A FOUR-VALUED APPROACH FOR HANDLING INCONSISTENCY IN PRIORITIZED KNOWLEDGE-BASES

Ofer Arieli
Department of Computer Science
School of Mathematical Sciences
Tel-Aviv University
Tel-Aviv 69978, Israel.
Email: ofera@math.tau.ac.il

Abstract

The use of priorities among formulae is an important tool to appropriately revise inconsistent knowledge-bases. We present a four-valued semantical approach for recovering consistent data from prioritized knowledge-bases. This approach is nonmonotonic and paraconsistent in nature.

1 Introduction

There are many cases in which a knowledge-base contains formulae with different importance. For instance, rules that state default assumptions are usually considered as less reliable than rules without exceptions. Also, inference rules are usually given a lower priority than explicit data. These kinds of considerations are particularly common when revising inconsistent knowledge-bases; If some formulae are more certain than others, one would probably like to reject the least certain first.

Many different approaches for resolving conflicts in prioritized knowledge-bases have been proposed in the literature (see, e.g., [5, 6, 8, 9, 11, 12]). A lot of these methods draw conclusions based on maximal consistent subsets of the knowledge-base under consideration (see, e.g., [6] for a survey). However, the semantics of the maximal consistent sets might not correspond to that of the original knowledge-base. For example, none of the maximal consistent subsets of the simplest inconsistent knowledge-base $\{p, \neg p\}$ reflects its intended meaning. Moreover, each maximal consistent set even *contradicts an explicit assertion* of the original knowledge-base.

The method that is presented here also considers consistent sets of formulae as representing "recovered" data. However, instead of insisting on maximality, our major concern here is to preserve the original semantics of the assertions that are assigned high priority in the knowledge-base. Roughly speaking, the idea is to construct consistent subsets of formulae that reflect the semantics of the higher priority data, and to choose one of them according to some preference criteria. Then it is possible to apply classical inferences for making plausible conclusions from the recovered set of assertions. This kind of approach is called *coherent* [6] or *conservative* [13], since it treats contradictory data as useless, and regards the remaining data unaffected.

2 Recovery of knowledge-bases

2.1 Preliminaries

Our method is based on Belnap well-known logic [3, 4], which consists of four truth values: the classical ones (t, f), a value that intuitively represents lack of information $(\bot)$, and a value that indicates inconsistency $(\top)$. The two latter values make Belnap logic particularly useful for reasoning with uncertainty. $(\{t, f, \top, \bot\}, \leq)$ is a distributive lattice in which f is the minimal element, t is the maximal one, and $\bot, \top$ are two intermediate values that are incomparable (see Figure 1). We shall denote the meet and the join of this lattice by $\wedge$ and $\vee$, respectively. It is also possible to define an involution $\neg$ on this lattice, for which $\neg\top = \top$ and $\neg\bot = \bot$. The truth values t and $\top$ are called the *designated* elements, since they intuitively represent formulae known to be true. The other semantic notions are natural generalizations of the similar classical ones: A *valuation* ν is a function that assigns a truth value from $\{t, f, \top, \bot\}$ to each atomic formula. Any valuation is extended to complex formulae in the obvious way. The set of valuations on a set S of formulae is denoted $\mathcal{V}(S)$. We

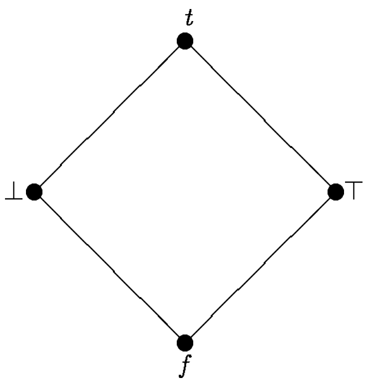

Figure 1: The four-valued lattice

will usually write $\psi : b \in \nu$ instead of $\nu(\psi) = b$. A valuation ν ***satisfies*** ψ iff $\nu(\psi) \in \{t, \top\}$. A valuation that satisfies every formula in S is a *model* of S. The set of all the models of S is denoted $mod(S)$.

Definition 2.1 Let S be a set of formulae and ψ – a formula. $S \models \psi$ if every model of S is a model of ψ.

