Approximate Coherence-Based Reasoning

Frédéric Koriche

JOURNAL OF APPLIED NON-CLASSICAL LOGICS

Approximate Coherence-Based Reasoning

Frédéric Koriche

LIRMM, UMR 5506, Université Montpellier II CNRS
161, rue Ada. 34392 Montpellier Cedex 5, France
koriche@lirmm.fr

ABSTRACT. *It has long been recognized that the concept of inconsistency is a central part of commonsense reasoning. In this issue, a number of authors have explored the idea of reasoning with maximal consistent subsets of an inconsistent stratified knowledge base. This paradigm, often called "coherent-based reasoning", has resulted in some interesting proposals for paraconsistent reasoning, non-monotonic reasoning, and argumentation systems. Unfortunately, coherent-based reasoning is computationally very expensive. This paper harnesses the approach of approximate entailment by Schaerf and Cadoli [SCH 95] to develop the concept of "approximate coherent-based reasoning". To this end, we begin to present a multi-modal propositional logic that incorporates two dual families of modalities: $\Box_S$ and $\Diamond_S$ defined for each subset S of the set of atomic propositions. The resource parameter S indicates what atoms are taken into account when evaluating formulas. Next, we define resource-bounded consolidation operations that limit and control the generation of maximal consistent subsets of a stratified knowledge base. Then, we present counterparts to existential, universal, and argumentative inference that are prominent in coherence-based approaches. By virtue of modalities $\Box_S$ and $\Diamond_S$, these inferences are approximated from below and from above, in an incremental fashion. Based on these features, we show that an anytime view of coherent-based reasoning is tenable.*

KEYWORDS: *coherence-based reasoning, approximate reasoning, anytime computation, multi-modal logics, four-valued logic.*

1. Introduction

It has long been recognized that the concept of inconsistency is an essential, pervasive and central part of commonsense reasoning. As observed by Perlis in [PER 97], inconsistencies arise in nearly all human-like intellectual activities, such as learning, cooperation, belief change and merging multiple opinions. Contradictions can be due to ambiguous data, conflicting information, or mutually-contradictory beliefs. But we are so accustomed to cope with this in everyday life that we do not notice it, most such contradictions being quickly resolved. Thus, since the real world force knowledge bases to cope with inconsistencies, it is essential to formalize some pragmatical and efficient ways of responding to them.

It is well known that classical, monotonic logic is inappropriate to handle the pragmatic issue of inconsistency. Indeed, any formula can be derived from contradiction. Nonmonotonic systems, based on symbolic or numeric structures for ordering pieces of knowledge, may offer a suitable way to handle inconsistency : the reasoner is prepared to the possibility of having an inconsistent set of knowledge and can derive appropriate conclusions from it without falling into triviality. In this context, an important part of the research has been influenced by the paradigm of *coherence based reasoning*, notably investigated by Rescher and Manor [RES 70], Pinkas and Loui [PIN 92], Nebel [NEB 91, NEB 98], and Benferhat and his colleagues [BEN 93, BEN 95, BEN 99]. The main idea is to start with an inconsistent knowledge base and to apply two successive mechanisms, namely, a *consolidation operation* which generates and selects several consistent subsets of the knowledge base and an *entailment relation* which uses classical logic on the consistent subsets in order to deduce nontrivial conclusions.

As noticed by Nebel in [NEB 91], an important advantage of coherence-based approaches is their *flexibility*. Different classes of consolidation operations can be distinguished according to the importance or relevance of formulas stored in the knowledge base. In particular, if priorities attached to formulas are available, then a preference ordering may be defined on the consistent subsets of the base and hence, the consolidation task has a more fined control over what formulas are discarded and what formulas are going to stay [FAG 83]. In an orthogonal way, different classes of entailment operations can be distinguished according to the cautiousness of reasoning. Three types of cautious entailment relation have received a great deal of interest in the literature: the so-called *existential*, *universal* and *argumentative* consequences [PIN 92, BEN 95, BEN 99]. The first two relations are defined as follows: a formula α is an existential (respectively universal) consequence of a knowledge base A if, and only if, α is classically inferred by at least one (respectively all) preferred consistent subset(s) of A. The third relation is a mild type of inference which can be specified as follows: a formula α is an argumentative consequence of a knowledge base A, if the latter contains an argument (i.e preferred consistent subset) that supports α, but no argument that supports its negation.

However, one of the main drawbacks of coherence-based approaches is their high computational complexity. Removing conflicts from a knowledge base is difficult and expensive since, as we know, inconsistencies may not lie on the surface and in most cases there is no single solution to eliminate them. As an extreme case, if we consider a knowledge base composed with n atomic propositions and their negation, and with no priority ordering on the literals, then there are exactly 2^n maximal consistent subsets generated by the consolidation operation. In fact, as stated in [CAY 98, NEB 98], the complexity of reasoning in the propositional case lies at least at the second level of the polynomial hierarchy. This is due to the interaction of two sources of complexity, namely, propositional entailment which is known to be coNP-complete, and the number of preferred consistent subsets which can grow exponentially. For this reason, we cannot expect to arrive at a polynomial algorithm when eliminating only one source of complexity, for instance, by restricting the knowledge base to Horn logic.

Approximate reasoning is a technique which is used in many areas of artificial intelligence to deal with the computational intractability of problems. This paradigm extends the conventional notion of an algorithm, which always returns optimal solutions to a given problem, by allowing it to provide approximate solutions that can be computed more efficiently. In the setting of knowledge representation, the overall approach is to develop a form of logic that allows weaker inferential power but that remains computationally feasible even with a full expressiveness of the representation language [LEV 84, CRA 89, LAK 94, FAG 95, SCH 95, DAL 98]. In this context, an approximate solution is a *maybe* answer which provides a middle ground between the exact *yes* and *no* answers. In a form of approximate reasoning called *sound* reasoning we have two possible answers: *yes* and *maybe no*. In the dual form, called *complete* reasoning, the two possible answers are *no* and *maybe yes*.

A particularly interesting form of approximate reasoning is *anytime reasoning* which produce better and better answers in an incremental fashion [DEA 88, ZIL 96]. In the setting of knowledge representation, the idea is to define a family of entailment relations that approximate classical entailment, by relaxing soundness or completeness of reasoning. The knowledge base can provide partial solutions even if stopped prematurely; the accuracy of the solution improves with the time used in computing the solution and may eventually converge to the exact answer. From this point of view, anytime reasoning offers a compromise between the time complexity needed to compute answers by means of approximate entailment relations and the quality of these answers. Based on this paradigm, Schaerf and Cadoli [SCH 95] present a general technique for approximating deduction problems. Their framework include a parameter S, a set of atomic propositions, which captures the quality of approximation. Based on this parameter, the authors define two dual families of entailment relations, which are respectively sound but incomplete and complete but unsound with respect to classical entailment. A logical characterization of their framework is presented in [KOR 98]. Recently, several extensions have been proposed in the literature, including notably default logic and circumscription [CAD 96], modal logics [MAS 98] and first-order logic [KOR 01].

The purpose of this article is to propose a model that handles both inconsistency and tractability. To this end, our study lies at the intersection of coherence-based reasoning and approximate reasoning. Our framework is based on a multi-modal propositional logic, presented in [KOR 98], and used to specify *approximate monotonic reasoners*. In this study, we extend our previous work in order to specify *approximate nonmonotonic reasoners*. Starting from a knowledge base A and a priority ordering on A, we introduce the notion of *approximate consolidation*, an operation which generates and selects approximate preferred consistent subsets of A. Then, we define three classes of *approximate entailment relations*, which respectively incorporate the existential principle, the universal principle and the argumentative principle. Based on these operations, we show that an "anytime" view of coherence-based reasoning is tenable. Specifically, our framework includes the following features:

– The logic is semantically founded on the notion of *resource* which reflects both the accuracy and the computational cost of the approximations.

– The framework enables *improvable reasoning*: the quality of approximations is a nondecreasing function of the resources that have been spent.

– The framework covers *dual reasoning*: both sound but incomplete and complete but unsound approximations are returned at any step.

