



HAL
open science

A Default Logic Patch for Default Logic

Philippe Besnard, Éric Grégoire, Sébastien Ramon

► **To cite this version:**

Philippe Besnard, Éric Grégoire, Sébastien Ramon. A Default Logic Patch for Default Logic. 10th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (EC-SQARU'09), 2009, Verona, Italy. pp.578-589. hal-00866491

HAL Id: hal-00866491

<https://hal.science/hal-00866491>

Submitted on 26 Sep 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Default Logic Patch for Default Logic

Philippe Besnard¹, Éric Grégoire², and Sébastien Ramon²

¹ IRIT, F-31062 Toulouse
CNRS UMR 5505, F-31062
118 route de Narbonne, F-31062 Toulouse France
`besnard@irit.fr`

² Université Lille - Nord de France, Artois, F-62307 Lens
CRIL, F-62307 Lens
CNRS UMR 8188, F-62307
rue Jean Souvraz SP18, F-62307 Lens France
`{gregoire,ramon}@cril.fr`

Abstract. This paper is about the fusion of multiple information sources represented using default logic. More precisely, the focus is on solving the problem that occurs when the standard-logic knowledge parts of the sources are contradictory, as default theories trivialize in this case. To overcome this problem, it is shown that replacing each formula belonging to Minimally Unsatisfiable Subformulas by a corresponding supernormal default allows appealing features. Moreover, it is investigated how these additional defaults interact with the initial defaults of the theory. Interestingly, this approach allows us to handle the problem of default theories containing inconsistent standard-logic knowledge, using the default logic framework itself.

Keywords: Default logic, logic-based fusion, inconsistency tolerance, MUS, Minimally Unsatisfiable Subformulas.

1 Introduction

In the Artificial Intelligence (A.I.) research community, one of the most popular tools to handle forms of defeasible reasoning remain Reiter's default logic [1] and its major variants (e.g. [2], [3], [4] and [5] just to name a few other seminal papers). Default logic has been defined to allow forms of reasoning by default to be modelled. It permits an inference system to jump to default conclusions and to retract them when new information shows that these conclusions now lead to inconsistency.

For example, default logic is a very convenient framework to encode patterns of reasoning like “Given an employee x , by default we should allow x to access the database unless this would contradict security rules. If some further additional information makes such contradictions occur then the permission must be retracted”.

A *default-logic theory* is made of two parts: a set of first-order logic formulas representing knowledge and a set of default rules, i.e. a sort of inference rules capturing patterns of defeasible reasoning as in the above example.

In this paper, we investigate how several¹ default theories in Reiter's default logic should be fused when it is assumed that each default theory represents the knowledge of an agent or of a community of agents. More precisely, it is shown that merging these theories is not an issue that is to be taken as granted when the set-theoretical union of the standard-logic formulas to be fused is inconsistent. Indeed, keeping all such formulas would make the whole language to be the set of acceptable inferences because when the standard-logic knowledge part of a default theory is inconsistent, the default theory itself trivializes.

Quite surprisingly, to the best of our knowledge, this trivialization property of default logic has not been addressed so far in the literature. In this respect, the goal of this paper is to revisit default logic in such a way that trivialization is avoided in the presence of inconsistent premises, sharing the concerns of the large research effort from the A.I. research community to study how to reason in the presence of inconsistent knowledge and to develop inconsistency tolerance techniques (see e.g. [6]). In particular, when several information sources are to be aggregated, a single, possibly minor contradiction between two sources should not cause the whole system to collapse.

In the paper, a family of approaches in that direction are discussed. Mainly, they rely on the study of MUSes (Minimally Unsatisfiable Subformulas) in the standard-logic formulas. Accordingly, a series of reasoning paradigms are investigated. Specifically, it is shown that replacing each formula in the set of MUSes by a corresponding default rule is an appealing solution. As a special case, it offers a powerful way to recover from the inconsistencies that might occur in sets of standard-logic formulas. Interestingly, this latter technique can easily be exported to the main variants of default logic, like e.g. constrained [2], rational [3], justified [4] and cumulative default logic [5], of which some ensure that general default theories have at least one extension.

The paper is organized as follows. In the next section, MUSes and the way according to which they can be computed are presented. In Sections 3 and 4, an approach to replace MUSes by additional default rules is introduced and studied in the context of recovering from inconsistency in standard Boolean logic. Section 5 is devoted to how these additional rules interact with the default ones of the initial theories. In Section 6, a complexity analysis of this technique is provided, together with possible approximation techniques.

Throughout the paper, we use the following standard notations: \neg , \vee , \wedge and \supset represent the standard negation, disjunction, conjunction and material implication connectives, respectively. When Ω is a set of first-order formulas, $Cn(\Omega)$ denotes the deductive closure of Ω . Also, let us recall that in the Boolean case a *CNF* is a finite conjunction of clauses, where a clause is a disjunction of signed Boolean variables.