The language we treat here is the standard propositional one. Given a set S of propositional formulae, we shall denote by $\mathcal{A}(S)$ the set of the atomic formulae that appear in the language of S, and by $\mathcal{L}(S)$ the set of the literals that appear in some formula of S. The formulae considered here are clauses, i.e.: disjunctions of literals. A set of clauses is called a *knowledge-base*, and is denoted by KB. As the following lemma shows, representing the formulae in a clause form does not reduce the generality.

Lemma 2.2 For every formula ψ there is a finite set S of clauses such that for every valuation ν, $\nu(\psi) \in \{\top, t\}$ iff $\nu(\phi) \in \{\top, t\}$ for every $\phi \in S$.

Definition 2.3 Let $\nu \in \mathcal{V}(KB)$. Define:
$Inc_\nu(KB) = \{p \in \mathcal{A}(KB) \mid \nu(p) = \top\}$.

Definition 2.4 Let $M, N \in mod(KB)$. M is *more consistent* than N iff $Inc_M(KB) \subset Inc_N(KB)$. M is a *most consistent* model of KB (mcm, for short) if there is no model of KB which is more consistent than M. The set of the most consistent models of KB is denoted $mcm(KB)$.

2.2 Recovered sets

As we have noted before, a drawback of using maximal consistent subsets is that none of these sets necessarily corresponds to the intended semantics of the original information. An approach that "salvages" consistent data from "polluted" knowledge-bases and still preserves their semantics is presented in [1, 2]. In what follows we briefly review this method and in the next section we generalize it for prioritized knowledge-bases.

Definition 2.5 A model M of a set of clauses KB is *consistent* if $Inc_M(KB) = \emptyset$. A knowledge-base is consistent if is has a consistent model.

Proposition 2.6 A knowledge-base is consistent iff it is classically consistent.

Definition 2.7 A *recovered set* S of KB is a subset of KB with a consistent model M s.t. there is a (not necessarily consistent) model M' of KB and $M(p) = M'(p)$ for every $p \in \mathcal{A}(S)$.

Definition 2.8 Let $\nu \in \mathcal{V}(KB)$. The set that is *associated with* ν is defined as follows:
$S_\nu(KB) = \{\psi \in KB \mid \nu(\psi) = t,\ \mathcal{A}(\psi) \cap Inc_\nu(KB) = \emptyset\}$.

Example 2.9 Consider the knowledge-base $KB = \{p,\ q,\ \neg p \vee \neg q\}$. $S_1 = \{p\}$ and $S_2 = \{q\}$ are the recovered sets of KB. These sets are associated with the (most consistent) models $\{p{:}t,\ q{:}\top\}$ and $\{p{:}\top,\ q{:}t\}$, respectively. Note that S_1 is no longer a recovered set of $KB' = KB \cup \{\neg p\}$, since there is no *consistent* model of S_1 that is expandable to a model of KB'.

Proposition 2.10 [1, 2] Every set that is associated with a model of KB is a recovered set of KB.

Given an inconsistent knowledge-base, the idea is to choose one of its maximal recovered sets (see [2]), and to treat this set as the relevant knowledge-base for deducing classical inferences. As the following proposition shows, there is a strong connection between maximal recovered sets of a knowledge-base and its mcms:

Proposition 2.11 [1, 2] Every maximal[1] recovered set of KB is associated with some mcm of KB.

3 Prioritized knowledge-bases

Consider a knowledge-base $KB = S \cup \{p, \neg p\}$ where p is an atomic formula, and S is a consistent set of clauses. For simplicity assume that $p \notin \mathcal{A}(S)$. The approach described in the previous section considers S as the recovered set of KB and ignores $p, \neg p$. Any larger consistent set will not properly reflect the intended semantics of KB.

This state of things should completely be changed if we know, for example, that p is more usual than

[1] The maximality is taken w.r.t. the containment relation.

$\neg p$. In such a case we would like to include p in the recovered set after all, since now the intended semantics is affected not only by the assertions in the knowledge-base, but also by some "meta-knowledge" that is provided with the original information. In this case p is given a higher priority than $\neg p$, and it seems reasonable that the recovered knowledge-base would contain p as well. The idea is, therefore, to distinguish between higher priority formulae and those with lower priority, and to assure that it will not be possible to draw any conclusion that contradicts formulae with high priority. In what follows we formalize this intuition.

