Permissive Belief Revision

(preliminary report)

Maria R. Cravo, João P. Cachopo, Ana C. Cachopo, João P. Martins
{mrcravo,jcachopo,acardoso,jpm}@gia.ist.utl.pt
Instituto Superior Técnico
Dep. Eng. Informática, GIA
Av. Rovisco Pais
1049-001 Lisboa — Portugal

Abstract

We point out that current belief revision operations can be used to revise non-monotonic theories and we propose a new operation, called permissive belief revision. The underlying idea of permissive belief revision consists of instead of abandoning some beliefs during a revision, transforming those beliefs into weaker ones, while still keeping the resulting belief set consistent. This framework allows us to keep more beliefs than what is usual using existent belief base-based revision theories.

1 Introduction

In this paper we are concerned with belief revision and with non-monotonic logics. In this section, we introduce the terminology that we use for belief revision and the main concepts in one non-monotonic logic formalism (default logic).

Although our presentation uses default logic, the basic ideas can be applied to other non-monotonic logic. Here, default logic should be taken as a general language for speaking about non-monotonic logics. It should not be understood as the chosen logic, to which we intend to apply our work. Further details would have to be specified for each particular non-monotonic logic.

1.1 Notation

Lower case greek letters (α, β, ...) represent meta-variables that range over single formulas; lower case roman letters (a, b, ...) represent single atomic formulas; upper case roman letters (A, B, ...) represent sets of formulas; $\mathcal{L}$ represents the language of classical logic (either propositional or first-order logic).

1.2 Belief revision

One of the main sources of inspiration in belief revision, the AGM theory, follows the work of [Alchourrón *et al.*, 1985]. AGM deals with deductively closed sets of sentences, called sets of beliefs. According to AGM, there are three operations on sets of beliefs: expansions, contractions, and revisions.

AGM presents a drawback from a computational point of view, since it deals with infinite sets of beliefs. Both [Nebel, 1989; 1990] and [Fuhrmann, 1991] modified AGM by working with a finite set of propositions, called a *belief base*, B, and using the set of consequences of B, defined as $Cn(B) = \{\phi : B \vdash \phi\}$.[1]

We use Nebel's notions of belief revision, a finite set of beliefs, and we only address the revision operation. The revision of a *consistent* belief base B with a formula ϕ, represented by $(B * \phi)$, consists in changing B in such a way that it contains ϕ and is consistent (if B is consistent). The case of interest is when $B \cup \{\phi\}$ is inconsistent, because, otherwise, ϕ can just be added to B.

To perform the revision $(B * \phi)$ when $B \cup \{\phi\}$ is inconsistent, we have to remove something from B, before we can add ϕ. In other words, in a revision $(B * \phi)$ some formulas *must* be discarded from B.

1.3 Non-monotonic logics

Most non-monotonic formalisms are composed of classical formulas (formulas that either belong to propositional logic or to first order logic, depending on the formalism) together with a way of expressing rules with exceptions (defaults) and exceptions to those rules: *Default Logic* [Reiter, 1980] is composed by a set of classical formulas, W, and by default rules, D, that are expressed as rules of inference; *Auto-epistemic Logic* [Moore, 1983; 1985; 1988] extends the language of first-order logic with the modal operator B (the intuitive interpretation of $B\alpha$ is "I believe in α"); *Circumscription* [McCarthy, 1980; 1986], uses the predicate Ab to identify exceptions to rules and provides a way of "jumping to conclusions" and to infer certain properties about the objects that satisfy certain relations.

In this paper we present examples based on Default Logic. Here, we recall the main concepts of Default Logic. Default Logic uses the language of classical logic, and, besides the classical rules of inference, it uses *rules of inference* of the form: $\alpha(\vec{x}) : \beta_1(\vec{x}), ..., \beta_m(\vec{x}) / \gamma(\vec{x})$, where $\alpha(\vec{x})$, $\beta_1(\vec{x}), ..., \beta_m(\vec{x})$, and $\gamma(\vec{x})$ are formulas whose free variables belong to the vector $\vec{x} = (x_1, ..., x_n)$. Such a rule of inference, *default rule*, is interpreted in the fol-

[1] $\vdash$ represents the classical derivability operation.

lowing way: from $\alpha(\vec{x_0})$,[2] if it is consistent to assume $\beta_1(\vec{x_0})$,...,$\beta_m(\vec{x_0})$, then we can infer $\gamma(\vec{x_0})$. Default rules can be looked at as suggestions with respect to what we should believe, in addition to what is dictated by classical logic.