¹ On the other hand, the following applies to a single default theory, too.

In the following, we assume that the reader is familiar with default logic [1]. A brief reminder about default logic is provided in Appendix A.

2 MUSes

Assume that Σ is a set of Boolean formulas. A *Minimally Unsatisfiable Subformulas (MUS)* Φ of Σ is defined as follows:

- $\Phi \subseteq \Sigma$,
- Φ is unsatisfiable,
- $\forall \Psi \subset \Phi$, Ψ is satisfiable.

Accordingly, a MUS of Σ is a subset of Σ that is contradictory and that becomes satisfiable whenever any of its formulas is removed. Thus, a MUS of Σ describes a contradiction within Σ using a set of formulas of Σ that cannot be made smaller.

Note 1. The set of all MUSes of a set of formulas Σ is denoted $MUS(\Sigma)$. The set of all formulas occurring in the MUSes of Σ is denoted $\cup MUS(\Sigma)$.

Example 1. Let $\Sigma = \{a, a \supset b, \neg b, a \supset (c \vee d), \neg d, c \supset b, d \supset e, (c \wedge e) \supset a\}$. Clearly, Σ is unsatisfiable and contains two MUSes, namely $\Phi_1 = \{a, a \supset b, \neg b\}$ and $\Phi_2 = \{a \supset (c \vee d), a, \neg d, c \supset b, \neg b\}$.

This example also illustrates that MUSes can share non-empty intersections.

Many techniques to handle contradictions in logic-based systems have been discussed in the literature (see e.g [6] and [7] for surveys in that matter). One family of approaches amount to recovering satisfiability by dropping MUSes or parts of MUSes. Indeed, removing one formula in each MUS allows consistency to be recovered. Two extreme approaches can thus be proposed in that direction. On the one hand, we might drop the set-theoretical union of all MUSes, thus removing every minimal (w.r.t. the number of involved formulas) cause of inconsistency. On the other hand, we might prefer a minimal change policy, which requires us to drop at most one formula per MUS.

3 How to Handle Default Theories Containing Contradictory Standard-Logic Knowledge

In the following, we assume that Σ is a set of Boolean formulas and we are mostly interested in default theories $\Gamma = (\Delta, \Sigma)$ where Σ is inconsistent. In such a case, Γ has a unique extension, which is the whole logical language.

We distinguish between skeptical and credulous reasonings from a default theory Γ : a formula f can be *skeptically* (resp. *credulously*) inferred from a default theory Γ iff f belongs to all (resp. some) extensions of Γ .

Now, since a default theory consists of two parts, namely a set of defaults and a set of facts, the fusion of default theories amounts to merging sets of facts and

merging sets of defaults. In the following, we assume that facts (resp. defaults) are unioned by this fusion process.

Assume that we are given n default theories $\Gamma_i = (\Delta_i, \Sigma_i)$ ($i \in [1..n]$) to be fused, that are such that the set-theoretical union of their standard logic parts, namely $\cup_{i=1}^n \Sigma_i$, is inconsistent. One direct way to address the trivialization of the resulting aggregated default theory consists of removing enough formulas from $\cup_{i=1}^n \Sigma_i$ so that the resulting subset becomes consistent. However, dropping formulas is unnecessarily destructive.

Indeed, a credulous reasoner might be interested in exploring the various extensions that could be obtained if regarding as being acceptable the various maximal consistent subsets of the various MUSes of $\cup_{i=1}^n \Sigma_i$. Also, a skeptical reasoner might want to explore what would belong to all those extensions. In this respect, if we replace each formula f in the MUSes of $\cup_{i=1}^n \Sigma_i$ by a corresponding supernormal default $\frac{f}{f}$, we get a new default theory where the reasoner is considering that *each formula f in the MUSes could be inferred if f could be consistently assumed*. However, since the set-theoretical union of the consequents of these new defaults is inconsistent, default logic forbids the acceptance of all such f within the same extension. Let us stress that this policy does not enforce by itself any priority between the replaced formulas since all of those are treated in a uniform way. Interestingly, this approach allows us to handle the problem of default theories containing inconsistent standard-logic knowledge, using the default logic framework itself.

Definition 1 (fused default theory). *Let us consider a non-empty set of n default theories of the form $\Gamma_i = (\Delta_i, \Sigma_i)$ to be fused. The resulting fused default theory is given by $\Gamma = (\Delta, \Sigma)$ where:*

- $\Sigma = \cup_{i=1}^n \Sigma_i \setminus \cup MUS(\cup_{i=1}^n \Sigma_i)$,
- $\Delta = \cup_{i=1}^n \Delta_i \cup \{ \frac{f}{f} \mid f \in \cup MUS(\cup_{i=1}^n \Sigma_i) \}$.