The rest of the paper is organized as follows. Section 2 presents the logical machinery for anytime monotonic reasoners. Our main contribution lies in section 3 which is devoted to the formalization of anytime nonmonotonic reasoners. Related work and future extensions are discussed in section 4.

2. Approximate monotonic reasoning

In this section, we focus on the formalization of approximate monotonic reasoners. For this purpose, we present a propositional logic, named $\mathbf{ARL}$, for approximate reasoning. We first define the syntax and semantics of the logic, next we introduce the notion of approximate entailment and then, we examine its computational properties.

2.1. *Syntax*

The linguistic basis of $\mathbf{ARL}$ consists in a set atoms P. The language of *propositions* is the smallest set built from P and closed under the connectives $\wedge$ and $\neg$. The connectives $\vee$, $\supset$ and $\equiv$ are defined in terms of $\neg$ and $\wedge$; that is, the proposition $\alpha \vee \beta$ is an abbreviation of $\neg(\neg\alpha \wedge \neg\beta)$, the proposition $\alpha \supset \beta$ is an abbreviation of $\neg\alpha \vee \beta$, and the proposition $\alpha \equiv \beta$ is an abbreviation of $(\alpha \supset \beta) \supset (\beta \supset \alpha)$. Given a proposition α, the set of atoms that occur in α is denoted $P(\alpha)$. A *literal* is an atom or its negation. A *clause* is a finite disjunction of literals. A *knowledge base* is a finite conjunction of clauses. When there is no risk of confusion, we shall model knowledge bases as sets of clauses.

Following [SCH 95], the concept of *computational resource* is captured by a parameter S, a subset of P. The language of $\mathbf{ARL}$ is defined by the smallest set of *sentences* built from the following rules: if α is a proposition then α is a sentence, if α is a sentence then $\neg\alpha$ is a sentence, if α and β are sentences then $\alpha \wedge \beta$ is a sentence, and if α is a proposition and S is a subset of P then $\Box_S\, \alpha$ is a sentence.

The sentence $\Diamond_S\, \alpha$ is an abbreviation of $\neg\Box_S\, \neg\alpha$. Intuitively, a sentence such as $\Box_S\, \alpha$ is read "the agent knows α given the resources S". Dually, $\Diamond_S\, \alpha$ is read "the agent considers α as possible given the resources S".

2.2. *Semantics*

The basic building block of the semantics is a domain of truth values which determines the interpretation of sentences. In the context of limited reasoning, the four-valued semantics, first proposed by Belnap [BEL 77], and notably studied in [FAG 95, CAD 96] meets our needs. It is a simple modification of classical interpretation in which sentences can take as truth-values subsets of $\{0, 1\}$, instead simply 0 or 1 alone. Based on this domain, we define a *valuation* as a mapping v from P to the powerset of $\{0, 1\}$. The space of valuations generated from P is denoted V. Given two valuations v and v', we say that v is *less specific* than v', written $v \subseteq v'$, if for every atom $p \in P$, $v(p) \subseteq v'(p)$. Interestingly, we remark that the poset $(V, \subseteq)$ is a complete and distributive lattice.

The concept of negation is semantically captured by an order reversing involution in the poset $(V, \subseteq)$, often called *adjunction* [FAG 95]. Given a valuation v, the *adjunct* of v, denoted v^*, is defined according to the following conditions: $1 \in v^*(p)$ iff $0 \notin v(p)$ and $0 \in v^*(p)$ iff $1 \notin v(p)$. A *world* is a valuation w such that $w^* = w$. The space of worlds generated from P is denoted W.

The notion of resource is semantically represented by an equivalence relation between valuations. A *relative equivalence relation* for $\mathbf{ARL}$ is a map R from the powerset of P into binary relations of V such that $(v, v') \in R(S)$ iff for every atom $p \in P$, if $p \in S$ then $v(p) = v'(p)$. In the following, $R(S)(v)$ denotes the equivalence class of v with respect to S. Interestingly, it can be easily observed that for every parameter S and every valuation v, $R(S)(v)$ is a sublattice of V. In particular, $R(S)(v)$ has a unique minimal valuation and a unique maximal valuation, which respectively correspond to the meet and the join of the sublattice.

Intuitively, a relation $R(S)$ induces a partition of the set V into equivalence classes whose granularity captures the accuracy of approximation. When S increases, the partition becomes "finer" and the approximation more precise. Namely, it can be shown that, for any parameters S and S' such that $S \subseteq S'$, we have $R(S') \subseteq R(S)$. The "coarsest" partition is obtained when S is the empty set; in this case, we observe that $R(\varnothing)$ is $V \times V$. Conversely, the "finest" partition is given when S is the set P; in this case $R(P)$ is the identity relation over V.

We have now all notions in hand to assign truth values to sentences. Given a valuation v, we define the support relation $v \models \alpha$ by induction on the structure of α.

$$
\begin{array}{l}
v \models p \;\text{ iff }\; 1 \in v(p), \\
v \models \neg\alpha \;\text{ iff }\; v^* \not\models \alpha, \\
v \models \alpha \wedge \beta \;\text{ iff }\; v \models \alpha \text{ and } v \models \beta, \\
v \models \Box_S \,\alpha \;\text{ iff }\; v' \models \alpha, \text{ for all } v' \in V \text{ such that } v' \in R(S)(v).
\end{array}
$$

The specificity ordering between valuations captures an important structural property of the support relation. Namely, it can be proved that for any proposition α and any valuations v and v' such that $v \subseteq v'$, if $v \models \alpha$, then $v' \models \alpha$. This property will be frequently used in the remaining sections.

A sentence α is *satisfiable* iff there exists a possible world w such that $w \models \alpha$. We say that α is *valid*, and write $\Vdash \alpha$, iff for every $w \in W$, $w \models \alpha$ holds. Based on the fact that $w = w^*$ for every world w, it can be easily shown that a sentence is valid iff its negation is unsatisfiable. Given two sentences α and β, we say that α *entails* β, and write $\alpha \Vdash \beta$, iff $\alpha \supset \beta$ is valid.

2.3. *Entailment operations*

After an excursion into the logic $\mathbf{ARL}$, we now apply our results to the formalization of approximate entailment operations. In the setting suggested by our approach, an approximate monotonic reasoner can be specified as a function that takes as input a knowledge base A, a resource parameter S and a proposition α, and returns as output "yes" if the agent knows that A implies α given the resources S, "no" if the agent considers impossible that A implies α given the resources S, and "maybe" otherwise. In formal terms, the entailment operations are defined as follows.

$$
\begin{array}{lll}
(A,S) \Vdash_{\Box} \alpha & \text{iff} & \Vdash \Box_S \,(A \supset \alpha), \\
(A,S) \Vdash_{\Diamond} \alpha & \text{iff} & \Vdash \Diamond_S \,(A \supset \alpha).
\end{array}
$$

Interestingly, our model can be shown *improvable* and *dual*. Specifically, anytime reasoning may be defined by an increasing sequence of parameters $S_0 = \varnothing \cdots \subset S_k \cdots \subset S_n = P$ that approximate the problem of deciding whether A entails α, or not, by means of two dual families of tests $(A, S_k) \Vdash_{\Box} \alpha$ and $(A, S_k) \Vdash_{\Diamond} \alpha$. For any index k, if the reasoner returns "yes" using the relation $\Vdash_{\Box}$ then A entails α. Dually, if the reasoner answers "no" using the relation $\Vdash_{\Diamond}$ then A does not entail α. This stepwise process has the important advantage that the iteration may be stopped when a confirming answer is already obtained for a small index k. These considerations are clarified by the following theorem.

