On Anytime Coherence-Based Reasoning

Frédéric Koriche

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

Abstract. A great deal of research has been devoted to nontrivial reasoning in inconsistent knowledge bases. Coherence-based approaches proceed by a consolidation operation which selects several consistent subsets of the knowledge base and an entailment operation which uses classical implication on these subsets in order to conclude. An important advantage of these formalisms is their flexibility : consolidation operations can take into account the priorities of declarations stored in the base, and different entailment operations can be distinguished according to the cautiousness of reasoning. However, one of the main drawbacks of these approaches is their high computational complexity. The purpose of our study is to define a logical framework which handles this difficulty by introducing the concepts of anytime consolidation and anytime entailment. The framework is semantically founded on the notion of resource which captures both the accuracy and the computational cost of anytime operations. Moreover, a stepwise procedure is included for improving approximations. Finally, both sound approximations and complete ones are covered. Based on these properties, we show that an anytime view of coherence-based reasoning is tenable.

1 Introduction

A great deal of research has been devoted to nontrivial reasoning from inconsistency. This problems arises in a number of areas in artificial intelligence, e.g., in merging knowledge bases [1, 2], defeasible reasoning [16, 5] and belief revision [13, 14]. Most of the research in this issue is influenced by work in nonmonotonic reasoning, in particular by Nebel [13, 14], Pinkas and Loui [15], and Benferhat and his colleagues [3, 4], who developed the so-called *coherence-based approaches.* The main idea of these techniques is to start with a knowledge base and to apply two successive mechanisms, namely, a *consolidation operation* which generates and selects several consistent subsets of the base and an *entailment relation* which uses classical logic on the consistent subsets in order to conclude.

As noticed by Nebel in [14], 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 declarations stored in the knowledge base. For example, if priorities attached to declarations 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 declarations are discarded and what declarations are going to stay. In an orthogonal way, different classes of entailment operations can be distinguished according to the cautiousness of reasoning. For example, the following kind of entailment is considered in [2, 5] : "a knowledge base A entails the declaration α if, and only if, α is classically inferred by all the preferred consistent subsets of A". A taxonomy of entailment operations, from credulous to skeptical ones, can be found in [15].

Unfortunately, one of the main drawbacks of coherence-based approaches is their high computational complexity. As stated in [7], 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 satisfiability and the selection of preferred consistent subsets. For this reason, one cannot expect to arrive at a polynomial algorithm when eliminating only one source, e.g., by restricting the base to Horn logic.

Anytime reasoning is a technique which is used in many areas of artificial intelligence to deal with the computational intractability of problems [20]. This paradigm extends the traditional notion of reasoner by allowing it to return many possible answers to any given query. An original method, primarily due to Schaerf and Cadoli in [17], and recently generalized in [9], has received a great deal of interest in the knowledge representation community. The basic idea of the method is to define a family of inference relations 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. Several extensions of this method have been proposed in the fields of modal logics [12] and first-order logic [10]. However, despite few exceptions (e.g. [6]), most of the studies in anytime reasoning have concentrated to the monotonic case. In particular, it is necessary to make formal steps in the direction of coherence-based reasoning.

The purpose of this paper is to develop a logic oriented framework for anytime coherence-based reasoning. Our formalism is based on a multi-modal propositional logic, presented in [9], and used to specify *anytime monotonic reasoners*. In this study, we extend our previous work in order to specify *anytime nonmonotonic reasoners*. Starting from a knowledge base A and a preordering on A, we introduce the notion of *anytime consolidation*, an operation which generates and selects approximate preferred consistent subsets of A. Then, we define three classes of *anytime entailment relations*, which respectively incorporate the credulous principle, the skeptical 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 *incremental 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. Finally, section 4 suggests some topics for future research.

2 Anytime monotonic reasoning

In this section, we focus on the formalization of anytime monotonic reasoners. For this purpose, we present a propositional logic, named **ARL**, for anytime reasoning. We begin to define its syntax, next we examine its semantics and then we present some interesting properties of the logic.