This definition thus corresponds to a policy that requires a uniform treatment of formulas inside MUSes. On the contrary, alternative definitions could make use of selection operators *select* to deliver a subset of $\cup MUS(\cup_{i=1}^n \Sigma_i)$ such that $\cup_{i=1}^n \Sigma_i \setminus select(\cup MUS(\cup_{i=1}^n \Sigma_i))$ is consistent, and such that each formula from $select(\cup MUS(\cup_{i=1}^n \Sigma_i))$ is to be replaced by a corresponding supernormal default in the fused theories.

4 Addressing the Trivialization Issue in the Standard Boolean Case

First, let us consider the basic situations where the set of defaults $\cup_{i=1}^n \Delta_i$ is empty. Obviously enough, this coincides with the problem of fusing sets of Boolean formulas that are such that their set-theoretical union is inconsistent. The above definition thus provides an original approach to address this issue.

Example 2. Let $\Gamma_1 = (\emptyset, \{-a \vee b, \neg b\})$ and $\Gamma_2 = (\emptyset, \{a\})$ two default theories to be fused. Clearly, $\cup_{i=1}^2 \Sigma_i = \{-a \vee b, a, \neg b\}$ is inconsistent. The fused default

theory $\Gamma = (\{\frac{ia}{a}, \frac{-a \vee b}{-a \vee b}, \frac{-b}{-b}\}, \emptyset)$ exhibits the extensions $E_1 = Cn(\{a, -a \vee b\})$, $E_2 = Cn(\{-b, -a \vee b\})$ and $E_3 = Cn(\{a, -b\})$; each of them containing two of the consequents of the three defaults.

Interestingly, it is possible to characterize the set-theoretic intersection of all extensions of the fused default theory, when the initial theories do not contain any default.

To do that, we resort to (usual) choice functions θ for a finite family of non-empty sets $\Xi = \{\Omega_1, \dots, \Omega_n\}$, which “pick” an element in every Ω_i of the family. In the limiting case that Ξ is empty, $\theta(\Xi)$ is empty.

Note 2. Let Θ denote the subclass of choice functions θ for $\Xi = \{\Omega_1, \dots, \Omega_n\}$ such that for $i \neq j$, $\theta(\Omega_i) \in \Omega_j \Rightarrow \theta(\Omega_j) = \theta(\Omega_k)$ for some $k \neq j$ (but k may, or may not, be i). This subclass is reduce to choices functions θ whose image is minimal s.t. if $\theta \in \Theta$ then $\nexists \theta' \in \Theta$ s.t. $\theta'(\Xi) \subset \theta(\Xi)$.

Proposition 1. *Let $n > 1$. Consider n finite default theories of the form $\Gamma_i = (\Delta_i, \Sigma_i)$ to be fused. If Δ_i is empty for $i = 1..n$, then the set-theoretic intersection of all extensions of the resulting fused default theory $\Gamma = (\Delta, \Sigma)$ is $Cn(\{\psi\})$ where:*

$$\psi = \bigvee_{\theta \in \Theta} \bigwedge ((\cup_{i=1}^n \Sigma_i) \setminus \theta(MUS(\cup_{i=1}^n \Sigma_i)))$$

Remark 1. It is essential that the default theories to be fused are *finite* for Proposition 1 to hold. Otherwise, $MUS(\cup_{i=1}^n \Sigma_i)$ can be infinite. Then, not only would the axiom of countable choice be needed, but even worse, an infinite disjunction would be needed (which is outside classical logic). For example, assume that the default theories to be fused are $\Gamma_1 = (\emptyset, \Sigma_1)$, $\Gamma_2 = (\emptyset, \Sigma_2)$, and $\Gamma_3 = (\emptyset, \Sigma_3)$ where:

$$\begin{aligned} \Sigma_1 &= \{p_1, q_1, r_1, \dots\}, \\ \Sigma_2 &= \{p_2, q_2, r_2, \dots\}, \\ \Sigma_3 &= \{\neg p_1, \neg p_2, \neg q_1, \neg q_2, \neg r_1, \neg r_2, \dots\}. \end{aligned}$$

Clearly, $MUS(\cup_{i=1}^3 \Sigma_i) = \{\{p_1, \neg p_1\}, \{p_2, \neg p_2\}, \{q_1, \neg q_1\}, \{q_2, \neg q_2\}, \dots\}$ is infinite. Therefore, ψ would have infinitely many conjuncts and disjuncts. For instance, taking θ_1 to pick only negative literals yields the infinite conjunction $\bigwedge \{p_1, p_2, q_1, q_2, r_1, r_2, \dots\}$. The disjunction would also be infinite because infinitely many choice functions must be taken into account.