Theorem 2.1. *For any knowledge base A, any proposition α and any parameters S and S' such that $S \subseteq S'$,*

$$\text{if } (A,S) \Vdash_{\Box} \alpha \quad \text{then} \quad (A,S') \Vdash_{\Box} \alpha \text{ and } A \Vdash \alpha, \tag{1}$$

$$\text{if } (A,S) \nVdash_{\Diamond} \alpha \quad \text{then} \quad (A,S') \nVdash_{\Diamond} \alpha \text{ and } A \nVdash \alpha. \tag{2}$$

Proof. Let β be an abbreviation of $A \supset \alpha$. We begin to examine part [1]. Assume that $\Vdash \Box_S \beta$ holds. Thus, for every world w, we have $w \models \Box_S \beta$. Based on the definition of relative equivalence relations, we clearly have $R(S')(w) \subseteq R(S)(w)$. It follows that $w \models \Box_{S'} \beta$. Therefore, we obtain $\Vdash \Box_{S'} \beta$. Moreover, since $R(S)$ is reflexive, we have $w \in R(S)(w)$. It follows that $w \models \beta$. Hence, we obtain $\Vdash \beta$, as desired. Now let us turn to part [2]. Suppose we have $\nVdash \Diamond_S \beta$. Thus, there exists a world w such that $w \not\models \Diamond_S \beta$. This is equivalent to $w \not\models \neg\Box_S \neg\beta$. It follows that $w^* \models \Box_S \neg\beta$. Based on the fact that $w = w^*$, we obtain $w \models \Box_S \neg\beta$. Since $R(S')(w) \subseteq R(S)(w)$, we obtain $w \models \Box_{S'} \neg\beta$. It follows that $w \not\models \Diamond_{S'} \beta$. Hence, we have $\nVdash \Diamond_{S'} \beta$. Moreover, we know that $w \in R(S)(w)$. So, $w \models \neg\beta$. It follows that $w \not\models \beta$. Therefore, we obtain $\nVdash \beta$, as desired. $\Box$

Based on this monotony result, it is interesting to remark the convergence of approximations is always guaranteed whenever $S = P$. In particular, for any proposition α, the sentences $\Box_P \alpha \equiv \alpha$ and $\Diamond_P \alpha \equiv \alpha$ are valid in the logic $\mathbf{ARL}$.

2.4. *Computational Properties*

We now investigate the computational aspects of approximate entailment relations. The following result states that tractability is ensured by limiting the size of the parameter S. From this point of view, the precision of the inference depends on the computational effort that has been spent.

Theorem 2.2. *For any knowledge base A, any proposition α and any parameter S, there is an algorithm for deciding whether $(A,S) \Vdash_{\Box} \alpha$ holds and $(A,S) \Vdash_{\Diamond} \alpha$ holds which runs in $O((|A| + |\alpha|) \cdot 2^{|S|})$ time.*

Proof. Let us examine $(A,S) \Vdash_{\Box} \alpha$. Let β be an abbreviation of $A \supset \alpha$. $\Vdash \Box_S \beta$ iff $w \models \Box_S \beta$ for every world w. Thus, $\Vdash \Box_S \beta$ iff $R(S)(w) \models \beta$ for every equivalence class $R(S)(w)$, and $R(S)(w) \models \beta$ iff $v \models \beta$ for every valuation $v \in R(S)(w)$. Let w_{min} be the valuation defined as follows: for every $p \in P$, $w_{min}(p) = w(p)$ if $p \in S$ and $w_{min}(p) = \varnothing$ otherwise. Clearly, w_{min} is the meet of $R(S)(w)$ under the specificity ordering $\subseteq$. It follows that $R(S)(w) \models \beta$ iff $w_{min} \models \beta$. This can be done in $O(|\beta|)$ time. Since there are $2^{|S|}$ equivalence classes $R(S)(w)$, checking $\Vdash \Box_S \beta$ can be done in $O(|\beta| \cdot 2^{|S|})$ time. Dual considerations holds for $(A,S) \Vdash_{\Diamond} \alpha$. $\Box$

The above complexity result is just the worst case upper bound of an enumeration algorithm. Actually, in the case of clausal knowledge bases, one may conceive a two-phase procedure which first simplifies the initial knowledge base and next explores the resulting search space. The simplification phase proceeds as follows. In the scope of the modality $\Diamond_S$, the algorithm deletes all clauses of α that contain a literal whose atom does not occur in S. Dually, in the scope of $\Box_S$, the algorithm eliminates in any clause of α all literals whose atom does not occur in S. Since any atom in the resulting theory occurs in S, the exploration phase consists in a standard (two-valued) satisfiability algorithm. Systematic methods such as depth first search [ZHA 00] can be used to compute at the same time the satisfiability of $\Box_S\, \alpha$ and the unsatisfiability of $\Diamond_S\, \alpha$. On the other hand, local search algorithms [SCH 01] can be exploited if we concentrate on the satisfiability of $\Box_S\, \alpha$. The role of the simplification phase is to reduce the dimensions of the formula, thus gaining efficiency in the exploration phase.

The correct choice of S is crucial for the usefulness of deduction. Taking to the extreme, when S is chosen incorrectly, approximate reasoning may end up as expensive as classical reasoning. From this perspective, several heuristics have been proposed in the literature. For example, the atoms of S may be dynamically chosen using the *diversity heuristic* advocated in [DEC 94]. The diversity of an atom p is the product of the number of positive occurrences by the number of negative occurrences of p in the theory. This notion is based on the observation that an atom is a potential source of unsatisfiability only when it appears both positively and negatively in different clauses. Thus, in the scope of the modality $\Diamond_S$, the strategy consists in choosing atoms whose diversity is maximal. Dually, in the scope of $\Box_S$, the algorithm iteratively selects atoms whose diversity is minimal.

Example 2.3. *Let* $A = \{(\neg a \vee b \vee c), (a \vee b \vee d), (\neg a \vee \neg b \vee d), (\neg a \vee \neg b \vee c)\}$. *We want to show that* A *is satisfiable. We need to find a subset* S *of* $\{a, b, c, d\}$ *s.t.* $\Box_S\, A$ *is satisfiable. The diversity of the atoms* a, b, c *and* d *is* 3, 4, 0 *and* 0, *respectively. Starting with* $S = \varnothing$ *and using the minimal diversity heuristic, we add* c *and* d *to* S. *The simplification of* A *in the scope of the modality* $\Box_S$ *returns* $\{c, d\}$. *Obviously,* $\Box_S\, A$ *is satisfiable. Therefore,* A *is satisfiable.*

Example 2.4. *Suppose we want to show that* $a \supset c$ *is a logical consequence of the knowledge base* A, *defined above. We need to find a subset* S *such that the sentence* $\Diamond_S\, (A \wedge a \wedge \neg c)$ *is unsatisfiable. Now, the diversity of the atoms* a, b, c *and* d *is* 6, 4, 2 *and* 0, *respectively. Using the maximal diversity strategy, we iteratively add* a, b *and* c *to* S. *The simplification of* A *in the scope of the modal operator* $\Diamond_S$ *returns* $\{(\neg a \vee b \vee c), (\neg a \vee \neg b \vee c), a, \neg c\}$. *Clearly enough,* $\Diamond_S\, (A \wedge a \wedge \neg c)$ *is unsatisfiable. Therefore,* $a \supset c$ *is a logical consequence of* A.

3. Approximate nonmonotonic reasoning

In this section, we extend the concepts developed so far to the formalization of anytime nonmonotonic reasoners. These systems are defined in terms of *approximate consolidation* and *approximate entailment*. The quality of the information returned by these operations depends on the computational resources that have been spent. We begin to analyze approximate consolidation, next we present several classes of approximate entailment, and then we turn to the computational properties of our model.

3.1. *Consolidation operations*

A "standard" consolidation operation starts from a knowledge base and a priority ordering on the propositions of the base, and selects the preferred consistent subsets of the base. The purpose of "approximate" consolidation is to control the generation of these subsets by the notion of resource parameter.

To this end, we need some additional definitions. A *prioritized knowledge base* is a pair $(A, \leq)$ where A is a knowledge base and $\leq$ is a total preorder on A. It is equivalent to consider that A is stratified in a collection $(A_1, \cdots, A_n)$, where A_1 contains the propositions of lowest priority and A_n those of highest priority. Each knowledge base A_i is called a *stratum* of A. The structure $(A, \leq)$ is called *flat* if the relation $\leq$ is symmetric, or equivalently, if A contains an unique stratum.