A *default theory* is a pair (D, W), composed of a set of default rules, D, and by a set of closed formulas, $W \subset \mathcal{L}$. Given a default theory (D, W), we want to compute its consequences, the sets of formulas derivable from W using the rules of inference of classical logic and the default rules in D. In Default Logic, there may be more than one of these sets or even none. Each one of these sets is called an *extension* of the default theory (D, W). Each extension may be interpreted as a reasonable set of beliefs generated from W, using the default rules in D.

1.4 Contributions

We address the revision operation in relation with non-monotonic theories and we use it in two aspects:

1. We use the classical revision "$*$" applied to a non-monotonic theory. This application is done over the classical formulas of the theory and does not change the default rules and exceptions of the theory;
2. We define a revision operation, called permissive revision and denoted by "$\circledast$", that adds to the results of the classical revision a weaker version of the formulas removed by the classical revision operation. Given two formulas f and w, we say that w is weaker than f if and only if everything that is derivable from w was also derivable from f, but not the inverse.[3]

This second aspect can either be applied to a monotonic or to a non-monotonic theory. The main difference between the two types of theories results from the fact that in a non-monotonic theory, universal rules may be weakened by being turned into default rules.

In this paper we give greater emphasis to permissive revision using a monotonic logic, although we point out research directions towards non-monotonic logics.

2 Revision of non-monotonic theories

Most belief revision theories consider classical logic as their underlying logic. We refer to these as *classical belief revision theories.* In this section we discuss how classical belief revision can be used to determine what the revision of a default theory with a formula should be (but not the contraction of a default theory with a formula). Given a default theory, (D, W), we can distinguish two types of consequences:

- *Classical consequences.* The formulas which are a classical consequence of W. The classical consequences are the set $\{\phi : W \vdash \phi\}$.
- *Defeasible consequences.* The consequences that were reached by using at least one default rule. If E is an extension of (D, W), the defeasible consequences in E are the set $E - \{\phi : W \vdash \phi\}$.

When we revise a default theory (D, W), with a formula ϕ, two things may happen:[4]

1. If ϕ is consistent with the consequences (classical or defeasible) of (D, W), nothing needs to be done, the revised theory is $(D, (W \cup \{\phi\}))$.
2. If ϕ is inconsistent with some consequence of (D, W) we must distinguish between two cases (notice that these are not mutually exclusive):
   (a) If ϕ is inconsistent with a classical consequence of (D, W), we can use a classical belief revision to determine $(W * \phi)$, and the set of default rules does not need to be changed. In other words, the revision $((D, W) * \phi)$ is $(D, (W * \phi))$.
   (b) If ϕ is inconsistent with a defeasible consequence of (D, W), nothing needs to be done, the logic will take care of the problem.

In summary, the revision of a default theory with a formula is defined by:

$$((D, W) * \phi) = (D, (W * \phi))$$

Note that the same symbol "$*$", is being used for the revision of a default theory, in $((D, W) * \phi)$, and for classical revision, in $(D, (W * \phi))$. Since there is no possible confusion, we do not introduce a different symbol.

3 Permissive revision

3.1 Motivations

The idea of permissive revision is to transform the beliefs that were discarded in a classical revision into weaker versions and to add them to the result of the revision. Permissive revision, thus, corresponds to a "smaller" change in beliefs than classical revision, while keeping the goal of having a consistent result.

Conjunctions are the most obvious candidates to be weakened. This aspect was already recognized by [Lehmann, 1995], who discussed that revision theories sometimes require to give up too many beliefs, without providing a solution to the problem. While Lehmann only presents the problem regarding conjunctions, we

[2] $\vec{x_0}$ is an instance of $\vec{x}$.

[3] This is based on the definition presented in [Quine and Ullian, 1978].