In this example, Σ is empty but it is easy to alter it to make Σ infinite:

$$\begin{aligned} \Sigma_1 &= \{p_1, q_1, r_1, s_1, \dots\}, \\ \Sigma_2 &= \{p_2, q_2, r_2, s_2, \dots\}, \\ \Sigma_3 &= \{\neg p_1, p_2, q_1, \neg q_2, \neg r_1, r_2, s_1, \neg s_2, \dots\}. \end{aligned}$$

Interestingly, the following proposition shows us that any formula in the set $\cup MUS(\cup_{i=1}^n \Sigma_i)$ belongs to at least one extension of the fused theory. Accordingly, no formula is lost in the fusion process in the sense that each non-contradictory formula that is replaced by a default – and that would be dropped in standard fusion approaches – can be found in at least one extension of the fused theory.

Proposition 2. *Let $n > 1$. Consider n finite default theories $\Gamma_i = (\Delta_i, \Sigma_i)$ such that $\cup_{i=1}^n \Delta_i$ is empty and $\cup_{i=1}^n \Sigma_i$ is inconsistent. Let Γ denote the resulting fused default theory. There exists no extension of Γ that contains $\cup MUS(\cup_{i=1}^n \Sigma_i)$ but for any satisfiable formula f in $\cup MUS(\cup_{i=1}^n \Sigma_i)$, there exists an extension of Γ containing f .*

Based on the above proposition, it could be imagined that the intersection of all extensions will merely coincide with the extensions of the default theory $\Gamma' = (\emptyset, \cup_{i=1}^n \Sigma_i \setminus \cup MUS(\cup_{i=1}^n \Sigma_i))$. As the following example shows, this is not the case since the computation of the multiple extensions mimics a case analysis process that allows inferences to be entailed that would simply be dropped if $\cup MUS(\cup_{i=1}^n \Sigma_i)$ were simply removed from $\cup_{i=1}^n \Sigma_i$.

Example 3. Let us consider $\Gamma_1 = (\emptyset, \{a \wedge b, c \supset d\})$ and $\Gamma_2 = (\emptyset, \{\neg a \wedge c, b \supset d\})$. Clearly, $\{a \wedge b, \neg a \wedge c\}$ is a MUS. If we simply drop the MUS, we get $\Gamma' = (\emptyset, \{c \supset d, b \supset d\})$. Clearly, Γ' has a unique extension $Cn(\{c \supset d, b \supset d\})$ that does not contain d . This is quite inadequate since the contradiction is explained by the co-existence of a and $\neg a$. Assume that a is actually *true*. Then, from $a \wedge b$ and $b \supset d$ we should be able to deduce d . Similarly, if a is actually *false* then we should also be able to deduce d . Now, $\Gamma = (\{\frac{a \wedge b}{a \wedge b}, \frac{\neg a \wedge c}{\neg a \wedge c}\}, \{c \supset d, b \supset d\})$ exhibits two extensions, each of them containing d . Accordingly, d can be inferred using a skeptical approach to default reasoning.

This last example also shows us that this treatment of inconsistency permits more (legitimate) conclusions to be inferred than would be by removing MUSes or parts of MUSes. This is not surprising since by weakening formulas into default rules, we are dropping less information than if we were merely removing them. An alternative approach to allow such a form of case-analysis from inconsistent premises can be found in [8].

Applying Proposition 1 to Example 3 shows what the consequences of the resulting fused default theory $\Gamma = (\Delta, \Sigma)$ are:

Example 4 (con'd). $MUS(\cup_{i=1}^2 \Sigma_i) = \{\{a \wedge b, \neg a \wedge c\}\}$ because $\Sigma_1 \cup \Sigma_2$ has only one MUS, that is, $\{a \wedge b, \neg a \wedge c\}$. Hence, there are only two choice functions over $\{a \wedge b, \neg a \wedge c\}$. One picks $a \wedge b$ and the other picks $\neg a \wedge c$. Let us denote them θ_1 and θ_2 respectively. Then, the formula ψ in Proposition 1 becomes:

$$\psi = \bigvee_{\theta \in \{\theta_1, \theta_2\}} \bigwedge ((\cup_{i=1}^2 \Sigma_i) \setminus \theta(MUS(\cup_{i=1}^2 \Sigma_i))).$$

That is:

$$\psi = \bigvee_{\theta \in \{\theta_1, \theta_2\}} \bigwedge ((\Sigma_1 \cup \Sigma_2) \setminus \theta(\{\{a \wedge b, \neg a \wedge c\}\})).$$

