+1}^b$, where the sequence (S_i^b) is inductively defined by $S_0^b = S_0$ and ($\forall i \in 1, \ldots, k-1$) $S_i^b = S_{i-1}^b \cup Var(\Delta_i)$ if $|S_{i-1}^b \cup Var(\Delta_i)| \leq b$, $S_i^b = S_{i-1}^b$ otherwise. We define S^b as $S_{i_B^b}^b$.*

Finally, once both S_0 and a flat set S are available, a stratification of S can be easily derived from the given stratification of B. Basically, once S_0 has been taken into account, it is reasonable to give more importance to the variables of S that belong to the more plausible beliefs. Loosely speaking, we want to reason as much as classically as possible from the most plausible beliefs in order to take into account their, possibly conflicting, classical consequences: inconsistencies must not be ignored when supported by plausible beliefs.

Definition 4.5 *Let $B = (\Delta_1, \ldots, \Delta_k)$ be a SBB, S a subset of PS and S_0 a subset of S. The stratification of S induced by B given S_0 is $S = (S_0, \ldots, S_{k+1})$ s.t. $S_i = (S \cap Var(\Delta_i)) \setminus \bigcup_{l=0}^{i-1} S_l$ for every $i \in 1..k$, and $S_{k+1} = S \setminus \bigcup_{l=0}^{k} S_l$.*

There are many other ways to derive a stratification of S. For instance, in the situation where we are ready to give more credit to facts than to rules, we can take advantage of the notion of degree of definiteness[6] of a variable (as defined in [Besnard and Schaub, 1998]) to achieve this goal from a CNF belief base B.

4.3 Logical properties and cautiousness

Some logical properties. As in the standard coherence-based approach, it is easy to give a preferential models semantics to each of our relations $\approx_S^{\mathcal{P},S_0}$. Indeed, $B \approx_S^{\mathcal{P},S_0} \gamma$ holds iff every preferred (S_0-)3-model of B w.r.t. $\mathcal{P}$ is a (S_0-)3-model of γ, where the preferred (S_0-)3-models of B w.r.t. $\mathcal{P}$ are exactly the $S' - 3$-models of B with $S' \in \mathcal{S}_{\mathcal{P}}(B, S, S_0)$. Thus, when $S = (\emptyset, PS)$, the preferred 3-models of B w.r.t. $\mathcal{IP}$ are the 3-models of it encoding worlds that are "as normal as possible" (i.e., as close as possible to 2-interpretations).

[4] Upstream to our work are several techniques that aim at deriving a SBB from a "flat" belief base, see [Benferhat, 2000] for a survey.

[5] At a first glance, the computation of the inference level of a belief base looks like the computation of the inconsistency level for possibilistic belief bases. However, the former is guided by computational motivations, only; especially, it is not the case that $B \not\models_{S^b}^3 false$ in the general case.

[6] Roughly, every variable x of S occurring in B can be associated with the size of one of the smallest clauses of B containing x; a stratification of S is obtained by considering x before y whenever the score attached to x is smaller than the score attached to y.

Another strong point of our approach is that our inference relations exhibit a limited form of syntax-sensitivity. Indeed, the binary connectives $\wedge$ and $\vee$ of our language behave in a classical manner: the set of $S-3$ models of $\phi \wedge \psi$ (resp. $\phi \vee \psi$) is the intersection (resp. union) of the set of $S-3$ models of ϕ and the set of $S-3$ models of ψ. As a consequence, given a stratification $S = (S_0, \ldots, S_j)$, every formula from B (and B itself!) can be turned into a CNF formula without modifying the corresponding inference relations $\approx_S^{\mathcal{P},S_0}$. The possibility of normalizing the belief base under CNF is particularly important for readibility reasons and the practical computing of the inference relations $\approx_S^{\mathcal{P},S_0}$. Moreover, this valuable property is usually not satisfied by paraconsistent logics (e.g., the one given in [da Costa, 1974], the one given in [Lin, 1987], or the standard coherence-based approach). Contrastingly, like any paraconsistent inference relation, our relations do not satisfy left logical equivalence. For the same reason, they do not satisfy right weakening (since they satisfy reflexivity). Finally, unlike $S-3$ inference [Schaerf and Cadoli, 1995] , they are not monotonic[7] (neither in S nor in B) in the general case.