Definition 3.1 Let $\nu \in \mathcal{V}(KB)$ and $S \subseteq KB$. The *reduction* of ν to S is the set $\nu \downarrow S = \{\nu(p) \mid p \in \mathcal{A}(S)\}$.

Definition 3.2 A *ranking* of a knowledge-base KB is a function r from the clauses in KB to $\{1, 2, \ldots, n\}$.

The ranking function determines a preference relation on the clauses of a knowledge-base; Intuitively, a clause with a lower rank has a higher priority.

Notation 3.3 $KB^i = \{\psi \in KB \mid r(\psi) \leq i\}$.

Definition 3.4 Let $R^i = \{S_\nu(KB) \mid \nu \in \mathcal{V}(KB),\ \nu \downarrow KB^i \in mcm(KB^i)\}$. Denote the maximal elements of R^i by $\mathcal{R}^i$. I.e., $\mathcal{R}^i = \{S \in R^i \mid \neg\exists T \in R^i \text{ s.t. } S \subset T\}$.

The $\mathcal{R}^i$'s are the candidates to be the set of the recovered knowledge-bases of KB. Following [5], we provide some criteria for choosing the preferred set. The index of this set determines what would be considered as a high ranking level:

- *set cardinality:* $\mathcal{R}^i \geq^{\text{sc}} \mathcal{R}^j$ iff $\forall S \in \mathcal{R}^i\ \exists T \in \mathcal{R}^j$ s.t. $|T| \leq |S|$.
- *set inclusion:* $\mathcal{R}^i \geq^{\text{si}} \mathcal{R}^j$ iff $\forall S \in \mathcal{R}^i\ \exists T \in \mathcal{R}^j$ s.t. $T \subseteq S$.
- *cardinality of consistent consequences:* $\mathcal{R}^i \geq^{\text{cc}} \mathcal{R}^j$ iff $\forall S \in \mathcal{R}^i\ \exists T \in \mathcal{R}^j$ s.t. $|\{l \in \mathcal{L}(KB) \mid T \models l, T \not\models \bar{l}\}| \leq |\{l \in \mathcal{L}(KB) \mid S \models l, S \not\models \bar{l}\}|$.[2]
- *inclusion of consistent consequences:* $\mathcal{R}^i \geq^{\text{ci}} \mathcal{R}^j$ iff $\forall S \in \mathcal{R}^i\ \exists T \in \mathcal{R}^j$ s.t. $\{l \in \mathcal{L}(KB) \mid T \models l, T \not\models \bar{l}\} \subseteq \{l \in \mathcal{L}(KB) \mid S \models l, S \not\models \bar{l}\}$.

Definition 3.5 The *optimal recovery level* of KB w.r.t. $\leq^{\text{sc}}$ is $i_0 = \max\{i \mid \neg\exists j \neq i \text{ s.t. } \mathcal{R}^j \geq^{\text{sc}} \mathcal{R}^i\}$. The optimal recovery levels of KB w.r.t. $\leq^{\text{si}}$, $\leq^{\text{cc}}$, and $\leq^{\text{ci}}$, are defined similarly.

[2] Where $\bar{l}$ denotes the complement of l.

Definition 3.6 Let i_0 be the optimal recovery level of a prioritized knowledge-base KB. The *recovered knowledge-bases* of KB are the elements of $\mathcal{R}^{i_0}$.

Definition 3.6 generalizes the notion of recovered sets (Definition 2.7) to the case of prioritized knowledge-bases:

Proposition 3.7 If all the clauses in KB have the same priority, then S is a recovered set of KB iff $S \in \mathcal{R}^1$.
Proof: Immediate from Proposition 2.11 and Definition 3.4, since $KB = KB^1$. □

Preference criteria like inclusion of consistent consequences (see above) or maximal information [10] might be applied to the elements of $\mathcal{R}^{i_0}$ for choosing the "best" recovered knowledge-base. For other preference criteria see, e.g., [6].

Before considering an example we extend the discussion to a language with predicates and variables. It is possible to do so in a straightforward way, provided that each clause that contains variables is considered as universally quantified. Consequently, a knowledge-base containing a non-grounded formula, ψ, will be viewed as representing the corresponding set of ground formulae formed by substituting each variable in ψ with every possible member of the Herbrand universe, U. Formally: $KB^U = \{\rho(\psi) \mid \psi \in KB,\ \rho : var(\psi) \to U\}$.