Different methods have been proposed to use the priority relation in order to select "preferred" consistent subsets (see e.g. [BEN 93, NEB 98]). In this study, we focus on the so-called *inclusion-based preference ordering*, denoted $\preceq$, whose strict part is defined as follows: $B \prec C$ iff $\exists i : B \cap A_i \subset C \cap A_i$ and $\forall j : i < j \leq n, B \cap A_j = C \cap A_j$. By extension, $B \preceq C$ iff $B \prec C$ or $B = C$. We remark that the preference ordering extends set containment, that is, $B \prec C$ whenever $B \subset C$. Based on these definitions, the standard consolidation operation, denoted $\triangle$, is a mapping that takes as input a prioritized knowledge base $(A, \leq)$ and returns as output the set of maximally coherent elements of A, under the corresponding preference ordering $\preceq$.

$$\triangle(A, \leq) \quad = \quad max(\{B \subseteq A : B \text{ is satisfiable}\}, \preceq).$$

Now we incorporate the notion of computational resource. A parameter S is said *acceptable* for a prioritized knowledge base $(A, \leq)$ iff the following condition holds: if $\exists i : S \cap P(A_i) \neq \varnothing$ then $\forall j : i < j \leq n$, $P(A_j) \subseteq S$. Intuitively, the acceptability condition imposes a restriction on the choice of computational resources: if an acceptable parameter contains at least one atom of any given stratum then it must contain all atoms of strata of higher priority. To this point, it is interesting to remark that if the structure $(A, \leq)$ is flat, then every subset of P is acceptable for $(A, \leq)$.

The formalization of approximate consolidation is realized by parameterizing the operation $\triangle$ by means of two operations $\Box$ and $\Diamond$, the first one being sound, while the second one being complete, with respect to standard consolidation. Each operation returns a set of maximally coherent elements of a prioritized knowledge base whose quality is dependent on the computational resources that have been spent. The two dual *approximate consolidation operations* are defined as follows:

$$\Box(A, \leq, S) \quad = \quad max(\{B \subseteq A : \Box_S B \text{ is satisfiable } \}, \preceq),$$
$$\Diamond(A, \leq, S) \quad = \quad max(\{B \subseteq A : \Diamond_S B \text{ is satisfiable } \}, \preceq).$$

The two lemmas below capture important algebraic properties of consolidation operations. From an intuitive point of view, the first lemma states that the consolidation process is *improvable*: the quality of maximal elements generated by the operations $\Box$ and $\Diamond$ improves with the accuracy of the parameter S. The second result is even stronger than improvability; it shows that consolidation operations are *incremental*: the process only needs to expand the maximal elements generated in previous steps and does not require to perform all computations from scratch.

Lemma 3.1. *For any prioritized knowledge base* $(A, \leq)$ *and any acceptable parameters* S *and* S' *such that* $S \subseteq S'$*:*

$$\forall B \in \Box(A, \leq, S) \quad \exists C \in \Box(A, \leq, S') \quad \text{such that } B \subseteq C, \qquad [1]$$
$$\forall B \in \Diamond(A, \leq, S') \quad \exists C \in \Diamond(A, \leq, S) \quad \text{such that } B \subseteq C. \qquad [2]$$

Proof. We only examine part 1, since a dual argument holds for part 2. Suppose that there exists a base $B \in \Box(A, \leq, S)$ such that for every base $C \in \Box(A, \leq, S')$, we have $B \not\subseteq C$. We show that this leads to a contradiction. By definition, if $B \in \Box(A, \leq, S)$, then $\Box_S B$ is satisfiable. Thus, $\not\Vdash \Diamond_S \neg B$. By application of theorem 2.1[2], it follows that $\not\Vdash \Diamond_{S'} \neg B$. Therefore, $\Box_{S'} B$ is satisfiable. Since $B \notin \Box(A, \leq, S')$, there must exist a base $C \in \Box(A, \leq, S')$ such that $B \prec C$. By definition of the inclusion-based preference ordering, $\exists i : B \cap A_i \subset C \cap A_i$ and $\forall j : i < j \leq n, B \cap A_j = C \cap A_j$. By assumption, we know that $B \not\subseteq C$. So, $\exists k < i : B \cap A_k \not\subset C \cap A_k$. It follows that $B \cap A_k \neq \varnothing$. Since $\Box_S B$ is satisfiable, we must have $S \cap P(A_k) \neq \varnothing$. Moreover, since S is an acceptable parameter for $(A, \leq)$, it follows that $S \subseteq P(A_{k'})$ for every $k' > k$. Let B' be the base $\bigcup\{C \cap A_{k'} : k' > k\}$. Obviously, $\Box_{S'} B'$ is satisfiable. Moreover, since $B \cap A_i \subset B' \cap A_i$ and $\forall j : i < j \leq n$, $B \cap A_j = B' \cap A_j$, we thus have $B \prec B'$. Therefore, $B \notin \Box_S(A, \leq, S)$, hence contradiction. $\Box$

Lemma 3.2. *For any prioritized knowledge base $(A, \leq)$ and any acceptable parameters S and S' such that $S \subseteq S'$:*

$$\forall B \in \Box(A, \leq, S') \quad \exists C \in \Box(A, \leq, S) \quad \text{such that } C \subseteq B, \qquad [1]$$

$$\forall B \in \Diamond(A, \leq, S) \quad \exists C \in \Diamond(A, \leq, S') \quad \text{such that } C \subseteq B. \qquad [2]$$

Proof. As for the previous lemma, we only examine part 1, since a dual strategy holds for part 2. Suppose that there exists a base B in $\Box(A, \leq, S')$ such that for every member C in $\Box(A, \leq, S)$, we have $C \not\subseteq B$. We show that this leads to a contradiction.

Let α be the clause $\bigvee\{\beta \in A/B : \Box_S \beta$ is satisfiable $\}$. Obviously, $\Box_{S'} (B \cup \{\alpha\})$ is unsatisfiable. By contraposition of theorem 2.1[2], it follows that $\Box_S (B \cup \{\alpha\})$ is unsatisfiable. Let B' be a subset of B constructed by the following procedure. First, we assign the empty set to B' and α to a temporary clause γ. Second, we choose a literal l of γ such that its atom is a member of S and we choose a clause β of B such that the negation of l is in β. Third, we add β to B' and we assign the resolvent of γ and β to γ. If $\Box_S \gamma$ is satisfiable, then we go back to the second step. Otherwise, we return B'. We remark that $\Box_S B'$ is satisfiable and that $\Box_S (B' \cup \{\alpha\})$ is unsatisfiable.

Clearly enough, the set B' can be extended to a maximally preferred base D of $\Box(B, \leq, S)$. Since $D \notin \Box(A, \leq, S)$, there exists a base $C \in \Box(A, \leq, S)$ such that $D \prec C$. By definition of the inclusion-based preference ordering, $\exists i : D \cap A_i \subset C \cap A_i$ and $\forall j : i < j \leq n, D \cap A_j = C \cap A_j$. Suppose that $D \not\subset C$. In this case, $\exists k < i : D \cap A_k \not\subset C \cap A_k$. It follows that $D \cap A_k \neq \varnothing$. Since $\Box_S D$ is satisfiable, we must have $S \cap P(A_k) \neq \varnothing$. Moreover, since S is an acceptable parameter for $(A, \leq)$ we obtain $S \subseteq P(A_{k'})$ for every $k' > k$. It follows that $D \cap A_{k'} = B \cap A_{k'}$ for every $k' > k$. Therefore $B \prec C$. Moreover, since $\Box_S C$ is satisfiable, by application of theorem 2.1[2], it follows that $\Box_{S'} C$ is satisfiable. Therefore $B \notin \Box(A, \leq, S')$, hence contradiction. So, we must have $D \subset C$. By assumption, we know that $C \not\subseteq B$. Thus, there exists a clause $\beta \in C$ such that $\beta \in A/B$. However, since $\Box_S (D \cup \alpha)$ is unsatisfiable, it follows that $\Box_S (D \cup \beta)$ is unsatisfiable. Thus, $\Box_S C$ is unsatisfiable. Therefore, $C \notin \Box(A, \leq, S)$, hence contradiction. $\Box$