Example 4.2 *Let $B = \{a, (\neg a \vee b), (\neg a \vee c), \neg c\}$, $B' = \{a, (\neg a \vee b), (\neg a \vee c), \neg c, \neg a\}$, $S = (\emptyset, \{a\})$, and $S' = (\emptyset, \{a, c\})$. We have $B \approx_S^{\mathcal{IP},S_0} b$ but $B' \not\approx_S^{\mathcal{IP},S_0} b$ and $B \not\approx_{S'}^{\mathcal{IP},S_0} b$. Similar counter-examples can be easily found for the other policies.*

This is not very surprising: while monotonicity is a highly desirable property when B is consistent, this is not the case otherwise. It is natural that the epistemic status of beliefs changes when more evidence is got through the incorporation of additional explicit beliefs or through an improvement of inferential capabilities. Specifically, increasing the computational effort can lead to discover inconsistencies that were hidden before.

Cautiousness. In the spectrum of inference relations in our framework, one extreme bound corresponds to the case $S = \emptyset$, while at the other extreme bound, we have $S = PS$. None of these two extreme cases can be considered as a "reasonable" inference relation for common-sense reasoning. On the one hand, the latter case corresponds to an ideal agent, with unlimited inferential capabilities: $\models_{PS}^3$ coincides with $\models$. Accordingly, if B is 2-consistent, all our inference relations $B \approx_{PS}^{\mathcal{P},S_0}$ coincide with $B \models$ (admittedly, this contrasts with many paraconsistent inference relations). On the other hand, the former case corresponds to an agent with very limited reasoning capabilities. Indeed, if $S = \emptyset$, all our inference relations correspond to Levesque's 3-entailment [Levesque, 1984]. This inference relation is tractable and avoids trivialization (as long as B does not contain any occurrence of $false$). However, with this inference relation, disjunctive syllogism can never be applied: it is not possible to conclude b from $a \wedge (\neg a \vee b)$. This suggests that cautiousness must be taken into account as an important criterion in the choice of a relation.

[7]Nevertheless, we have a "stratumwise" monotonicity property: for $\mathcal{P} \in \{\mathcal{PO}, \mathcal{LO}, \mathcal{IP}, \mathcal{LE}\}$, for every $i \in 0, \ldots, j-1$, if $B \approx_{(S_0,\ldots,S_i)}^{\mathcal{P},S_0} \gamma$, then $B \approx_{(S_0,\ldots,S_{i+1})}^{\mathcal{P},S_0} \gamma$. This property is helpful to check whether $B \approx_S^{\mathcal{P},S_0} \gamma$ in an incremental fashion.

In the case where S is a consistent subset, all our inference relations coincide with $\models_S^3$, so they are just as cautious as $\models_S^3$, and the size of S can be viewed as a degree of cautiousness. In the remaining case, our inference relations are more cautious relations than $\models_S^3$ but this is what is expected since trivialization must be avoided. In this situation, our inference relations do not coincide in the general case, and results similar to the corresponding ones in the standard coherence-based approach can be obtained; especially, the less cautious inference relations correspond to the $\mathcal{LO}$ and $\mathcal{LE}$ policies.

Proposition 4.2

- *$\approx_S^{\mathcal{PO},S_0}$ is more cautious than $\approx_S^{\mathcal{LO},S_0}$ and $\approx_S^{\mathcal{IP},S_0}$, but the converse does not always hold.*
- *$\approx_S^{\mathcal{IP},S_0}$ is more cautious than $\approx_S^{\mathcal{LE},S_0}$ but the converse does not always hold.*
- *$\approx_S^{\mathcal{LO},S_0}$ cannot be compared with $\approx_S^{\mathcal{IP},S_0}$ and $\approx_S^{\mathcal{LE},S_0}$ w.r.t. cautiousness in the general case.*