So:

$$\psi = \left(\bigwedge \{\neg a \wedge c, c \supset d, b \supset d\} \right) \vee \left(\bigwedge \{a \wedge b, c \supset d, b \supset d\} \right).$$

Applying various logical laws, ψ becomes the conjunction of the following four formulas:

$$\begin{aligned} &(\neg a \wedge c) \vee (a \wedge b), \\ &((c \supset d) \wedge (b \supset d)) \vee (\neg a \wedge c), \\ &((c \supset d) \wedge (b \supset d)) \vee (a \wedge b), \\ &((c \supset d) \wedge (b \supset d)) \vee ((c \supset d) \wedge (b \supset d)). \end{aligned}$$

Of course, the latter disjunction is equivalent with $(c \supset d) \wedge (b \supset d)$ and subsumes the preceding two formulas. As a consequence:

$$\psi = ((\neg a \wedge c) \vee (a \wedge b)) \wedge (c \supset d) \wedge (b \supset d).$$

Finally, the set-theoretic intersection of all extensions of the resulting fused default theory $\Gamma = (\Delta, \Sigma)$ is $Cn(\{b \supset d, c \supset d, b \vee c, \neg a \vee b, a \vee c\})$.

Observe that $Cn(\{b \supset d, c \supset d, b \vee c, \neg a \vee b, a \vee c\})$ can be simplified as $Cn(\{b \supset d, c \supset d, \neg a \vee b, a \vee c\})$. In any case, both $b \vee c$ and d are in $Cn(\{b \supset d, c \supset d, \neg a \vee b, a \vee c\})$.

Now, another interesting feature of the fusion process given by Definition. 1 is that a skeptical reasoner will be able to infer (at least) all formulas that it would be able to infer if the MUSes of $\cup_{i=1}^n \Sigma_i$ were simply dropped.

Proposition 3. *Let $n > 1$. Consider n finite default theories $\Gamma_i = (\Delta_i, \Sigma_i)$ to be fused. Let $\cap_j E_j$ denote the set-theoretic intersection of all extensions of the resulting fused default theory $\Gamma = (\Delta, \Sigma)$. Let E denote the unique extension of $\Gamma' = (\emptyset, \Sigma)$. If Δ_i is empty for $i = 1..n$, then $E \subseteq \cap_j E_j$.*

Example 5. Assume $\Gamma_i = (\Sigma_i, \Delta_i)$ where $\cup_{i=1}^n \Sigma_i = \{a, \neg a \vee \neg b, b, c\}$ and $\cup_{i=1}^n \Delta_i = \emptyset$. $\cup MUS(\cup_{i=1}^n \Sigma_i) = \{a, \neg a \vee \neg b, b\}$. The resulting fused default theory is $\Gamma = (\Delta, \Sigma)$ where $\Sigma = \{c\}$ and $\Delta = \{\frac{a}{a}, \frac{\neg a \vee \neg b}{\neg a \vee \neg b}, \frac{b}{b}\}$. The extensions of Γ are $E_1 = Cn(\{c, a, \neg a \vee \neg b\})$, $E_2 = Cn(\{c, \neg a \vee \neg b, b\})$ and $E_3 = Cn(\{c, a, b\})$. The unique extension of the default theory where all formulas from $\cup MUS(\cup_{i=1}^n \Sigma_i)$ are dropped is $E = Cn(\{c\})$. We have $\cap_{j=1}^n E_j = Cn(\{c, a \vee b\})$ and $E \subseteq \cap_{j=1}^n E_j$.

Let us now consider the general case where theories contain defaults, and study how new defaults interact with defaults of the initial theories.

5 How the New Defaults Interact with the Defaults of the Theories to be Fused

First, it is well-known that normal default theories enjoy interesting properties, like semi-monotonicity [1]. This property ensures that whenever we augment a normal default theory Γ with an additional normal default, every extension of Γ is included in an extension of the new theory. Accordingly, we can insure that the extension of Proposition 2 to normal default theories holds since we only add supernormal defaults to the set-theoretical union of initial theories.

On the other hand, the extension of Proposition 3 to normal default theories does not hold: as the following example shows, the unique extension of $\Gamma' = (\cup_{i=1}^n \Delta_i, \Sigma)$ is not necessarily contained in the set-theoretical intersection of all the extensions of the resulting fused theory.