2.1 Syntax

Throughout this paper, we consider a nonempty and finite set of atomic propositions (atoms for short) P. The language of *declarations* is the smallest set built from P and closed off under the connectives $\wedge$, $\vee$ and $\neg$. The connective $\supset$ is defined in terms of $\neg$ and $\vee$; that is, $\alpha \supset \beta$ is an abbreviation of $\neg\alpha \vee \beta$. Given a declaration α, the set of atoms that occur in α is denoted $P(\alpha)$. A *literal* is an atom or its negation and 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 [17], the concept of *computational resource* is captured by a parameter S, a subset of P. Intuitively, the parameter S corresponds to a limited and controlled exploration in the space of possibilities defined from P.

The main contribution of the logic relies on two families of modalities $\Box_S$ and $\Diamond_S$, defined for each subset S of P. The operator $\Box_S$ is to capture sound but incomplete inference and $\Diamond_S$ to capture complete but unsound inference. The language of **ARL** is defined by the smallest set of *sentences* built from the following rules: if α is a declaration then α is a sentence, if α and β are sentences then $\neg\alpha$, $\alpha \wedge \beta$ and $\alpha \vee \beta$ are sentences, and if α is a declaration and S is a subset of P then $\Box_S\,\alpha$ and $\Diamond_S\,\alpha$ are sentences. 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

In the context of limited reasoning, the four valued semantics first proposed by Belnap and notably studied in [11] meets our needs. The domain T of truth values is the powerset of $\{0, 1\}$. So, in the logic **ARL**, sentences can be valued to be true, false, both, or neither. Based on this structure, we define a *valuation* as a total function v from P to T. The space of valuations generated from P is denoted V. A *possible world* is a valuation which maps every atom p of P into $\{1\}$ or $\{0\}$. The space of possible worlds generated from P is denoted W.

The notion of resource is semantically represented by an equivalence relation between valuations. Given a parameter S, we say that two valuations v and v' are *S-equivalent* and write $v \sim_S v'$, iff for every atom $p \in P$, if $p \in S$ then $v(p) = v'(p)$. Intuitively, a relation of S-equivalence 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. The "coarsest" partition is obtained when S is the empty set; in this case, $\sim_S$ is the total relation over V. Conversely, the "finest" partition is given when S is the set P; in this case $\sim_S$ is the identity relation over V.

An *interpretation* of **ARL** consists of a *truth support relation* $\models_1$ and a *falsity support relation* $\models_0$ inductively defined by the following conditions:

$$v \models_1 p \text{ iff } 1 \in v(p), \quad v \models_0 p \text{ iff } 0 \in v(p), \tag{1}$$

$$v \models_1 \neg\alpha \text{ iff } v \models_0 \alpha, \quad v \models_0 \neg\alpha \text{ iff } v \models_1 \alpha, \tag{2}$$

$$v \models_1 \alpha \wedge \beta \text{ iff } v \models_1 \alpha \text{ and } v \models_1 \beta, \quad v \models_0 \alpha \wedge \beta \text{ iff } v \models_0 \alpha \text{ or } v \models_0 \beta, \tag{3}$$

$$v \models_1 \alpha \vee \beta \text{ iff } v \models_1 \alpha \text{ or } v \models_1 \beta, \quad v \models_0 \alpha \vee \beta \text{ iff } v \models_0 \alpha \text{ and } v \models_0 \beta, \tag{4}$$

$$v \models_1 \Box_S \alpha \text{ iff } \forall v' \in V, \text{ if } v \sim_S v' \text{ then } v' \models_1 \alpha, \quad v \models_0 \Box_S \alpha \text{ iff } v \not\models_1 \Box_S \alpha, \tag{5}$$

$$v \models_1 \Diamond_S \alpha \text{ iff } \exists v' \in V \text{ such that } v \sim_S v' \text{ and } v' \models_1 \alpha, \quad v \models_0 \Diamond_S \alpha \text{ iff } v \not\models_1 \Diamond_S \alpha. \tag{6}$$

A sentence α is *satisfiable* iff there exists a possible world w such that $w \models_1 \alpha$. We say that α is *valid*, and write $\models \alpha$, iff for every $w \in W$, $w \models_1 \alpha$ holds. Given two sentences α and β, we say that β is a *logical consequence* of α iff $\models \alpha \supset \beta$ holds. A sound and complete axiomatization for **ARL** can be found in [9].