4.4 A general framework

Our framework is quite general. As evoked previously, it includes $S-3$ inference as a specific case. As a consequence, it also includes the inference relation $\vdash_{LP}$ of the three-valued logic LP [8] [Priest, 1989] with $\{1, \frac{1}{2}\}$ as designated truth values, defined by:

Definition 4.6

- *An LP-interpretation I over PS is a mapping from PS to $THREE = \{0, \frac{1}{2}, 1\}$.*
- *The semantics of a formula in an LP-interpretation I is defined inductively by:*
  - $I(true) = 1$ *and* $I(false) = 0$*;*
  - $I(\neg\phi) = 1 - I(\phi)$*;*
  - $I(\phi \wedge \psi) = min(I(\phi), I(\psi))$*;*
  - $I(\phi \vee \psi) = max(I(\phi), I(\psi))$.
- *I is an LP-model of Σ iff $I(\Sigma) > 0$.*
- *$\Sigma \vdash_{LP} \gamma$ holds iff every LP-model of Σ is an (LP-)model of γ.*

Indeed, it is not very difficult to prove that $\vdash_{LP}$ coincides with $\models_\emptyset^3$. More interestingly, let $\leq$ be the preference ordering over LP-interpretations defined by $I \leq J$ iff $\{x \in PS \mid I(x) = \frac{1}{2}\} \subseteq \{x \in PS \mid J(x) = \frac{1}{2}\}$; then, the relation $\vdash_{LP_m}$ defined by $\Sigma \vdash_{LP_m} \gamma$ iff every minimal LP-model of Σ w.r.t. $\leq$ is an (LP-)model of γ, can be recovered as a specific case of our $\approx_S^{\mathcal{IP},S_0}$ relation:

Proposition 4.3 *Let Σ and γ be two formulas from $PROP_{PS}$. We have $\Sigma \vdash_{LP_m} \gamma$ iff $\{\Sigma\} \approx_{PS}^{\mathcal{IP},\emptyset} \gamma$.*

Our work is also closely related to the approach to paraconsistent reasoning given in [Besnard and Schaub, 1998]. In this work, every formula Σ of $PROP_{PS}$ is associated to a default theory $\langle \Sigma^{\pm}, D_{PS} \rangle$ where:

[8]LP is also known asJ_3, see [Epstein, 1990].

- $\Sigma^{\pm}$ is a formula in the language $PROP_{PS^{\pm}}$ where $PS^{\pm} = \{x^{+} \mid x \in PS\} \cup \{x^{-} \mid x \in PS\}$; $\Sigma^{\pm}$ is obtained by replacing in Σ every occurrence of a positive literal x by the positive literal x^{+} and every occurrence of a negative literal $\neg x$ by the positive literal x^{-}.
- $D_{PS} = \{\delta_x \mid x \in PS\}$ is a set of default rules

$$\delta_x = \frac{: x^{+} \Leftrightarrow \neg x^{-}}{(x \Leftrightarrow x^{+}) \wedge (\neg x \Leftrightarrow x^{-})}$$

Among the inference relations that can be defined from $\langle \Sigma^{\pm}, D_{PS} \rangle$ is skeptical signed inference; a formula γ of $PROP_{PS}$ is a skeptical signed consequence of Σ, noted $\Sigma \vdash_s^{\pm} \gamma$ iff $\gamma^{\pm}$ belongs to every extension of $\langle \Sigma^{\pm}, D_{PS} \rangle$.

Interestingly, the relation $\vdash_s^{\pm}$ can be recovered as a specific case of our $\approx_S^{\mathcal{IP},S_0}$ relation.

Proposition 4.4 *Let Σ and γ be two formulas from $PROP_{PS}$. We have $\Sigma \vdash_s^{\pm} \gamma$ iff $\{\Sigma\} \approx_{PS}^{\mathcal{IP},\emptyset} \gamma$.*

Finally, our approach encompasses as well the skeptical inference relations $\mid\!\sim^{\mathcal{P}}$ of the standard coherence-based approach[9] as defined in [Pinkas and Loui, 1992] [Benferhat *et al.*, 1993] [Benferhat *et al.*, 1995].