Example 3.3. *The figure 1 below illustrates the space of maximally coherent elements generated from a flat knowledge base $A = \{a, (\neg a \vee b), \neg b, c, \neg c\}$ and a sequence of parameters $(\varnothing, \{a\}, \{a, b\}, \{a, b, c\})$. The nodes and the edges represent the elements and the inclusion relation, respectively. The elements in lower side of the figure are generated by the operation $\Box$, while those on the upper side are generated by $\Diamond$. The elements in the center correspond to the standard maximally coherent subsets of A.*

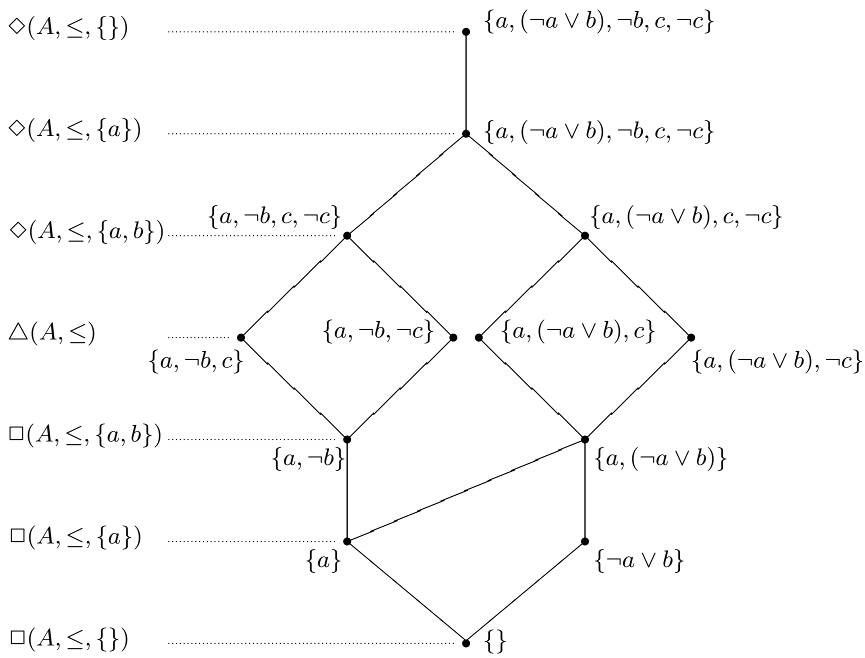

Figure 1. *Anytime consolidation used in example 3.3*

3.2. *Entailment operations*

In the setting of coherence based-reasoning, a "standard" entailment relation takes as input a knowledge base and a query and returns as output a cautious conclusion that handles the potential contradictions in the knowledge base. A taxonomy of several entailment principles has been established in [PIN 92] according to their cautiousness. In this study, we are interested in three of them: the existential principle, the universal principle and the argumentative principle. We begin to present these different classes of entailment relations and then we examine their corresponding approximations.

The first two entailment principles, introduced by Rescher and Manor in [RES 70], are the most commonly used in presence of contradictory knowledge bases (see e.g. [BRE 89, BAR 91, BAR 92]). They can be defined as follows:

$$(A, \leq) \Vdash^{\exists} \alpha \quad \text{iff} \quad \exists B \in \triangle(A, \leq) \text{ such that } \models B \supset \alpha,$$
$$(A, \leq) \Vdash^{\forall} \alpha \quad \text{iff} \quad \forall B \in \triangle(A, \leq), \models B \supset \alpha.$$

Obviously, universal entailment is more cautious than existential entailment, since each conclusion obtained from $(A, \leq)$ using $\Vdash^{\forall}$ is also obtained by $\Vdash^{\exists}$. In fact, universal entailment is often too conservative and hence rather unproductive while existential entailment is often too permissive and may lead to pairs of mutually exclusive conclusions. The notion of argumentative entailment, suggested for instance in [PIN 92, BEN 95, BEN 99], is based on an intermediate principle which is more productive than universal entailment but does not lead to contradictory conclusions. It consists in keeping only the consequences obtained by the existential principle whose negation cannot be inferred. In formal terms:

$$(A, \leq) \Vdash^{\mathcal{A}} \alpha \quad \text{iff} \quad (A, \leq) \Vdash^{\exists} \alpha \text{ and } (A, \leq) \nVdash^{\exists} \neg\alpha.$$

In the remaining paper, the symbol x will be used to refer to one of the entailment principles denoted by the symbols $\exists$, $\forall$ and $\mathcal{A}$.

We now turn to the formalization of approximate nonmonotonic entailment. The idea is to approximate a standard nonmonotonic relation, say $\Vdash^{x}$, by means of two dual families of relations $\Vdash^{x}_{\Box}$ and $\Vdash^{x}_{\Diamond}$, the first one being sound, while the second one being complete with respect to $\Vdash^{x}$. The notions of *approximate existential entailment* and *approximate universal entailment* are defined as follows:

$$\begin{array}{lll}
(A, \leq, S) \Vdash^{\exists}_{\Box} \alpha & \text{iff} & \exists B \in \Box(A, \leq, S) \text{ such that } \models \Box_S (B \supset \alpha), \\
(A, \leq, S) \Vdash^{\exists}_{\Diamond} \alpha & \text{iff} & \exists B \in \Diamond(A, \leq, S) \text{ such that } \models \Diamond_S (B \supset \alpha), \\
(A, \leq, S) \Vdash^{\forall}_{\Box} \alpha & \text{iff} & \forall B \in \Box(A, \leq, S), \models \Box_S (B \supset \alpha), \\
(A, \leq, S) \Vdash^{\forall}_{\Diamond} \alpha & \text{iff} & \forall B \in \Diamond(A, \leq, S), \models \Diamond_S (B \supset \alpha).
\end{array}$$

Technically, the entailment relations $\Vdash^{\exists}_{\Box}$ and $\Vdash^{\forall}_{\Box}$ are specified in terms of necessity operators, while $\Vdash^{\exists}_{\Diamond}$ and $\Vdash^{\forall}_{\Diamond}$ are defined in terms of possibility operators. The relations of *approximate argumentative entailment* use both necessity and possibility operators. They are formalized as follows:

$$\begin{array}{lll}
(A, \leq, S) \Vdash^{\mathcal{A}}_{\Box} \alpha & \text{iff} & (A, \leq, S) \Vdash^{\exists}_{\Box} \alpha \text{ and } (A, \leq, S) \nVdash^{\exists}_{\Diamond} \neg\alpha, \\
(A, \leq, S) \Vdash^{\mathcal{A}}_{\Diamond} \alpha & \text{iff} & (A, \leq, S) \Vdash^{\exists}_{\Diamond} \alpha \text{ and } (A, \leq, S) \nVdash^{\exists}_{\Box} \neg\alpha.
\end{array}$$

We are now in position to provide a specification tool for approximate nonmonotonic reasoning. Formally, an approximate nonmonotonic reasoner can be defined as a function that takes as input a prioritized knowledge base $(A, \leq)$, an entailment principle x, an acceptable parameter S, and a proposition α (i.e. the query) and returns as output "yes" if $(A, \leq, S) \Vdash^{x}_{\Box} \alpha$, "no" if $(A, \leq, S) \nVdash^{x}_{\Diamond} \alpha$, and "maybe" otherwise.

The nonmonotonic reasoning process can be shown anytime by using an increasing sequence of parameters ($S_0 = \varnothing \cdots \subset S_k \cdots \subset S_n = P$) that approximate the problem of deciding whether $(A, \leq) \Vdash^x \alpha$ holds, or not, by means of two dual families of entailment tests $(A, \leq, S_k) \Vdash^x_\Box \alpha$ and $(A, \leq, S_k) \Vdash^x_\Diamond \alpha$. If the reasoner returns "yes" for a given index k, then $(A, \leq) \Vdash^x \alpha$ holds. On the other hand, if the reasoner answers "no" for a given k, then $(A, \leq) \Vdash^x \alpha$ does not hold. These considerations are clarified by the following properties.