2.3 Properties

After an excursion into the logic **ARL**, we now focus on its main properties. In this purpose, we specify an *anytime monotonic reasoner* as a function that takes in input a knowledge base A, parameter S and a declaration α, and returns in output "yes" if $\models \Box_S (A \supset \alpha)$, "no" is $\not\models \Diamond_S (A \supset \alpha)$ and "unknown" otherwise.

Interestingly, our model can be shown *incremental* and *dual*. Specifically, the reasoning process 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 α is a logical consequence of A, or not, by means of two dual families of tests $\models \Box_{S_k} (A \supset \alpha)$ and $\models \Diamond_{S_k} (A \supset \alpha)$. If the reasoner returns "yes" using any operator $\Box_{S_k}$ then α is a consequence of A. Dually, if the reasoner answers "no" using any operator $\Diamond_{S_k}$ then α is not a consequence of A. 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.

Theorem 1. *For any declaration α and any parameters S and S' s.t. $S \subseteq S'$,*

$$\text{if } \models \Box_S\,\alpha \text{ then } \models \Box_{S'}\,\alpha \text{ and hence } \models \alpha, \tag{1}$$

$$\text{if } \not\models \Diamond_S\,\alpha \text{ then } \not\models \Diamond_{S'}\,\alpha \text{ and hence } \not\models \alpha. \tag{2}$$

Lemma 1. *For any declaration α,*

$$\models \Box_S\,\alpha \text{ iff } \Diamond_S\,\neg\alpha \text{ is unsatisfiable}, \tag{1}$$

$$\not\models \Diamond_S\,\alpha \text{ iff } \Box_S\,\neg\alpha \text{ is satisfiable}. \tag{2}$$

Theorem 2. *For any declaration α and any S, there is an algorithm for deciding whether $\Box_S\,\alpha$ is satisfiable and $\Diamond_S\,\alpha$ is satisfiable which runs in $O(|\alpha| \cdot 2^{|S|})$.*

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-phases 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 occurs in S. Dually, in the scope of $\Box_S$, the algorithm eliminates in any clause of α all literals whose atom occurs 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 enumeration [19] 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 [18] can be exploited if we concentrate on the satisfiability of $\Box_S\,\alpha$. In a nutshell, 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, anytime 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 [8]. 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 1. Let $A = \{(\neg a \vee b \vee c), (a \vee b \vee \neg d), (a \vee \neg b \vee d), (\neg a \vee \neg b \vee c)\}$. 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. Starting with $S = \varnothing$ and using the minimal diversity heuristic, we gradually add c and a to S. This is sufficient for proving that A is satisfiable.

Example 2. 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. Using the maximal diversity strategy, we iteratively add a, b and c to S. This is sufficient for proving that $(A \wedge a \wedge \neg c)$ is unsatisfiable. So, $a \supset c$ is indeed a logical consequence of A.

3 Anytime nonmonotonic reasoning

In this section, we extend the concepts developed so far to the formalization of anytime nonmonotonic reasoners. In the setting suggested by our approach, these systems are defined in terms of *anytime consolidation* and *anytime entailment*. The quality of result of each operation depends on the computational resources that have been spent. We begin to define the concept of anytime consolidation, next we present three classes of anytime entailment, and then we examine the computational properties of our framework.

3.1 Anytime consolidation operations

As considered for instance in [13, 14], a "standard" consolidation operation starts from a knowledge base and a priority ordering on this base and selects the preferred consistent subsets of the base. The purpose of "anytime" 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 declarations of highest priority and A_n those of lowest 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. [3]). In this study, we focus on the *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 $\forall j : 1 \leq j < i, B \cap A_j = C \cap A_j$. By extension, $B \preceq C$ iff $B \prec C$ or $B = C$. Based on these considerations, the standard consolidation operation, denoted $\triangle$, is defined as follows:

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

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 $\exists i : S \cap P(A_i) \neq 0$ then $\forall j : 1 \leq j < i$, $P(A_j) \subseteq S$. 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. In particular, it is interesting to remark that if the structure $(A, \leq)$ is flat, then every subset of P is acceptable for $(A, \leq)$. The anytime view of consolidation is realized by parameterizing the operation $\triangle$ by means of two families of operations $\Box$ and $\Diamond$, the first one being sound, while the second one being complete with respect to standard consolidation. The corresponding *anytime consolidation operations* are defined as follows:

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

The following lemmas capture important properties of anytime consolidation. They will be frequently used in the remaining paper.