Proposition 4.5 *Let $B = (\Delta_1, \ldots, \Delta_k)$ be a SBB. We associate to B in polynomial time the belief base $B_{new} = \{new_{i,j}, \neg new_{i,j} \vee \phi_{i,j} \mid \phi_{i,j} \in \Delta_i\}$ and the stratification $S = (Var(B), \{new_{1,j} \mid \phi_{1,j} \in \Delta_1\}, \ldots, \{new_{k,j} \mid \phi_{k,j} \in \Delta_k\})$ where all the $new_{i,j}$ are new variables (not occurring in B). For every selection policy $\mathcal{P}$ among $\mathcal{PO}$, $\mathcal{LO}$, $\mathcal{IP}$, $\mathcal{LE}$ and every formula γ not containing a new variable, we have $B \mid\!\sim^{\mathcal{P}} \gamma$ iff $B_{new} \approx_S^{\mathcal{P},Var(B)} \gamma$.*

Focusing on variables instead of beliefs gives less syntax-sensitivity and some flexibility that is hardly achieved by the standard coherence-based approach. For instance, let us assume that B gathers beliefs stemming from two different sources of information. The first source states $a \wedge b$ while the second source states $\neg a \wedge \neg b$. The two sources are equally reliable; however, the first source is more reliable to what concerns b while the second source is more reliable to what concerns a. With $\mathcal{P} \in \{\mathcal{IP}, \mathcal{LE}\}$, the standard coherence-based approach gives $(\{a \wedge b, \neg a \wedge \neg b\}) \mid\!\sim^{\mathcal{P}} a \vee \neg b$, but the expected conclusion is just the *negation* of $a \vee \neg b$! Indexing variables with sources, we can express in our framework that b from source 1. is more reliable than a from source 1. and the opposite holds when source 2. is considered. Indeed, with $B = \{(\neg a_1 \vee a) \wedge (\neg a \vee a_1) \wedge (\neg a_2 \vee a) \wedge (\neg a \vee a_2) \wedge (\neg b_1 \vee b) \wedge (\neg b \vee b_1) \wedge (\neg b_2 \vee b) \wedge (\neg b \vee b_2) \wedge a_1 \wedge b_1 \wedge \neg a_2 \wedge \neg b_2\}$, $\mathcal{P} \in \{\mathcal{PO}, \mathcal{LO}, \mathcal{IP}, \mathcal{LE}\}$, and $S = (\{a, b\}, \{b_1, a_2\}, \{a_1, b_2\})$, we have $B \approx_S^{\mathcal{P},S_0} \neg a \wedge b$ and we do **not** have $B \approx_S^{\mathcal{P},S_0} a \vee \neg b$.

4.5 Dealing with intractability

Despite the generality of our framework, our inference relations are just as hard as the corresponding ones in the standard coherence-based approach when the size of S is not bounded.

[9] A converse polynomial encoding based on $\pm$ transformation exists as well; we omit it here due to space limitations.

$\mathcal{P}$	CLAUSE / LITERAL $\approx_S^{\mathcal{P},S_0}$
$\mathcal{PO}$	$\Delta_2^p[\mathcal{O}(log\ n)]$-complete
$\mathcal{LO}$	Δ_2^p-complete
$\mathcal{IP}$	Π_2^p-complete
$\mathcal{LE}$	Δ_2^p-complete

Table 1: Complexity of inference.

Definition 4.7 *For every selection policy $\mathcal{P}$,* CLAUSE $\approx_S^{\mathcal{P},S_0}$ *is the following decision problem:*

- **Input:** *A belief base B, a subset S of PS, a subset S_0 of S s.t. $B \not\models_{S_0}^{3} false$ and a CNF formula γ.*
- **Query:** *Does $B \approx_S^{\mathcal{P},S_0} \gamma$ hold ?*

LITERAL $\approx_S^{\mathcal{P},S_0}$ *is the restriction of this problem to the case where γ is a conjunction of literals.*

Proposition 4.6 *The complexity of* CLAUSE $\approx_S^{\mathcal{P},S_0}$ *and* LITERAL $\approx_S^{\mathcal{P},S_0}$ *is reported in Table 1.*

Interestingly, limiting the size of S is sufficient to ensure tractability. The next proposition shows that our second objective is reached in this case:

Proposition 4.7 *The complexity of* CLAUSE $\approx_S^{\mathcal{P},S_0}$ *from a belief base B given a stratification $S = (S_0, \ldots, S_j)$ of a* **bounded** *set of propositional variables is in* P.