Example 3.8 (Tweety Dilemma) Consider the following well-known puzzle:

$bird(x) \to fly(x)$,

$penguin(x) \to \neg fly(x)$,

$penguin(x) \to bird(x)$,

$bird(Tweety),\ \ bird(Fred),\ \ penguin(Tweety)$

Denote the above knowledge-base by KB, and abbreviate the predicates *bird*, *penguin*, *fly* with b, p, and f (respectively)[3]. Also, T, F will stand for the individuals Tweety and Fred. KB has three mcms (see Figure 2) and three corresponding associated sets:

$$S_{M_1}(KB) = KB^U \setminus \{\psi \in KB^U \mid f(T) \in \mathcal{A}(\psi)\}$$
$$S_{M_2}(KB) = KB^U \setminus \{\psi \in KB^U \mid p(T) \in \mathcal{A}(\psi)\}$$
$$S_{M_3}(KB) = KB^U \setminus \{\psi \in KB^U \mid b(T) \in \mathcal{A}(\psi)\}$$

[3] Note that the symbol f has double meanings here: abbreviating the predicate *fly*, and representing the truth value FALSE. Each appearance of f will be understood by the context.

mcm	$b(T)$	$p(T)$	$f(T)$	$b(F)$	$p(F)$	$f(F)$
M_1	t	t	$\top$	t	f	t
M_2	t	$\top$	t	t	f	t
M_3	$\top$	t	f	t	f	t

Figure 2: The mcms of KB

There is no maximal recovered set of KB that entails all the properties of Tweety that one would expect to infer (i.e., that it is a bird, a penguin, and cannot fly). We claim, however, that this state of things is due to the fact that the actual representation of the problem does not properly reflect the intuitive understanding of this particular puzzle: While according to the above representation every rule is given the same importance, usually an explicit data (i.e. rules with empty bodies) is assigned a higher priority than other inference rules. Also, the first rule represents default assumption, and unlike the other rules it has exceptions, so it should be given a lower priority. In other words, we claim that a more accurate representation of this problem should be accompanied with some mechanism for making precedences among the rules. In our case this is a ranking function r. A possible ranking of KB is $r(b(T)) = r(b(F)) = r(p(T)) = 1$, $r(p(x) \to \neg f(x)) = r(p(x) \to b(x)) = 2$, and $r(b(x) \to f(x)) = 3$. By Proposition 3.10 below it follows that the optimal recovery level w.r.t. either $\leq^{\mathrm{ci}}$ or $\leq^{\mathrm{cc}}$ is $i=2$, and $KB^2 = \{b(T),\ b(F),\ p(T),\ p(x) \to \neg f(x),\ p(x) \to b(x)\}$. The most consistent models of KB^2 are given in the table of Figure 3.

mcm	$b(T)$	$p(T)$	$f(T)$	$b(F)$	$p(F)$	$f(F)$
M_1^2	t	t	f	t	f	f
M_2^2	t	t	f	t	f	t
M_3^2	t	t	f	t	f	$\perp$
M_4^2	t	t	f	t	t	f
M_5^2	t	t	f	t	$\perp$	f

Figure 3: The mcms of KB^2

It follows that $\mathcal{R}^2 = \{S_{M_2^2}(KB)\}$, so the recovered knowledge-base of KB is the following:

$$S_{M_2^2}(KB) = KB^U \setminus \{b(T) \to f(T)\}.$$

This recovered knowledge-base is associated with M_2^2, which coincides with the expected conclusions: Tweety is a bird, a penguin, and cannot fly, while Fred is a bird that can fly and it is not a penguin. $KB^U \setminus \{b(T) \to f(T)\}$ is also the (single) recovered knowledge-base obtained when taking $\leq^{\mathrm{sc}}$ or $\leq^{\mathrm{si}}$ as the preference order, or when the ranking is the following: $r(b(T)) = r(b(F)) = r(p(T)) = r(p(x) \to \neg f(x)) = r(p(x) \to b(x)) = 1$, $r(b(x) \to f(x)) = 2$. $\square$

We conclude with some basic properties of recovered knowledge-bases:

Proposition 3.9 Every recovered knowledge-base S of KB has a model that assigns classical values to the element of $\mathcal{A}(S)$.