Lemma 2. *For any prioritized knowledge base $(A, \leq)$ and any acceptable parameters S and S' such that $S \subseteq S'$:*

$$\forall B \in \Box(A, \leq, S) \;\; \exists C \in \Box(A, \leq, S') \;\; \text{such that } B \subseteq C, \tag{1}$$

$$\forall B \in \Diamond(A, \leq, S') \;\; \exists C \in \Diamond(A, \leq, S) \;\; \text{such that } B \subseteq C. \tag{2}$$

Proof. Let us examine part (1). Assume that there exists a knowledge 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. If $B \in \Box(A, \leq, S)$ then $\Box_S B$ is satisfiable. By application of theorem 1 and lemma 1, it follows that $\Box_{S'} B$ is satisfiable. Since $B \notin \Box(A, \leq, S')$ there must exist a base $C \in \Box(A, \leq, S')$ such that $B \prec C$. Therefore, $\exists i : B \cap A_i \subset C \cap A_i$ and $\forall j : 1 \leq j < i, B \cap A_j = C \cap A_j$. By assumption, we know that $B \not\subset C$. So, $\exists k > i : B \cap A_k \not\subset C \cap A_k$. Thus, it follows that $B \cap A_k \neq \varnothing$. Since $\Box_S B$ is satisfiable, we must have $S \cap P(A_k) \neq \varnothing$. By acceptability condition, it follows that $\forall k' < k$, $P(A_{k'}) \subseteq S$. Let B' denotes the set $\bigcup \{C \cap A_{k'} : k' < k\}$. Obviously, $\Box_S B'$ is satisfiable. Moreover, it is clear that $B \prec B'$. Therefore, we obtain $B \notin \Box(A, \leq, S)$, but this contradicts the initial hypothesis. A dual argument applies to part (2).

Lemma 3. *For any knowledge base A, any clause α and any parameters S and S' such that $S \subseteq S'$,*

1. *if $\Box_{S'} A$ is satisfiable and $\Box_{S'} A \cup \{\alpha\}$ is unsatisfiable, then there exists a subset B of A such that $\Box_S B$ is satisfiable and $\Box_S B \cup \{\alpha\}$ is unsatisfiable.*
2. *if $\Diamond_S A$ is satisfiable and $\Diamond_S A \cup \{\alpha\}$ is unsatisfiable, then there exists a subset B of A such that $\Diamond_{S'} B$ is satisfiable and $\Diamond_{S'} B \cup \{\alpha\}$ is unsatisfiable.*

Proof. Let us examine part (1), If $\Box_S \alpha$ is unsatisfiable then $B = \varnothing$ and we have demonstrated the property. Now, suppose that $\Box_S \alpha$ is satisfiable. Thus, there exists a literal l in α such that its atom belongs to S. Moreover, since $\Box_{S'} A$ is satisfiable and $\Box_{S'} A \cup \{\alpha\}$ is unsatisfiable, there exists a clause β in A such that the negation of l belongs to β. So, $\Box_S \beta$ is satisfiable. Let γ denotes the resolvent of α and β. If $\Box_S \gamma$ is unsatisfiable, then $B = \{\beta\}$. Otherwise, there exists a literal l' in γ such that its atom belongs to S. Thus, there exists a clause β' in A that contains the negation of l'. Since γ does not contain any occurrence of l, it is clear that $P(l') \cap P(l) = \varnothing$. Therefore, $\Box_S \beta \wedge \beta'$ is satisfiable. Let γ' denotes the resolvent of γ and β'. If $\Box_S \gamma'$ is unsatisfiable then $B = \{\beta, \beta'\}$. Otherwise, we iteratively apply the same method until we obtain all the clauses of A. In this case, $\Box_S A$ is satisfiable. An analogous strategy applies to part (2).