Theorem 3.4. *For any prioritized knowledge base $(A, \leq)$, any proposition α and any acceptable parameters S and S' such that $S \subseteq S'$,*

$$\text{if } (A, \leq, S) \Vdash^\exists_\Box \alpha \quad \text{then} \quad (A, \leq, S') \Vdash^\exists_\Box \alpha \text{ and } (A, \leq) \Vdash^\exists \alpha, \quad [1]$$

$$\text{if } (A, \leq, S) \nVdash^\exists_\Diamond \alpha \quad \text{then} \quad (A, \leq, S') \nVdash^\exists_\Diamond \alpha \text{ and } (A, \leq) \nVdash^\exists \alpha. \quad [2]$$

Proof. Let us examine part [1]. We begin to focus on the first implication. Suppose that $(A, \leq, S) \Vdash^\exists_\Box \alpha$ holds. Then, there exists a base $B \in \Box(A, \leq, S)$ such that $\models \Box_S (B \supset \alpha)$ holds. By application of lemma 3.1[1], there exists a base $C \in \Box(A, \leq, S')$ such that $B \subseteq C$. By using the monotonicity property of conjunction, it follows that $\models \Box_S (C \supset \alpha)$. By application of theorem 2.1[1], it follows that $\models \Box_{S'} (C \supset \alpha)$. Therefore, we obtain $(A, \leq, S') \Vdash^\exists_\Box \alpha$, as desired. Now we turn to the second implication. Suppose that $(A, \leq, S) \Vdash^\exists_\Box \alpha$ holds. Since $S \subseteq P$, it follows that $(A, \leq, P) \Vdash^\exists_\Box \alpha$. Moreover, we know that for any proposition β, the sentence $\Box_P \beta \equiv \beta$ is a validity of our logic. So, it follows that $\Box(A, \leq, P) = \triangle(A, \leq)$. Moreover, we have $\models \Box_P (A \supset \alpha)$ iff $\models A \supset \alpha$. Therefore, we obtain $(A, \leq) \Vdash^\exists \alpha$, as desired. A dual strategy holds for part [2]. □

Theorem 3.5. *For any prioritized clausal knowledge base $(A, \leq)$, any proposition α and any acceptable parameters S and S' such that $S \subseteq S'$,*

$$\text{if } (A, \leq, S) \Vdash^\forall_\Box \alpha \quad \text{then} \quad (A, \leq, S') \Vdash^\forall_\Box \alpha \text{ and } (A, \leq) \Vdash^\forall \alpha, \quad [1]$$

$$\text{if } (A, \leq, S) \nVdash^\forall_\Diamond \alpha \quad \text{then} \quad (A, \leq, S') \nVdash^\forall_\Diamond \alpha \text{ and } (A, \leq) \nVdash^\forall \alpha. \quad [2]$$

Proof. We only examine the first implication of part [1]. Suppose we are given the assertions $(A, \leq, S) \Vdash^\forall_\Box \alpha$ and $(A, \leq, S') \nVdash^\forall_\Box \alpha$. We show that this leads to a contradiction. From the second assertion, there exists a base $B \in \Box(A, \leq, S')$ such that we have $\not\models \Box_{S'} (B \supset \alpha)$. By contraposition of theorem 2.1[1], it follows that $\not\models \Box_S (B \supset \alpha)$. Moreover, since $B \in \Box(A, \leq, S')$, by application of lemma 3.2[1], there exists a base $C \in \Box(A, \leq, S)$ such that $C \subseteq B$. By the monotonicity property of conjunction, it follows that $\not\models \Box_S (C \supset \alpha)$. Therefore, $(A, \leq, S) \Vdash^\forall_\Box \alpha$ does not hold, hence contradiction. □

Theorem 3.6. *For any prioritized knowledge base* $(A, \leq)$*, any proposition* α *and any acceptable parameters* S *and* S' *such that* $S \subseteq S'$*,*

$$\text{if } (A, \leq, S) \Vdash^{\mathcal{A}}_{\Box} \alpha \quad \text{then} \quad (A, \leq, S') \Vdash^{\mathcal{A}}_{\Box} \alpha \text{ and } (A, \leq) \Vdash^{\mathcal{A}} \alpha, \quad [1]$$
$$\text{if } (A, \leq, S) \nVdash^{\mathcal{A}}_{\Diamond} \alpha \quad \text{then} \quad (A, \leq, S') \nVdash^{\mathcal{A}}_{\Diamond} \alpha \text{ and } (A, \leq) \nVdash^{\mathcal{A}} \alpha. \quad [2]$$

Proof. We only examine the first implication of part [1]. Suppose that we are given $(A, \leq, S) \Vdash^{\mathcal{A}}_{\Box} \alpha$. Then, $(A, \leq, S) \Vdash^{\exists}_{\Box} \alpha$ and $(A, \leq, S) \nVdash^{\exists}_{\Diamond} \neg\alpha$. From the first assertion and by application theorem 3.4[1], it follows that $(A, \leq, S') \Vdash^{\exists}_{\Box} \alpha$. In a dual way, from the second assertion and by application of theorem 3.4[2], it follows that $(A, \leq, S') \nVdash^{\exists}_{\Diamond} \neg\alpha$. Thus, we obtain $(A, \leq, S') \Vdash^{\mathcal{A}}_{\Box} \alpha$, as desired. $\Box$

3.3. *Computational properties*

We now turn to computational considerations. For this, recall that coherence-based reasoning is characterized by two interacting sources of complexity, namely, propositional satisfiability and the selection of preferred consistent subsets. We intend to state that, in approximate nonmonotonic reasoning, both sources of complexity are bounded by the same resource parameter S. This statement is characterized by the next lemma and the complexity result below.

Lemma 3.7. *For any prioritized knowledge base* $(A, \leq)$ *and any acceptable parameter* S*, the cardinality of the sets* $\Box(A, \leq, S)$ *and* $\Diamond(A, \leq, S)$ *is bounded by* $2^{|S|}$.

Proof. We only consider the set $\Box(A, \leq, S)$ since a dual argument holds for the set $\Diamond(A, \leq, S)$. We split the demonstration in two parts. We first examine the non-prioritized case. Let $\Box(A, S)$ be the set of maximal subsets B of A such that $\Box_S B$ is satisfiable. We show that $\Box(A, S)$ is bounded by $2^{|S|}$. Let $V^{\Box}_S$ be the set of valuations v such that for every $p \in P$, $v(p) = \{0\}$ or $v(p) = \{1\}$ if $p \in S$, and $v(p) = \{\}$ otherwise. We remark that for each base $B \in \Box(A, S)$ there exists at least one value $v \in V^{\Box}_S$ such that $v \models B$. Let f be an application from $\Box(A, S)$ to $V^{\Box}_S$ which sends each member $B \in \Box(A, S)$ to one corresponding value $f(B)$ such that $f(B) \models B$. For every distinct bases B and B' in $\Box(A, S)$, $\Box_S (B \cup B')$ is unsatisfiable. So, $f(B)$ and $f(B')$ are distinct and hence, f is injective. Therefore, we have $|\Box(A, S)| \leq |V^{\Box}_S|$. Since $|V^{\Box}_S| = 2^{|S|}$, it follows that $|\Box(A, S)|$ is bounded by $2^{|S|}$. The worst case is obtained when A is set of all literals defined from S. In this case, f is bijective, and hence we obtain $|\Box(A, S)| = 2^{|S|}$.