Example 6. Let us assume that $\cup_{i=1}^n \Sigma_i = \{a, \neg a \vee \neg b, b, c\}$ and $\cup_{i=1}^n \Delta_i = \{\frac{a:d}{d}, \frac{c:\neg d}{\neg d}\}$. The resulting fused default theory is $\Gamma = (\Delta, \Sigma)$ where $\Sigma = \{c\}$ and $\Delta = \{\frac{a}{a}, \frac{\neg a \vee \neg b}{\neg a \vee \neg b}, \frac{b}{b}, \frac{a:d}{d}, \frac{c:\neg d}{\neg d}\}$. The extensions of Γ are $E_1 = Cn(\{c, a, \neg a \vee \neg b, d\})$, $E_2 = Cn(\{c, \neg a \vee \neg b, b, \neg d\})$ and $E_3 = Cn(\{c, a, b, d\})$. The unique extension of the default theory $\Gamma' = (\cup_{i=1}^n \Delta_i, \Sigma)$ is $E = Cn(\{c, \neg d\})$ while $\cap_{j=1}^n E_j = Cn(\{c\})$. Thus $E \not\subseteq \cup_{j=1}^n E_j$.

Indeed, removing MUSes prevents the application of initial (normal) defaults whose prerequisite belongs to MUSes. Accordingly, we derive the following proposition.

Proposition 4. *Let $n > 1$. Consider n finite normal default theories $\Gamma_i = (\Delta_i, \Sigma_i)$ to be fused and $\Gamma' = (\cup_{i=1}^n \Delta_i, \cup_{i=1}^n \Sigma_i \setminus \cup MUS(\cup_{i=1}^n \Sigma_i))$. For any extension E of Γ' , there exists an extension of the resulting fused default theory that contains E .*

Interestingly, this proposition ensures that whenever we iterate the fusion process of normal default theories, we are always ensured that any extension can only be extended in the process.

Now, in the general case, replacing MUSes or subparts of MUSes by corresponding defaults does not ensure that we shall obtain supersets of the extensions that would be obtained if those MUSes or some of their subparts were removed: the semi-monotonicity does not hold.

Example 7. Let us consider $\Gamma = (\Delta, \{a, \neg a\})$ where $\Delta = \{\frac{b}{b}, \frac{a:c}{\neg b}, \frac{\neg a:c}{\neg b}\}$. The default theory $\Gamma' = (\Delta, \emptyset)$ exhibits one extension, which is $Cn(\{b\})$. On the contrary, $(\Delta \cup \{\frac{a}{a}, \frac{\neg a}{\neg a}\}, \emptyset)$ does not contain any extension containing b .

Generalizing Proposition 2 to default theories with non-empty sets of defaults does not hold either: as shown by the following example, it may happen that consistent formulas of $\cup MUS(\cup_{i=1}^n \Sigma_i)$ are in no extension of the resulting fused default theory.

Example 8. Let us consider the default theory $\Gamma_1 = (\emptyset, \{a, c\})$ and $\Gamma_2 = (\{\frac{c:b}{\neg a}\}, \{\neg a\})$ to be fused. $\cup MUS(\cup_{i=1}^2 \Sigma_i) = \{a, \neg a\}$. The resulting fused default theory is $\Gamma = (\{\frac{a}{a}, \frac{\neg a}{\neg a}, \frac{c:b}{\neg a}\}, \{c\})$. The unique extension of Γ is $E = Cn(\{c, \neg a\})$, which does not contain a .

6 Complexity Issues and Approximation Techniques

In the Boolean case, computing MUSes is computationally heavy in the worst case since checking whether a clause belongs or not to the set of MUSes of a CNF is Σ_2^P -complete [9]. Accordingly, the whole process of finding and replacing contradictory formulas by corresponding defaults, and then achieving Boolean credulous default reasoning is not computationally harder than credulous default reasoning itself since, in the general case, the latter is also Σ_2^P -complete (whereas it is Π_2^P -complete in the skeptical case) [10].

Interestingly, recent algorithmic techniques make it possible to compute one MUS for many real-life problems [11]. However, the number of MUSes in a set of n clauses can be intractable too, since it is $C_n^{n/2}$ in the worst case. Fortunately, efficient techniques have also been defined recently to compute all MUSes for many benchmarks, modulo a possible exponential blow-up limitation [12].

However, in some situations we cannot afford to compute the set-theoretical union of all MUSes. In this context, several techniques can then be applied.

First, it should be noted that it is not required to replace all formulas in all MUSes by corresponding defaults to recover consistency. Indeed, one could first detect one MUS, replace all its formulas by defaults, then iterate this process until consistency is recovered. Such an approach may avoid us computing all MUSes; it has been studied in the clausal Boolean framework in the context of the detection of *strict inconsistent covers* [13].