Lemma 4. *For any prioritized knowledge base $(A, \leq)$ and any acceptable parameters S and S' such that $S \subseteq S'$:*

$$\forall B \in \Box(A, \leq, S') \;\; \exists C \in \Box(A, \leq, S) \;\; \text{such that } C \subseteq B, \tag{1}$$

$$\forall B \in \Diamond(A, \leq, S) \;\; \exists C \in \Diamond(A, \leq, S') \;\; \text{such that } C \subseteq B. \tag{2}$$

Proof. Let us examine part (1). Suppose we have $B \in \Box(A, \leq, S')$. If $B = A$ the demonstration is straightforward. Now, suppose that $B \subset A$. We know that $\Box_{S'} B$ is satisfiable. Moreover, for every clause β in A/B, $\Box_{S'} B \cup \{\beta\}$ is unsatisfiable. Let α denotes the clause $\bigvee\{\beta : \beta \in A/B\}$. Obviously, $\Box_{S'} B \cup \{\alpha\}$ is unsatisfiable. By application of lemma 3, there exists a subset B' of B such that $\Box_S B$ is satisfiable and $\Box_S B \cup \{\alpha\}$ is unsatisfiable. Clearly enough, B' can be extended to a set C such that $C \in \Box(B, \leq, S)$. Suppose that $C \notin \Box(A, \leq, S)$. Then, there must exist a set $C' \in \Box(A, \leq, S)$ such that $C \prec C'$. Therefore, $\exists i : C \cap A_i \subset C' \cap A_i$ and $\forall j : 1 \leq j < i, C \cap A_j = C' \cap A_j$. Clearly, $C \not\subset C'$. Suppose not. In this case, $\exists k > i$, such that $C \cap A_k \not\subset C' \cap A_k$. Thus, it follows that $C \cap A_k \neq \varnothing$. Since $\Box_S C$ is satisfiable, $P(A_k) \cap S \neq \varnothing$. Therefore, $\forall k' < k$, $P(A_{k'}) \subseteq S$. It follows that $\forall k' < k, C \cap A_{k'} = B \cap A_{k'}$. Thus, we obtain $B \cap A_i \subset C' \cap A_i$. So, $B \prec C'$. Since $\Box_{S'} C'$ is satisfiable, $B \notin \Box(A, \leq, S')$, but this contradicts the initial hypothesis. So, we can state that $C \subset C'$. Thus, $\exists \beta \in A/B$ such that $C \cup \{\beta\} \subseteq C'$. However, $\Box_S C \cup \{\beta\}$ is unsatisfiable. Therefore $C' \notin \Box(A, \leq, S)$. So, $C \in \Box(A, \leq, S)$. Moreover, since $C \in \Box(B, \leq, S)$, we obtain $C \subseteq B$, as desired. A dual argument applies to part (2).

3.2 Anytime entailment operations

In the setting of coherence based-reasoning, a "standard" entailment relation takes in input a collection of preferred consistent subsets and returns in output a set of cautious conclusions. A taxonomy of numerous entailment principles has been established in [15] 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 next we examine their corresponding approximations.

The first two entailment principles, introduced by Rescher and Manor in [16], are the most commonly used in presence of contradictory knowledge bases (see e.g. [2, 3]). They can be respectively described in the following way:

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

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 [4, 15], 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:

$$(A, \leq) \Vdash^{\mathcal{A}} \alpha \ \text{ iff } \ (A, \leq) \Vdash^{\exists} \alpha \text{ and } (A, \leq) \not\Vdash^{\exists} \neg\alpha.$$

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 anytime view of entailment relations. 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 *anytime existential entailment* and *anytime universal entailment* are defined as follows:

$$(A, \leq, S) \Vdash^\exists_\Box \alpha \text{ iff } \exists B \in \Box(A, \leq, S) \text{ such that } \models \Box_S (B \supset \alpha),$$
$$(A, \leq, S) \Vdash^\exists_\Diamond \alpha \text{ iff } \exists B \in \Diamond(A, \leq, S) \text{ such that } \models \Diamond_S (B \supset \alpha),$$

$$(A, \leq, S) \Vdash^\forall_\Box \alpha \text{ iff } \forall B \in \Box(A, \leq, S), \models \Box_S (B \supset \alpha),$$
$$(A, \leq, S) \Vdash^\forall_\Diamond \alpha \text{ iff } \forall B \in \Diamond(A, \leq, S), \models \Diamond_S (B \supset \alpha).$$