We now turn to the prioritized case. Let g be a mapping inductively defined as follows: for $i = n$, $g(A_i, S) = \Box(A_i, S)$, and for $i : 1 \leq i < n$, $g(A_i, \cdots, A_n, S) = \{C \in \Box(A_i \cup \cdots \cup A_n, S) : \exists B \in g(A_{i+1}, \cdots, A_n, S) \text{ and } B \subseteq C\}$. By construction, it is easy to verify that $g(A_1, \cdots, A_n, S) = \Box(A, \leq, S)$. Moreover, we have $g(A_1, \cdots, A_n, S) \subseteq \Box(A, S)$. Hence, $|\Box(A, \leq, S)|$ is bounded by $2^{|S|}$. $\Box$

Theorem 3.8. *For any prioritized knowledge base* $(A, \leq)$*, any acceptable parameter* S *and any proposition* α*, there exists an algorithm for deciding whether* $(A, \leq, S) \Vdash_{\Box}^{x} \alpha$ *holds and* $(A, \leq, S) \Vdash_{\Diamond}^{x} \alpha$ *holds which runs in* $O((|A| + |\alpha|) \cdot 2^{|S|} \cdot 2^{|S|})$ *time.*

Proof. We focus on the complexity analysis of $(A, \leq, S) \Vdash_{\Box}^{\exists} \alpha$. The demonstration is analogous for the other entailment relations. If the assertion holds, then there exists a base $B \in \Box(A, \leq, S)$ such that $\models \Box_S (B \supset \alpha)$. By application of theorem 2.2, the validity test of $\Box_S (B \supset \alpha)$ is in $O((|A| + |\alpha|) \cdot 2^{|S|})$. Since there are at most $2^{|S|}$ bases B, the entailment test is in $O((|A| + |\alpha|) \cdot 2^{|S|} \cdot 2^{|S|})$ time. □

Several algorithms can be used for approximate nonmonotonic reasoning. The key difficulty lies in the consolidation operation. From this perspective, one may conceive an algorithm which takes as input a prioritized clausal base $(A, \leq)$ and computes $\Box(A, \leq, S_k)$ by means of an increasing sequence S_k. For $k = 0$, the procedure simply returns the empty base. For $k > 0$, the procedure proceeds into two steps. First, for each subset B of $\Box(A, \leq, S_{k-1})$, the procedure computes the satisfiable expansions of B that take clauses containing the literal p_k or its negation $\neg p_k$. Second, the procedure selects the maximal expansions and add them to $\Box(A, \leq, S_k)$. As far as $\Diamond(A, \leq, S_k)$ is concerned, dual considerations hold. In order to improve the algorithm, several data structures such as *set enumeration trees* [RYM 92] or *binary decision diagrams* [CAY 98] can be advocated. It is interesting to remark that such an algorithm is indeed *anytime*: by exploiting lemmas 3.1 and 3.2, the procedure can be interrupted at any step in order to evaluate the query and, in case of unsatisfactory result, it only needs to expand the maximal subsets generated in previous steps.

The correct choice of S is crucial for the usefulness of anytime consolidation. This choice may be guided by the priority ordering $\leq$. Following the acceptability condition, the parameter is constructed by selecting the atoms from the stratum of highest priority, then the atoms of the next important stratum are added, and so on. Alternatively, inside each stratum, the choice of S may be heuristic. In this case, the letters are iteratively selected to minimize the predicted number of consistent subsets, using a strategy such as the *minimal diversity heuristic*.

Example 3.9. *Consider the flat base* $A = \{a, b, c, \neg c, \neg a \vee \neg b, \neg a \vee c, \neg a \vee \neg c, \neg b \vee d\}$*. Obviously,* A *is unsatisfiable. We want to show that* $A \Vdash^{\exists} d$*. Hence, we need to find a set* S *such that* $(A, S) \Vdash_{\Box}^{\exists} d$*. The minimal diversity of the atoms* a*,* b*,* c *and* d *is* 3*,* 2*,* 4 *and* 0*, respectively. Starting with* $S = \varnothing$*, we iteratively add* d *and* b *to* S*. Based on the following results, we indeed observe that* $(A, S) \Vdash_{\Box}^{\exists} d$*.*

S	$\Box(A, S)$
$\varnothing$	$\varnothing$
$\{d\}$	$\{\{\neg b \vee d\}\}$
$\{b, d\}$	$\{\{b, \neg b \vee d\}, \{\neg a \vee \neg b, \neg b \vee d\}\}$

Example 3.10. *Suppose we are given the prioritized base $A = (A_1, A_2)$ where $A_1 = \{c, \neg d, \neg a \vee b, \neg c \vee d\}$ and $A_2 = \{a, \neg a, e\}$. We want to show that $(A, \leq) \Vdash^{\mathcal{A}} b$. So, we need to find a set S such that $(A, \leq, S) \Vdash^{\mathcal{A}}_{\Box} b$. Starting with $S = \varnothing$ and using the acceptability condition, we first add the atoms a and e and next we select b. Based on the following results, we indeed obtain $(A, \leq, S) \Vdash^{\exists}_{\Box} b$ and $(A, \leq, S) \nVdash^{\exists}_{\Diamond} \neg b$.*

S	$\Box(A, \leq, S)$	$\Diamond(A, \leq, S)$
$\varnothing$	$\varnothing$	A
$\{a, e\}$	$\{\{a, e\}, \{\neg a, e, \neg a \vee b\}\}$	$\{\{a, e\} \cup A_2, \{\neg a, e\} \cup A_2\}$
$\{a, b, e\}$	$\{\{a, e, \neg a \vee b\}, \{\neg a, e, \neg a \vee b\}\}$	$\{\{a, e\} \cup A_2, \{\neg a, e\} \cup A_2\}$

4. Discussion

In this paper, we have studied the problem of reasoning from inconsistency focusing on the so-called coherence-based approaches. One of the main drawbacks of these methods is their high computational complexity. Our aim was to provide a logical framework which tackles this difficulty through the paradigm of anytime computation. We have illustrated that the framework integrates several major features: resource-bounded reasoning, improvability and dual reasoning.

A close look at the literature in artificial intelligence and logic shows that very few investigations have addressed *together* the problems of inconsistency and intractability. On the one hand, there have been a great number of proposals for constructing systems that allow nontrivial reasoning in presence of inconsistency. Among them are paraconsistent logics [HUN 98], argumentative logics [ELV 95], knowledge merging systems [LIN 95], and the so-called coherence based approaches [PIN 92]. However, for most of the part, existing frameworks are known to be typically intractable. On the other hand, many techniques have been developed to deal with the intractability of deduction problems. The most significant approaches are logics of explicit belief [LEV 84, LAK 94], access-limited reasoning [CRA 89], approximate knowledge compilation [SEL 96] and anytime reasoning [SCH 95, DAL 98]. Unfortunately, most of these approaches do not address the issue of inconsistency handling.

An important exception is the model of *paraconsistent resource-bounded inference* recently proposed by Marquis and Porquet in [MAR 01]. The key idea is to approximate deduction in knowledge bases by using "maximally preferred" resources. Technically, given a knowledge base A, a resource parameter S and a query α, we say that α is a *consequence* of A with respect to S, if $\Box_{S'} (A \supset \alpha)$ is valid for every maximal subset S' of S such that $\Diamond_{S'} A$ is satisfiable. The model can be extended to prioritized knowledge bases, using several preference orderings defined from inclusion-based, lexicographic or possibilistic policies.

Although our framework is in the spirit of Marquis and Porquet's approach, there are significant differences. From a conceptual point of view, their model is based on the idea of "approximate paraconsistency" which does not explicitly restore consistency but instead tolerates contradictions. In contrast, our framework advocates the idea of "approximate consolidation" which progressively restores consistency. At any step, the result can be stored in a data structure which can be later used for different queries and different entailment principles. From a logical point of view, their model only covers sound approximations, while ours includes both sound approximations and complete approximations. Finally, from a computational point of view, their approach is not improvable since paraconsistent approximations are nonmonotonic in essence. In contrast, our framework is *anytime*: the consolidation phase can be shown improvable and incremental by using increasing sequences of resource parameters and by exploiting computations from previous steps.

There are various avenues of research that come out of this work. Most prominent amongst these is the empirical analysis of approximate coherence-based reasoning. To this point, some benchmarks for coherence-based reasoning have recently been proposed in [CAY 98]. In this setting, the performance of our anytime technique should be compared with standard methods. A second interesting issue is to generalize the framework to other models of inconsistency handling proposed in the literature. For example, in the domain of diagnosis, consolidation operations are sometimes defined over cardinality-based or lexicographic preference orderings [GIN 86, LEH 95]. In knowledge merging, entailment relations often use a notion of majority vote inspired from social choice theory [BOR 84, LIN 95]. For all these approaches, the anytime view remains to be explored. Finally, other forms of approximate inference should be examined. In particular, Dalal in [DAL 98] presents a family of approximate entailment operations which are based on boolean constraint propagation. However, these operations trivialize in presence of inconsistency. This opens the door for interesting extensions of Dalal's framework in the setting of coherence-base reasoning.