The main reason is that when $|S| \leq b$, there is only 2^b subsets of S, which is a constant when b is a constant. Accordingly, the naive algorithm consisting in generating $\mathcal{S}_{\mathcal{P}}(B, S, S_0)$ by filtering out the preferred consistent elements of 2^S, and testing for every of the resulting elements S' whether or not $B \models_{S'} \gamma$ holds runs in time $\mathcal{O}(2^{2b}.|B|.|\gamma|)$, hence in $\mathcal{O}(|B|.|\gamma|)$ for a constant b, for every B and a CNF query γ.

Of course, this is a "trick" in some sense but we can hardly do better: unless P = NP, there is no tractable inference relation that is general enough to include $\models$ as a specific case. In addition, we argue that focusing on $B \models_S^3$ as a tractable approximation of $B \models$ in order to design a paraconsistent inference relation is not so bad, especially when compared with other possible choices. Indeed, let $B \vdash$ be a tractable approximation by below of $B \models$ upon which we want to build a tractable paraconsistent inference relation. Requiring that such a tractable paraconsistent relation coincides with $B \vdash$ whenever $B \vdash$ is not explosive and is obtained by weakening $B \vdash$ through the removal of some beliefs of B otherwise limits the spectrum of interesting relations $\vdash$. For instance, considering the restriction $B \vdash_{unit}$ of $B \models$ where only unit-refutation derivable consequences of B are considered would lead to a complex semantics with respect to which $\wedge$ and $\vee$ do not behave classically [Crawford and Etherington, 1998]; considering the restriction $B_{Horn} \models$ of $B \models$ where B_{Horn} is the subset of Horn clauses of B would not lead to a tractable relation [Cayrol *et al.*, 1998] [Nebel, 1998].

5 Other related work

A close look at the AI literature shows that both trivialization and intractability can be handled by considering *infraclassical* inference relations, i.e., approximations by below of $B \models$.

What is quite surprising is the fact that the two issues have usually **not** been addressed together.

On the one hand, the proof theory of any paraconsistent logic limits the set of admissible proofs to a strict subset (i.e., an approximation by below) of the classical proofs in the objective of avoiding trivialization. In the standard coherence-based approach, the focus is laid on some preferred consistent subbases of B, which can be considered as approximations of B by below. Preferred subbases can also be defined by splitting the belief base B into several micro-theories characterized by their sets of variables, and selected in a dynamic way by looking at the variables occurring in the queries [Chopra and Parikh, 1999]. Unfortunately, for all these approaches, the corresponding inference relations are typically at least as hard as classical inference in the general case.

On the other hand, many approximation techniques have been developed so far to deal with the intractability issue (e.g., [Levesque, 1984] [Crawford and Kuipers, 1989] [Schaerf and Cadoli, 1995] [Selman and Kautz, 1996] [Dalal, 1996] [Crawford and Etherington, 1998]); in such approaches, the inference relation or the knowledge base itself are approximated into computationally easier one(s). However, the trivialization issue is typically not addressed. For instance, the inference relation defined in [Selman and Kautz, 1996] is always explosive when B is inconsistent, while the ones given in [Schaerf and Cadoli, 1995] [Crawford and Etherington, 1998] are explosive in the general case when B is inconsistent.

6 Conclusion

The main contribution of this paper is a family of resource-bounded paraconsistent inference relations. In contrast to many approaches to inconsistency handling, the computational issue is very central in our framework. Especially, the computational resources can be tuned to give rise to realistic, computationally viable, inference relations. Nevertheless, our framework is general enough to encompass $S-3$ logic, the standard coherence-based approach to inconsistency handling, and some other systems for paraconsistent reasoning as specific cases.