[4] Thanks to the anonymous reviewer that pointed out that this may produce a theory with no extensions. Should we decide to particularize this work for default logic, something would have to be done, namely to revise the default rules of the theory (something in the line of the work of [Witteveen *et al.*, 1994]). However, that would be a completely different issue from the ones addressed in this paper, where only the revision of the classical formulas of a non-monotonic theory is contemplated. It should also be noted that the problem of getting a default theory with no extensions after a revision with a classical formula is just a manifestation of the more general fact that there are default theories with no extensions. This does not happen in all non-monotonic logics [Cravo, 1993b; Poole, 1988].

argue that this problem is more general and that it can arise with other kinds of formulas.

To illustrate the main idea behind the weakening of conjunctions, suppose, for instance, that some classical revision operation provides the result:

$$(\{a \wedge b, a \Rightarrow c\} * \neg c) = \{a \Rightarrow c, \neg c\}$$

Permissive revision weakens the abandoned formula, $a \wedge b$, to b, and adds this to the result of classical revision:

$$(\{a \wedge b, a \Rightarrow c\} \circledast \neg c) = \{b, a \Rightarrow c, \neg c\}$$

If we are using a non-monotonic logic, other obvious candidates to be weakened are universal rules. For instance, suppose that we are using Default Logic, and we have a theory (D, W), where

$$D = \{\}$$
$$W = \{Bird(T), \forall(x) Bird(x) \Rightarrow Flies(x)\}$$

Further, assume that some classical revision operation gives us the following result:

$$((D, W) * \neg Flies(T)) = (\{\}, \{Bird(T), \neg Flies(T)\})$$

Permissive revision weakens the abandoned formula, $\forall(x)Bird(x) \Rightarrow Flies(x)$, to $Bird(x) : Flies(x)/Flies(x)$, and adds this to the result of classical revision:

$$((D, W) \circledast \neg Flies(T)) = (\{Bird(x) : Flies(x)/Flies(x)\}, \{Bird(T), \neg Flies(T)\})$$

3.2 Formalization

By now, it should be clear that the main task in defining permissive revision is the definition of a function Wk, which weakens the formula that was removed during classical revision. Actually, since there may be more than one such formula, we consider the conjunction of all the removed formulas, and weaken it into a new formula which will then be added to the result of classical revision to obtain permissive revision.

The function Wk will have different definitions, depending on whether we are using classical logic, a non-monotonic logic or some other logic. For instance, weakening a universally quantified formula into a default will not make sense in first order logic. For now we will restrict ourselves to classical first order logic.

Weakening a formula depends, naturally, on the set of formulas into which we will be adding the result. Therefore, the function Wk will receive the formula to weaken and a set of formulas:

$$\mathrm{Wk} : \mathcal{L} \times 2^{\mathcal{L}} \rightarrow \mathcal{L}$$

$\mathrm{Wk}(\phi, W)$ can be interpreted as "Weaken the formula ϕ, in such a way that after the weakened formula is added to W, the resulting set is not inconsistent".

Given such a function, we can formally define the permissive revision of a set of formulas W with a formula ϕ, $(W \circledast \phi)$. Let *Abandoned* be the conjunction of all the formulas which were abandoned during the classical revision of W with ϕ, $Abandoned = \bigwedge(W - (W * \phi))$. Then, the permissive revision of W with ϕ is given by

$$(W \circledast \phi) = (W * \phi) \cup \{\mathrm{Wk}(Abandoned, (W * \phi))\}$$

Let's now see how a formula is weakened. Obviously, this depends on the type of formula in question. The two examples in the previous section convey the main ideas behind weakening conjunctions and universal rules. However, there are other logical symbols, besides conjunctions and the universal quantifier, and there may be nested occurrences of each of these. Considering the usual logical symbols, $\{\neg, \Rightarrow, \wedge, \vee, \exists, \forall\}$, we have the following definition for Wk.

$$\mathrm{Wk}(\phi, W) = \begin{cases} \phi & \text{if } W \cup \{\phi\} \text{ is consistent} \\ \mathrm{WkN}(\phi, W) & \text{if } \phi \text{ is a negation} \\ \mathrm{WkI}(\phi, W) & \text{if } \phi \text{ is an implication} \\ \mathrm{WkD}(\phi, W) & \text{if } \phi \text{ is a disjunction} \\ \mathrm{WkC}(\phi, W) & \text{if } \phi \text{ is a conjunction} \\ \mathrm{WkE}(\phi, W) & \text{if } \phi \text{ is an existential rule} \\ \mathrm{WkU}(\phi, W) & \text{if } \phi \text{ is a universal rule} \\ \top & \text{otherwise} \end{cases}$$