Definition 2 (strict inconsistent cover). *Let Σ be a set of Boolean formulas. $\Sigma' \subseteq \Sigma$ is a strict inconsistent cover of Σ iff $\Sigma \setminus \Sigma'$ is satisfiable and $\Sigma' = \cup \mathcal{A}$ for some $\mathcal{A} \subseteq \text{MUS}(\Sigma)$ such that, if $|\mathcal{A}| > 1$, any two members of \mathcal{A} are disjoint.*

Lemma 1. *A strict inconsistent cover of Σ is empty iff Σ is satisfiable.*

Lemma 2. *A strict inconsistent cover of Σ always exists.*

Lemma 3. *For all $M \in \text{MUS}(\Sigma)$, there exists a strict inconsistent cover of Σ that contains M .*

Strict inconsistent covers $IC(\Sigma)$ are thus minimal sets of formulas in Σ that can capture enough sources of contradiction in Σ to recover consistency if they were fixed. In [13] a technique to compute strict inconsistent covers in the Boolean clausal case has been introduced and proved efficient for many difficult benchmarks. Clearly, strict inconsistent covers is an approximation of the set-theoretical union of MUSes in the sense that all formulas of the cover always belong to this union but not conversely, and that dropping the cover causes consistency to be restored. The price to be paid for this approximation is that several different inconsistent covers can co-exist for a given set of MUSes.

Now, most of the time, it is possible to extract a super-set Ω of all MUSes of $\cup_{i=1}^n \Sigma_i$ very quickly, in such a way that retracting Ω would restore the consistency of $\cup_{i=1}^n \Sigma_i$. At the extreme, this super-set can be $\cup_{i=1}^n \Sigma_i$ itself. Accordingly, we could replace all formulas in Ω by corresponding supernormal defaults. Clearly, such a process would restore consistency. The price to be paid is that both uncontroversial and problematic information, i.e. both formulas belonging and not belonging to MUSes would be downgraded and treated in the same manner.

An alternative approach consists in replacing at most one formula per MUS. Clearly, from a practical point of view, detecting one such formula does not require us to compute one MUS exactly but simply a superset of a MUS, such that dropping the formula would make this superset consistent. Since MUSes can share non-empty intersections, let us note that it is however difficult to guarantee that a minimal number of formulas are replaced without computing all MUSes explicitly.

7 Conclusions and Future Works

In this paper, a “patch” to default logic, one of the most popular logics for representing defeasible reasoning, has been proposed. It allows a reasoner to handle theories involving contradictory standard-logic bases whereas standard default logic trivializes in this case. Interestingly, the new framework offers a powerful way to treat inconsistent standard logic theories as well. Such a default logic variant is of special interest when the fusion of several sources of knowledge is considered: indeed, without the patch, a default logic reasoner would be able to infer any conclusion (and its logical contrary) whenever two pieces of (standard logic) information appear to be contradicting one another in the sources.

In the basic approaches described in the previous section, no distinction is made between the defaults from the initial theories and the defaults that are introduced to replace MUSes or subparts of MUSes, as if all defaults were of the same epistemological nature. Indeed, the new defaults are introduced to *correct* and *weaken* some pieces of knowledge that exhibit some deficiencies. Our way to correct MUSes amounts to considering that formulas participating in MUSes should be accepted *by default*. In this respect, it can be argued that the role of the additional defaults is similar to the role of defaults of initial theories, which are normally intended to represent pieces of default reasoning.

On the contrary, it can be argued that new defaults should be given a higher (resp. lower) priority than defaults of initial theories. In these cases, we must resort to a form of prioritized default logic (see eg. [14]). We plan to investigate this issue in the future.

References

1. Reiter, R.: A logic for default reasoning. *Artificial Intelligence* 13, 81–132 (1980)
2. Schaub, T.: On constrained default theories. In: Neumann, B. (ed.) *European conference on Artificial Intelligence (ECAI 1992)*, New York, NY, USA, pp. 304–308. John Wiley & Sons, Inc., Chichester (1992)
3. Mikitiuk, A., Truszczyński, M.: Rational default logic and disjunctive logic programming. In: *Workshop on Logic Programming and Non-Monotonic Reasoning (LPNMR 1993)*, pp. 283–299. MIT Press, Cambridge (1993)
4. Lukaszewicz, W.: Considerations on default logic: an alternative approach. *Computational intelligence* 4(1), 1–16 (1988)
5. Brewka, G.: Cumulative default logic: in defense of nonmonotonic inference rules. *Artificial Intelligence* 50(2), 183–205 (1991)
6. Bertossi, L., Hunter, A., Schaub, T. (eds.): *Inconsistency Tolerance*. LNCS, vol. 3300. Springer, Heidelberg (2005)
7. Konieczny, S., Grégoire, E.: Logic-based information fusion in artificial intelligence. *Information Fusion* 7(1), 4–18 (2006)
8. Besnard, P., Hunter, A.: Quasi-classical logic: Non-trivializable classical reasoning from inconsistent information. In: Froidevaux, C., Kohlas, J. (eds.) *ECSQARU 1995*. LNCS, vol. 946, pp. 44–51. Springer, Heidelberg (1995)
9. Eiter, T., Gottlob, G.: On the complexity of propositional knowledge base revision, updates, and counterfactuals. *Artificial Intelligence* 57(2-3), 227–270 (1992)