Proof: Every recovered knowledge-base S is of the form $S_\nu(KB)$, where $\forall p \in \mathcal{A}(S)\ \nu(p) \in \{t, f, \perp\}$. Consider the valuation ν' s.t. $\nu'(p) = \nu(p)$ if $\nu(p) \in \{t, f\}$ and $\nu'(p) = t$ otherwise. By an induction on the structure of the clauses $\psi \in S$ it is easy to verify that $\nu'(\psi) \in \{t, \top\}$ whenever $\nu(\psi) \in \{t, \top\}$, thus ν' is a model of S. $\square$

Proposition 3.10 Let either $\leq^{\mathrm{cc}}$ or $\leq^{\mathrm{ci}}$ be the preferential relation defined on the sets $\mathcal{R}^j$, and suppose that KB^1 is consistent. Then:

a) The optimal recovery level is the maximal rank i s.t. KB^i is consistent.

b) Every recovered knowledge-base of KB is associated with a classical model on $\mathcal{A}(KB)$.

Proof: By the assumption on KB^1, there exists at least one $\mathcal{R}^j$ for which KB^j is consistent. Every such $\mathcal{R}^j$ is maximal w.r.t both $<^{\mathrm{cc}}$ and $<^{\mathrm{ci}}$, since by Proposition 2.6, $mcm(KB^j)$ consists only of consistent models of KB^j, which can be modified to classical models in the same way as in the proof of Proposition 3.9. These models can be extended to classical valuations ν_k^j on $\mathcal{A}(KB)$ by assigning classical values to every atom in $\mathcal{A}(KB \backslash KB^j)$. Each valuation ν_k^j has a set $S_{\nu_k^j}(KB)$ with which it is associated, and for *every* $p \in \mathcal{A}(KB)$, either p or $\neg p$ is in $\{l \in \mathcal{L}(KB) \mid S_{\nu_k^j}(KB) \models l, S_{\nu_k^j}(KB) \not\models \bar{l}\}$. Therefore, $S_{\nu_k^j}(KB) \in \mathcal{R}^j$, and so part (b) of the claim obtains. On the other hand, if KB^m is inconsistent, then for every model M of KB^m there is a $p_M \in \mathcal{A}(KB^j)$ s.t. $M(p_M) = \top$. Thus, if $S_M(KB) \in \mathcal{R}^m$, then neither p_M nor $\neg p_M$ is in the set $\{l \in \mathcal{L}(KB) \mid S_M(KB) \models l, S_M(KB) \not\models \bar{l}\}$. Therefore $\mathcal{R}^j >^{\mathrm{cc}} \mathcal{R}^m$ and $\mathcal{R}^j >^{\mathrm{ci}} \mathcal{R}^m$. $\square$

Proposition 3.11 Let i be the optimal recovery level of KB and S – a recovered knowledge-base of KB. Then:

a) $S \cap KB^i$ is a maximal recovered set of KB^i.

b) There is an $M^i \in mcm(KB^i)$ s.t. $S = S_{M^i}(KB)$.

c) If KB^1 is consistent, and the preference relation is either $\leq^{\mathrm{cc}}$ or $\leq^{\mathrm{ci}}$, then $S = KB^i \cup S_{M^i}(KB \setminus KB^i)$, where i is the maximal rank s.t. KB^i is consistent, and M^i is a (most) consistent model of KB^i.

Proof: (a) and (b) immediately follow from Proposition 2.11 and Definitions 3.4, 3.6. Part (c) follows from (b) and Proposition 3.10, since $S = S_{M^i}(KB^i) \cup S_{M^i}(KB \setminus KB^i) = KB^i \cup S_{M^i}(KB \setminus KB^i)$. $\square$

Each recovered knowledge-base is therefore a maximal set that is consistent (in the sense of Definition 2.5), and preserves the semantics of the clauses with the i-highest priorities (where i is an optimal recovery level w.r.t. some pre-defined criteria). To the maximal recovered set of KB^i we add clauses with lower priority than the optimal recovery level, provided that they are still true in the intended semantics and the consistency of the recovered set is not damaged.[4]