Note that, although Wk will only be used, in the context of permissive revision, to weaken a formula ϕ known to be inconsistent with W, the weakening process is recursive (on the structure of formulas), and there may be sub-formulas which are consistent with W. That's the reason for the first case. As for the last case, which means that ϕ is an atomic formula inconsistent with W, there is no weaker formula we can give than a valid formula.

Next, we define each of the weakening functions mentioned above. We should keep in mind that a good weakening function should allow us to keep as much information as possible. In order to do that for non-atomic formulas, we weaken each sub-formula and combine the results.

When $\phi = \neg\alpha$, for some atomic formula α, there is nothing we can retain of the weakening of ϕ. However, if α is a non-atomic formula, $a \vee b$, for instance, we can apply logical transformations to ϕ to bring to the surface a kind of formula we know how to handle. In this case $\neg(a \vee b)$ is logically equivalent to $(\neg a) \wedge (\neg b)$.

$$\mathrm{WkN}(\phi, W) = \begin{cases} \mathrm{Wk}(\neg\alpha \wedge \neg\beta, W) & \text{if } \phi = \neg(\alpha \vee \beta) \\ \mathrm{Wk}(\neg\alpha \vee \neg\beta, W) & \text{if } \phi = \neg(\alpha \wedge \beta) \\ \mathrm{Wk}(\alpha \wedge \neg\beta, W) & \text{if } \phi = \neg(\alpha \Rightarrow \beta) \\ \mathrm{Wk}(\alpha, W) & \text{if } \phi = \neg\neg\alpha \\ \mathrm{Wk}(\forall(x)\neg\alpha(x), W) & \text{if } \phi = \neg\exists(x)\alpha(x) \\ \mathrm{Wk}(\exists(x)\neg\alpha(x), W) & \text{if } \phi = \neg\forall(x)\alpha(x) \\ \top & \text{otherwise} \end{cases}$$

Weakening an implication is treated in a similar way, transforming the implication into the logically equivalent disjunction, and weakening the result instead.

$$\mathrm{WkI}(\alpha \Rightarrow \beta, W) = \mathrm{Wk}(\neg\alpha \vee \beta, W)$$

If $\phi = \alpha \vee \beta$, and it is inconsistent with W (otherwise WkD would not be used), then both α and β are inconsistent with W. So, to weaken ϕ we have to individually weaken both α and β, in W, and combine the results with the disjunction again.

$$\mathrm{WkD}(\alpha \vee \beta, W) = \mathrm{Wk}(\alpha, W) \vee \mathrm{Wk}(\beta, W)$$

Conjunction seems to be a more complex case. To help understand its definition we present some examples.

- First, we consider the case where only one of the conjuncts is inconsistent: consider the set $W = \{a \wedge b\}$ and its revision with $\neg a$. Using permissive revision, $\mathrm{WkC}(a \wedge b, \{\neg a\}) = b$. We just have to abandon one of the elements of the conjunction and keep the other.
- If each of the conjuncts per se is inconsistent, like when $W = \{a \wedge b\}$ and we want to revise it with $\neg a \wedge \neg b$, we have to abandon both conjuncts, so $\mathrm{WkC}(a \wedge b, \{\neg a \wedge \neg b\}) = \top$, that is, there is nothing we can keep from the initial set.
- If it is the conjunction itself that is inconsistent, for example when $W = \{a \wedge b\}$ and we want to revise it with $\neg(a \wedge b)$, we have that $\mathrm{WkC}(a \wedge b, \{\neg(a \wedge b)\}) = a \vee b$, that is, each of the conjuncts may be true, but not both at the same time.