Several additional families of inference relations could be defined in our framework. Analyzing their computational and logical properties is a topic for further research.

Acknowledgements

Many thanks to the reviewers for their helpful comments, and to the Région Nord / Pas-de-Calais and the IUT de Lens for their support. Nadège Porquet is granted by the Région Nord / Pas-de-Calais.

References

[Benferhat *et al.*, 1993] S. Benferhat, C. Cayrol, D. Dubois, J. Lang and H. Prade. Inconsistency management and prioritized syntax-based entailment. *Proc. IJCAI'93*, 640-645, 1993.

[Benferhat *et al.*, 1995] S. Benferhat, D. Dubois, and H. Prade. How to infer from inconsistent beliefs without revising? *Proc. IJCAI'95*, 1449-1455, 1995.

[Benferhat, 2000] S. Benferhat. Computing specificity in default reasoning. In *Handbook of DRUMS*, vol. 5, Kluwer Academic, 147-177, 2000.

[Besnard and Schaub, 1998] Ph. Besnard and T. Schaub. Signed systems for paraconsistent reasoning. *Journal of Automated Reasoning*, 20:191-213, 1998.

[Cadoli and Schaerf, 1996] M. Cadoli and M. Schaerf. On the complexity of entailment in propositional multivalued logics. *Annals of Mathematics and Artificial Intelligence*, 18:29–50, 1996.

[Cayrol *et al.*, 1998] C. Cayrol, M.-C. Lagasquie-Schiex and Th. Schiex. Nonmonotonic reasoning: from complexity to algorithms. *Annals of Mathematics and Artificial Intelligence*, 22(3-4):207–236, 1998.

[Chopra and Parikh, 1999] S. Chopra and R. Parikh. An inconsistency tolerant model for belief representation and belief revision. *Proc. IJCAI'99*, 192-197, 1999.

[Crawford and Etherington, 1998] J.M. Crawford and D.W. Etherington. A non-deterministic semantics for tractable inference. *Proc. AAAI'98*, 1998.

[Crawford and Kuipers, 1989] J.M. Crawford and B. Kuipers. Towards a theory of access-limited logic for knowledge representation. *Proc. KR'89*, 67–78, 1989.

[da Costa, 1974] N. da Costa. On the theory of inconsistent formal systems. *Notre Dame Journal of Formal Logic*, 15:497-510, 1974.

[Dalal, 1996] M. Dalal. Anytime families of tractable propositional reasoners. *Proc. AI&Math'96*, 42–45, 1996.

[Epstein, 1990] R.L. Epstein. The semantic foundation of logic. Vol. I: propositional logics. *Kluwer Academic Publisher*, 1990.

[Levesque, 1984] H.J. Levesque. A logic of implicit and explicit belief. *Proc. AAAI'84*, 198–202, 1984.

[Lin, 1987] F. Lin. Reasoning in the presence of inconsistency. *Proc. AAAI'87*, 139-143, 1987.

[Nebel, 1998] B. Nebel. How hard is it to revise a belief base? In *Handbook of DRUMS*, vol. 3, Kluwer Academic, 77-145, 1998.

[Papadimitriou, 1994] Ch. H. Papadimitriou. *Computational Complexity*, Addison-Wesley, 1994.

[Pinkas and Loui, 1992] G. Pinkas and R.P. Loui. Reasoning from inconsistency: a taxonomy of principles for resolving conflict. *Proc. KR'92*, 709-719, 1992.

[Priest, 1989] G. Priest. Reasoning about truth. *Artificial Intelligence* 39:231-244, 1989.

[Priest, 1991] G. Priest. Minimally inconsistent LP. *Studia Logica* 50:321-331, 1991.

[Schaerf and Cadoli, 1995] M. Schaerf and M. Cadoli. Tractable reasoning via approximation. *Artificial Intelligence* 74(2):249-310, 1995.

[Selman and Kautz, 1996] B. Selman and H.A. Kautz. Knowledge compilation and theory approximation. *Journal of the Association for Computing Machinery* 43(2):193-224, 1996.