The notion of *anytime argumentative entailment* is defined as follows:

$$(A, \leq, S) \Vdash^A_\Box \alpha \text{ iff } (A, \leq) \Vdash^\exists_\Box \alpha \text{ and } (A, \leq) \nVdash^\exists_\Diamond \neg\alpha,$$
$$(A, \leq, S) \Vdash^A_\Diamond \alpha \text{ iff } (A, \leq) \Vdash^\exists_\Diamond \alpha \text{ and } (A, \leq) \nVdash^\exists_\Box \neg\alpha.$$

We are now in position to provide a specification tool for anytime nonmonotonic reasoning. From this perspective, we define an *anytime nonmonotonic reasoner* as a function that takes in input a prioritized knowledge base $(A, \leq)$, an acceptable parameter S, a declaration α (i.e. the query) and an entailment principle x, and returns in output "yes" if $(A, \leq, S) \Vdash^x_\Box \alpha$, "no" if $(A, \leq, S) \nVdash^x_\Diamond \alpha$, and "unknown" otherwise. As for monotonic deduction, the nonmonotonic reasoning process can be modeled 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 $(A, \leq) \Vdash^x \alpha$ 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. *For any prioritized knowledge base $(A, \leq)$, any declaration α and any acceptable parameters S and S' such that $S \subseteq S'$,*

$$\text{if } (A, \leq, S) \Vdash^\exists_\Box \alpha \text{ then } (A, \leq, S') \Vdash^\exists_\Box \alpha \text{ and hence } (A, \leq) \Vdash^\exists \alpha, \quad (1)$$
$$\text{if } (A, \leq, S) \nVdash^\exists_\Diamond \alpha \text{ then } (A, \leq, S') \nVdash^\exists_\Diamond \alpha \text{ and hence } (A, \leq) \nVdash^\exists \alpha. \quad (2)$$

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, $\exists B \in \Box(A, \leq, S)$ such that $\models \Box_S (B \supset \alpha)$ holds. By lemma 2, we know that $\exists C \in \Box(A, \leq, S')$ such that $B \subseteq C$. By the monotonicity property of conjunction, it follows that $\models \Box_S (C \supset \alpha)$. By application of theorem 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 of part (1). As before, we assume that $(A, \leq, S) \Vdash^\exists_\Box \alpha$ holds. Since $S \subseteq P$, it follows that $(A, \leq, P) \Vdash^\exists_\Box \alpha$. By using the semantical properties of $\sim_P$, we can easily verify that $\Box(A, \leq, P) = \triangle(A, \leq)$, and that $\models \Box_P (A \supset \alpha)$ holds iff $\models A \supset \alpha$ holds. So, $(A, \leq, P) \Vdash^\exists_\Box \alpha$ is logically equivalent to $(A, \leq) \Vdash^\exists \alpha$. Therefore, it follows that $(A, \leq) \Vdash^\exists \alpha$ holds, as desired. A dual strategy holds for part (2).

Theorem 4. *For any prioritized clausal knowledge base* $(A, \leq)$*, any declaration* α *and any acceptable parameters* S *and* S' *such that* $S \subseteq S'$*,*

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

Proof. We only examine the first implication of part (1). Suppose we are given $(A, \leq, S) \Vdash^{\forall}_{\Box} \alpha$ and $(A, \leq, S') \nVdash^{\forall}_{\Box} \alpha$. From the second assertion, $\exists B \in \Box(A, \leq, S')$ such that $\not\models \Box_{S'} (B \supset \alpha)$. By contraposition of theorem 1, it follows that $\not\models \Box_S (B \supset \alpha)$. Moreover, since $B \in \Box(A, \leq, S')$, by application of lemma 4, $\exists C \in \Box(A, \leq, S)$ such that $C \subseteq B$. By the monotonicity property of conjunction, it follows that $\not\models \Box_S (C \supset \alpha)$. Therefore $(A, \leq, S) \nVdash^{\forall}_{\Box} \alpha$, hence contradiction.