Corollary 3.12 Let S be a recovered knowledge-base of KB, and i – the optimal recovery level of KB. Then:
a) If $S \models \psi$ then $S \not\models \neg\psi$.
b) If $\psi, \neg\psi \in KB$ and $r(\psi) \leq i$ then $\neg\psi \notin S$.
c) Under the conditions of Proposition 3.10, if $\psi \in KB^i$ there is no recovered knowledge-base S' of KB s.t. $S' \models \neg\psi$.
Proof: (a) – Otherwise S cannot be consistent. (b) – If $\psi \in S$ then from (a), $\neg\psi \notin S$. Suppose, then, that $\psi \notin S$. By 3.11(b) $S = S_{M^i}(KB)$. Since $\psi \in KB^i$ necessarily $\mathcal{A}(\psi) \cap Inc_{M^i}(KB) \neq \emptyset$, and so $\neg\psi \notin S$ as well. (c) – Otherwise, by 3.11(b) $S'_{M^i}(KB) \models \neg\psi$, where $M^i \in mcm(KB^i)$. Thus $M^i(\psi) = \top$, and so there exists some $p \in \mathcal{A}(\psi)$ s.t. $M^i(p) = \top$ – a contradiction to 3.10(a). $\square$

4 Conclusion

We introduced a nonmonotonic and paraconsistent [7] method that copes with inconsistency in prioritized knowledge-bases. This approach preserves the semantics of the most valuable information of the original knowledge-base. Clauses with lower priorities are also added to the recovered knowledge-base provided that it remains consistent and the correspondence to the semantics of the data with higher priority is kept.

This method might be applied in various applications, among which are satisfactions of theories with prioritized constraints, systems that reason with prioritized default rules, knowledge-bases with multiple sources that have different priorities, diagnostic systems that analyze devices with different reliability, etc.

[4]In Example 3.8, for instance, $b(x) \rightarrow f(x)$ is not part of the recovered knowledge-base, but the instance $b(F) \rightarrow f(F)$ is included in it, since this instance is true in the intended semantics, and its addition still preserves the consistency of the recovered knowledge-base.

References

[1] O.Arieli and A.Avron. *A bilattice-based approach to recover consistent data from inconsistent knowledge-bases.* Proc. 4th Bar-Ilan Symp. on Foundations of AI, (M.Koppel and E.Shamir, eds.), pages 14–23. AAAI Press, 1995.

[2] O.Arieli and A.Avron. *Automatic diagnoses for properly stratified knowledge-bases.* Proc. 8th Int. Conf. on Tools with AI, pages 392–399. IEEE Press, 1996.

[3] N.D.Belnap. *How computer should think.* Contemporary Aspects of Philosophy (G.Ryle, ed.), pages 30–56. Oriel Press, 1977.

[4] N.D.Belnap. *A useful four-valued logic.* Modern Uses of Multiple-Valued Logic (G.Epstein and J.M.Dunn, eds.), pages 7–37. Reidel Publishing Company, 1977.

[5] S.Benferhat, C.Cayrol, D.Dubois, J.Lang, and H.Prade. *Inconsistency management and prioritized syntax-based entailment.* Proc. 13th Int. Joint Conf. on AI, pages 640–645, 1993.

[6] S.Benferhat, D.Dubois, and H.Prade. *How to infer from inconsistent beliefs without revising?* Proc. 14th Int. Joint Conf. on AI, pages 1449–1455, 1995.

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

[8] D.Lehmann and M.Magidor. *What does a conditional knowledge base entail?* J. Artificial Intelligence, 55(1):1–60, 1992.

[9] V.Lifschitz. *Circumscriptive theories: a logic-basd framework for knowledge representation.* Proc. AAAI'87, pages 364–368, 1987.

[10] E.L.Lozinskii. *Recovering contradictions: a plausible semantics for inconsistent systems.* J. Automated Reasoning, 12:1–31, 1994.

[11] V.S.Subrahmanian. *Amalgamating knowledge-bases.* ACM Transactions on Database Systems, 19(2):291–331, 1994.

[12] M.H.Van Emden. *Quantitative deduction and its fixpoint theory.* J. Logic Programming, 3:37–53, 1986.

[13] G.Wagner. *Vivid logic.* Lecture Notes in AI No.764. Springer-Verlag, 1994.