Acknowledgements

Thanks to Samir Chopra, Sébastien Konieczny, Thomas Meyer, and Geneviève Simonet for fruitful discussions, and to the anonymous referees for their constructive suggestions and criticisms.

5. References

[BAR 91] BARAL C., KRAUS S., MINKER J., "Combining multiple knowledge bases", *IEEE Transactions on Knowledge and Data Engineering*, vol. 3, num. 2, 1991, p. 208–220.

[BAR 92] BARAL C., KRAUS S., MINKER J., SUBRAHMANIAN V. S., "Combining knowledge bases consisting of first-order theories", *Computational Intelligence*, vol. 8, num. 1, 1992, p. 45-71.

[BEL 77] BELNAP N. D., "A useful four-valued logic", DUNN J. M., EPSTEIN G., Eds., *Modern Uses of Multiple-Valued Logic*, Reidel, Dordrecht, 1977, p. 8–37.

[BEN 93] BENFERHAT S., CAYROL C., DUBOIS D., LANG J., PRADE H., "Inconsistency management and prioritized syntax-based entailment", *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, 1993, p. 640-645.

[BEN 95] BENFERHAT S., DUBOIS D., PRADE H., "How to infer from inconsistent beliefs without revising ?", *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, 1995, p. 1449–1455.

[BEN 99] BENFERHAT S., DUBOIS D., PRADE H., "An overview of inconsistency-tolerant inferences in prioritized knowledge bases", DUBOIS D., KLEMENT E. P., PRADE H., Eds., *Fuzzy Sets, Logics and Reasoning About Knowledge*, vol. 15 of *Applied Logic Series*, p. 395–418, Kluwer Academic Publishers, 1999.

[BOR 84] BORGIDA A., IMIELINSKI T., "Decision making in committees: A framework for dealing with inconsistency and nonmonotonicity", *Non-Monotonic Reasoning Workshop*, AAAI Press, 1984, p. 21–32.

[BRE 89] BREWKA G., "Preferred subtheories: An extended logical framework for default reasoning", *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, 1989, p. 1043–1048.

[CAD 96] CADOLI M., SCHAERF M., "On the complexity of entailment in propositional multivalued logics", *Annals of Mathematics and Artificial Intelligence*, vol. 18, num. 1, 1996, p. 29–50.

[CAY 98] CAYROL C., LAGASQUIE-SCHIEX M. C., "Nonmonotonic reasoning: from complexity to algorithms", *Annals of Mathematics and Artificial Intelligence*, vol. 22, 1998, p. 207-236.

[CRA 89] CRAWFORD J. M., KUIPERS B., "Towards a theory of access-limited logic for knowledge representation", *Proceedings of the First International Conference on Principles of Knowledge Representation and Reasoning*, 1989, p. 67–78.

[DAL 98] DALAL M., "Anytime clausal reasoning", *Annals of Mathematics and Artificial Intelligence*, vol. 22, num. 3-4, 1998, p. 297–318.

[DEA 88] DEAN T., BODDY M., "An analysis of time-dependent planning", *Proceedings of the Seventh National Conference on Artificial Intelligence*, 1988, p. 49–54.

[DEC 94] DECHTER R., RISH I., "Directional resolution: The Davis-Putnam procedure, revisited", *Proceedings of the Fourth International Conference on Principles of Knowledge Representation and Reasoning*, 1994, p. 134–145.

[ELV 95] ELVANG-GØRANSSON M., HUNTER A., "Argumentative logics: Reasoning with classical inconsistent information", *IEEE Transactions on Data and Knowledge Engineering*, vol. 16, 1995, p. 125–145.

[FAG 83] FAGIN R., ULLMAN J. D., VARDI M. Y., "On the semantics of updates in databases", *Proceedings of the Second ACM SIGACT-SIGMOD Symposium on Principles of Database Systems*, 1983, p. 352–365.

[FAG 95] FAGIN R., HALPERN J. Y., VARDI M. Y., "A nonstandard approach to the logical omniscience problem", *Artificial Intelligence*, vol. 79, num. 2, 1995, p. 203–240.

[GIN 86] GINSBERG M. L., "Counterfactuals", *Artificial Intelligence*, vol. 30, num. 1, 1986, p. 35–79.

[HUN 98] HUNTER H., "Paraconsistent logics", GABBAY D., SMETS P., Eds., *Handbook of Defeasible Reasoning and Uncertainty Management Systems*, p. 11–36, Kluwer Academic Publishers, 1998.

[KOR 98] KORICHE F., "A logic for anytime deduction and anytime compilation", DIX J., FURBACH U., DEL CERRO L., Eds., *Logics in Artificial Intelligence*, vol. 1489 of *Lecture Notes in Artificial Intelligence*, p. 324–342, Springer Verlag, 1998.

[KOR 01] KORICHE F., "A logic for approximate first-order reasoning", FRIBOURG L., Ed., *Computer Science Logic*, vol. 2142 of *Lecture Notes in Computer Science*, p. 247–261, Springer Verlag, 2001.

[LAK 94] LAKEMEYER G., "Limited reasoning in first-order knowledge bases", *Artificial Intelligence*, vol. 71, 1994, p. 213-255.

[LEH 95] LEHMANN D., "Another perspective on default reasoning", *Annals of mathematics and artificial intelligence*, vol. 15, 1995, p. 61–82.

[LEV 84] LEVESQUE H., "A logic of implicit and explicit belief", *Proceedings of the Sixth National Conference on Artificial Intelligence*, 1984, p. 198–202.

[LIN 95] LIN J., "Frameworks for dealing with conflicting information and applications", PhD thesis, Graduate Department of Computer Science. Univesity of Toronto, Toronto, Canada, 1995.

[MAR 01] MARQUIS P., PORQUET N., "Resource-bounded inference from inconsistent belief bases", *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence*, 2001, p. 103–108.

[MAS 98] MASSACCI F., "Anytime approximate modal reasoning", *Proceedings of the Fifteenth National Conference on Artificial Intelligence and of the Tenth Conference on Innovative Applications of Artificial Intelligence*, 1998, p. 274–279.

[NEB 91] NEBEL B., "Belief revision and default reasoning: Syntax-based approaches", *Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning*, 1991, p. 417–428.

[NEB 98] NEBEL B., "How hard is it to revise a belief base?", GABBAY D., SMETS P., Eds., *Handbook of Defeasible Reasoning and Uncertainty Management Systems*, vol. 3, p. 77–145, Kluwer Academic Publishers, 1998.

[PER 97] PERLIS D., "Sources of, and exploiting, inconsistency: preliminary report", *Journal of Applied Non-Classical Logics*, vol. 7, num. 1–2, 1997, p. 25–75.

[PIN 92] PINKAS G., LOUI R. P., "Reasoning from inconsistency: A taxonomy of principles for resolving conflict", *Proceedings of the Third International Conference on Principles of Knowledge Representation and Reasoning*, 1992, p. 709–719.

[RES 70] RESCHER N., MANOR R., "On inference from inconsistent premises", *Theory and Decision*, vol. 1, 1970, p. 179-217.

[RYM 92] RYMON R., "Search through systematic set enumeration", *Proceedings of the Third International Conference in Principles of Knowledge Representation and Reasoning*, 1992, p. 539–550.

[SCH 95] SCHAERF M., CADOLI M., "Tractable reasoning via approximation", *Artificial Intelligence*, vol. 74, 1995, p. 249-310.

[SCH 01] SCHUURMANS D., SOUTHEY F., "Local search characteristics of incomplete SAT procedures", *Artificial Intelligence*, vol. 132, num. 2, 2001, p. 121-150.

[SEL 96] SELMAN B., KAUTZ H. A., "Knowledge compilation and theory approximation", *Journal of the ACM*, vol. 43, num. 2, 1996, p. 193–224.

[ZHA 00] ZHANG H., STICKEL M. E., "Implementing the Davis-Putnam method", *Journal of Automated Reasoning*, vol. 24, num. 1/2, 2000, p. 277-296.

[ZIL 96] ZILBERSTEIN S., "Using anytime algorithms in intelligent systems", *AI Magazine*, vol. 17, num. 3, 1996, p. 73–83.