Theorem 5. *For any prioritized knowledge base* $(A, \leq)$*, any declaration* α *and any acceptable parameters* S *and* S' *such that* $S \subseteq S'$*,*

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

Proof. We only examine the first implication of part (1). Suppose that $(A, \leq, S) \Vdash^{\mathcal{A}}_{\Box} \alpha$. Then, $(A, \leq, S) \Vdash^{\exists}_{\Box} \alpha$ and $(A, \leq, S) \nVdash^{\exists}_{\Diamond} \neg\alpha$. From the first assertion and by theorem 3(1), it follows that $(A, \leq, S') \Vdash^{\exists}_{\Box} \alpha$. From the second assertion and by theorem 3(2), it follows that $(A, \leq, S') \nVdash^{\exists}_{\Diamond} \neg\alpha$. Thus, $(A, \leq, S') \Vdash^{\mathcal{A}}_{\Box} \alpha$.

3.3 Computational properties

We now turn to computational considerations. To this very point, we recall that coherence-based reasoning is characterized by two interacting sources of complexity, namely, propositional satisfiability and the selection of preferred consistent subsets. The following theorem states that both sources of complexity are bounded by the same resource parameter S.

Theorem 6. *For any prioritized knowledge base* $(A, \leq)$*, any declaration* α *and any parameter* S*, there exists an algorithm for deciding whether* $(A, \leq, S) \Vdash^{x}_{\Box} \alpha$ *holds and* $(A, \leq, S) \Vdash^{x}_{\Diamond} \alpha$ *holds which runs in* $O((|A| + |\alpha|) \cdot 2^{|S|} \cdot 2^{|S|})$.

Proof. We focus on the complexity analysis of $(A, \leq, S) \Vdash^{\exists}_{\Box} \alpha$. The demonstration is analogous for the other entailment relations. We begin to prove that the size of $\Box(A, \leq, S)$ is bounded by $2^{|S|}$. Let B and B' be two sets of $\Box(A, \leq, S)$. Obviously, $\Box_S (B \cup B')$ is unsatisfiable. Let $V^{\Box}_S$ denotes the set of valuations v such that $\forall p \in P$, $v(p) = \{0\}$ or $v(p) = \{1\}$ if $p \in S$, and $v(p) = \{\}$ otherwise. Moreover, given a declaration β, let $V^{\Box}_S(\beta)$ denotes the set of valuations v in $V^{\Box}_S$ such that $v \models_1 \beta$. Clearly, $\Box_S (B \cup B')$ is unsatisfiable iff $V^{\Box}_S(B) \cap V^{\Box}_S(B') = \varnothing$. Since there exists $2^{|S|}$ valuations in $V^{\Box}_S$, the maximum number of bases being locally satisfiable and pairwise unsatisfiable under the scope of $\Box_S$ is $2^{|S|}$. Now, let us examine the main result. Suppose that if $(A, \leq, S) \Vdash^{\exists}_{\Box} \alpha$ holds then $\exists B \in \Box(A, \leq, S)$ such that $\models \Box_S (B \supset \alpha)$. By application of lemma 1 and theorem 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 $O((|A| + |\alpha|) \cdot 2^{|S|} \cdot 2^{|S|})$.

Several algorithms can be used for anytime nonmonotonic reasoning. The key difficulty lies in the consolidation operation. To this end, one may conceive an algorithm which takes in 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. Such an algorithm is indeed *incremental*; by exploiting lemmas 2 and 4, the procedure only needs to expand the maximal subsets generated in previous steps and does not require to perform all computations from scratch.

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. Consider the flat base $A = \{a, b, c, \neg c, \neg a \vee \neg b, \neg a \vee c, \neg a \vee \neg c, \neg b \vee d\}$. We want to show that $A \Vdash^{\exists} d$. Hence, we need to find a set S such that $A \Vdash^{\exists}_{\Box} d$. Starting with $S = \varnothing$ and using the minimal diversity heuristic, we iteratively add d and b to S. Based on the following results, we observe that $A \Vdash^{\exists}_{\Box} d$.