10. Gottlob, G.: Complexity results for nonmonotonic logics. *Journal of Logic Computation* 2(3), 397–425 (1992)
11. Grégoire, E., Mazure, B., Piette, C.: Extracting muses. In: *European Conference on Artificial Intelligence (ECAI 2006)*, pp. 387–391. IOS Press, Amsterdam (2006)
12. Grégoire, E., Mazure, B., Piette, C.: Boosting a complete technique to find mss and mus thanks to a local search oracle. In: *International Joint Conference on Artificial Intelligence (IJCAI 2007)*, pp. 2300–2305 (2007)
13. Grégoire, E., Mazure, B., Piette, C.: Tracking muses and strict inconsistent covers. In: *ACM/IEEE Conference on Formal Methods in Computer Aided Design (FMCAD 2006)*, pp. 39–46 (2006)
14. Brewka, G., Eiter, T.: Prioritizing default logic. In: *Intellectics and Computational Logic (to Wolfgang Bibel on the occasion of his 60th birthday)*, Deventer, The Netherlands, pp. 27–45. Kluwer Academic Publishers, Dordrecht (2000)

Appendix A: Default Logic

The basic ingredients of Reiter’s default logic [1] are *default rules* (in short, *defaults*). A default d is of the form:

$$\frac{\alpha(\mathbf{x}) : \beta_1(\mathbf{x}), \dots, \beta_m(\mathbf{x})}{\gamma(\mathbf{x})},$$

where $\alpha(\mathbf{x}), \beta_1(\mathbf{x}), \dots, \beta_m(\mathbf{x}), \gamma(\mathbf{x})$ are first-order formulas with free variables belonging to $\mathbf{x} = \{x_1, \dots, x_n\}$, and are called the *prerequisite*, the *justifications* and the *consequent* of d , respectively.

Intuitively, d is intended to allow the reasoning “*Provided that the prerequisite can be established and provided that each justification can be separately consistently assumed w.r.t. what is derived, infer the consequent*”.

Accordingly, the example in the introduction could be encoded by:

$$\frac{\text{employee}(x) : \text{permit_access_DB}(x)}{\text{permit_access_DB}(x)}.$$

Such a default where the justification and consequent are identical is called a *normal default*. A normal default with an empty prerequisite is called a *super-normal default*. For a default d , we use $\text{pred}(d)$, $\text{just}(d)$, and $\text{cons}(d)$ to denote the prerequisite, the set of justifications and the consequent of d , respectively.

A *default theory* Γ is a pair (Δ, Σ) where Σ is a set of first-order formulas and Δ is a set of defaults. It is usually assumed that Δ and Σ are in skolemized form and that open defaults, i.e. defaults with free variables, represent the set of their closed instances over the Herbrand universe. A default theory with open defaults is *closed* by replacing open defaults with their closed instances. In the following, we assume that default theories are closed.

Defining and computing what should be inferred from a default theory is not a straightforward issue. First, there is a kind of circularity in the definition and computation of what can be inferred. To decide whether a consequent of a default should be inferred, we need to check the consistency of its justifications. However,

this consistency check amounts to proving that the opposites of the justifications cannot be inferred in turn. Actually, in the general case, fixpoints approaches are used to characterize what can be inferred from a default theory. Secondly, zero, one or several maximal sets of inferred formulas, called *extensions*, can be expected from a same default theory. One way to characterize extensions is as follows [1].

Let us define a series of sets of formulas E_i where $E_0 = Cn(\Sigma)$ and $E_{i+1} =$

$$Cn(E_i \cup \{\gamma \text{ s.t. } \frac{\alpha : \beta_1, \dots, \beta_m}{\gamma} \in \Delta \text{ where } \alpha \in E_i \text{ and } \neg\beta_1, \dots, \neg\beta_m \notin E_i\}),$$

for $i = 0, 1, 2, \text{etc.}$ Then, E is an extension of Γ iff $E = \bigcup_{i=0}^{\infty} E_i$.

A default d is called *generating* in a set of formulas Π , if $pred(d) \in \Pi$ and $\{\neg a \text{ s.t. } a \in just(d)\} \cap \Pi = \emptyset$. We note $GD(\Delta, E)$ the set of all defaults from Δ that are generating in E . It is also well-known that every extension of a default theory $\Gamma = (\Delta, \Sigma)$ is characterized through $GD(\Delta, E)$, i.e. $E = Cn(\Sigma \cup cons(GD(\Delta, E)))$, where $cons(\Delta') = \{cons(d) \text{ s.t. } d \in \Delta'\}$ for any set Δ' .