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 4. Suppose we are given the prioritized base $A = (A_1, A_2)$ where $A_1 = \{a, \neg a, e\}$ and $A_2 = \{c, \neg d, \neg a \vee b, \neg c \vee d\}$. We want to show that $(A, \leq) \Vdash^{\mathcal{A}} b$. So, we need to find a set S such that $(A, \leq) \Vdash^{\mathcal{A}}_{\Box} b$. 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 $(A, \leq) \Vdash^{\exists}_{\Box} b$ and $(A, \leq) \nVdash^{\exists}_{\Diamond} \neg b$.

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 Conclusion

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, incrementality and dual reasoning. Some of the future directions of this work include the empirical study of anytime coherence-based reasoning. To this point, some benchmarks for coherence-based reasoning have recently been proposed in [7]. An important issue is to compare the performances of the standard methods with our anytime technique.

References

1. C. Baral, S. Kraus, and J. Minker. Combining multiple knowledge bases. *IEEE Transactions on Knowledge and Data Engineering*, 3(2):208–220, 1991.
2. C. Baral, S. Kraus, J. Minker, and V. S. Subrahmanian. Combining knowledge bases consisting of first order theories. *Comp. Intelligence*, 8(1):45–71, 1992.
3. S. Benferhat, D. Dubois, and H. Prade. How to infer form inconsistent beliefs without revising ? In *Proc. IJCAI'95*, pages 1449–1455, 1995.
4. S. Benferhat, D. Dubois, and H. Prade. Some syntactic approaches to the handling of inconsistent knowledge bases: a comparative study, part 1: the flat case. *Studia Logica*, 58:17–45, 1997.
5. G. Brewka. Preferred subtheories: An extended logical framework for default reasoning. In *Proc. IJCAI'89*, pages 1043–1048. Morgan Kaufmann, 1989.
6. M. Cadoli and M. Schaerf. Approximate inference in default reasoning and circumscription. In *Proc. ECAI'92*, pages 319–323, 1992.
7. C. Cayrol and M. C. Lagasquie-Schiex. Nonmonotonic reasoning: from complexity to algorithms. *Annals of Mathematics and Artificial Intelligence*, 22:207–236, 1998.
8. R. Dechter and I. Rish. Directional resolution: The Davis-Putnam procedure, revisited. In *Proc. KR'94*, pages 134–145, 1994.
9. F. Koriche. A logic for anytime deduction and anytime compilation. In *Logics in Artificial Intelligence*, volume 1489, pages 324–342. Springer Verlag, 1998.
10. F. Koriche. A logic for approximate first-order reasoning. In *Proc. CSL'01*, 2001. To appear.
11. H. J. Levesque. A logic of implicit and explicit belief. In *Proc. AAAI'84*, pages 198–202, 1984.
12. F. Massacci. Anytime approximate modal reasoning. In *Proc. AAAI'98*, pages 274–279, 1998.
13. B. Nebel. Belief revision and default reasoning: Syntax - based approaches. In *Proc. KR'91*, pages 417–428, 1991.
14. B. Nebel. Base revision operations and schemes: Semantics, representation, and complexity. In *Mathematical and Statistical Methods in Artificial Intelligence*, pages 157–170. Springer-Verlag, 1995.
15. G. Pinkas and R. P. Loui. Reasoning from inconsistency: A taxonomy of principles for resolving conflict. In *Proc. KR'92*, pages 709–719, 1992.
16. N. Rescher and R. Manor. On inference from inconsistent premises. *Theory and Decision*, 1:179–217, 1970.
17. M. Schaerf and M. Cadoli. Tractable reasoning via approximation. *Artificial Intelligence*, 74:249–310, 1995.
18. B. Selman, H. Levesque, and D. Mitchell. A new method for solving hard satisfiability problems. In *Proc. AAAI'92*, pages 440–446, 1992.
19. H. Zhang and M. E. Stickel. Implementing the Davis-Putnam method. *Journal of Automated Reasoning*, 24(1/2):277–296, 2000.
20. S. Zilberstein. Using anytime algorithms in intelligent systems. *AI Magazine*, 17(3):73–83, 1996.