However, if each element is itself a non-atomic formula, the contradiction may be deeper inside in either one or in both of the elements of the conjunction. For instance, given $W = \{(a \wedge b) \wedge (c \wedge d)\}$ and revising it with $\neg(b \wedge c)$ we would like to get $(a \wedge (c \wedge d)) \vee ((a \wedge b) \wedge d)$, i.e., if it's not possible to have both b and c, then we would like to have either a, b and d or a, c and d. This is the result of $\mathrm{WkC}((a \wedge b) \wedge (c \wedge d), \{\neg(b \wedge c)\})$, according to the following definition.

$$\begin{aligned}
&\mathrm{WkC}(\alpha \wedge \beta, W) = \\
&\qquad (\mathrm{Wk}(\alpha, W) \wedge \mathrm{Wk}(\beta, W \cup \{\mathrm{Wk}(\alpha, W)\})) \vee \\
&\qquad (\mathrm{Wk}(\beta, W) \wedge \mathrm{Wk}(\alpha, W \cup \{\mathrm{Wk}(\beta, W)\}))
\end{aligned}$$

Handling existentially quantified formulas will be done through skolemization, weakening the formula which results from the elimination of the existential quantifier.

$$\begin{aligned}
&\mathrm{WkE}(\exists(x)\alpha(x), W) = \mathrm{Wk}(\alpha(p), W) \\
&\qquad\qquad \text{where } p \text{ is a Skolem constant}
\end{aligned}$$

Although in the motivation one of the presented examples deals with universally quantified formulas, that example assumes an underlying non-monotonic logic. In a classical setting, as we are now, weakening this kind of formula is somewhat more difficult. We briefly discuss this in the following section and for now use the simpler possible definition.

$$\mathrm{WkU}(\forall(x)\alpha(x), W) = \top$$

Before we prove some desirable properties of the weakening function, we present an example that shows the weakening of a disjunction of conjunctions. Given the set

$$W = \{(a \wedge b) \vee (c \wedge d), b \Rightarrow e, d \Rightarrow e, a \Rightarrow f, c \Rightarrow f\}$$

suppose

$$(W * \neg e) = \{b \Rightarrow e, d \Rightarrow e, a \Rightarrow f, c \Rightarrow f, \neg e\}$$

then

$$\begin{aligned}
&\mathrm{Wk}((a \wedge b) \vee (c \wedge d), (W * \neg e)) = \\
&\quad = \mathrm{WkD}((a \wedge b) \vee (c \wedge d), (W * \neg e)) = \\
&\quad = \mathrm{Wk}(a \wedge b, (W * \neg e)) \vee \mathrm{Wk}(c \wedge d, (W * \neg e)) = \\
&\quad = \mathrm{WkC}(a \wedge b, (W * \neg e)) \vee \mathrm{WkC}(c \wedge d, (W * \neg e)) = \\
&\quad = a \vee c
\end{aligned}$$

and

$$(W \circledast \neg e) = \{a \vee c, b \Rightarrow e, d \Rightarrow e, a \Rightarrow f, c \Rightarrow f, \neg e\}$$

Note that in the classical revision we can no longer derive, for instance f, but this is still a consequence of the permissive revision.

Even though we have not fully studied which properties the Wk function should have, there are some properties that we certainly want it to have: we don't want to produce an inconsistent set when adding the weakening of a formula to the result of the classical revision, nor do we want to be able to derive new conclusions from the weakening of a formula that were not derivable from the formula itself.

Theorem 3.1 guarantees that we do not produce an inconsistent set of beliefs by adding the weakening of any formula to a consistent set.

Theorem 3.1 *Let W be a consistent set of formulas, and ϕ any formula. Then $W \cup \{\mathrm{Wk}(\phi, W)\}$ is consistent.*

Proof. If ϕ is consistent with W, then $\mathrm{Wk}(\phi, W) = \phi$ and the result follows trivially. Otherwise, we will prove by induction on the structure of the formula ϕ that the weakening function produces a formula consistent with W.

If ϕ is a literal (an atomic formula or the negation of an atomic formula) or a universally quantified formula, then $\mathrm{Wk}(\phi, W) = \top$, and therefore $W \cup \{\top\}$ is consistent, provided that W is consistent.

The cases where ϕ is of the form $\neg\alpha$ or $\alpha \Rightarrow \beta$, reduce to one of the other cases, since the weakening of ϕ in these cases reduces to the weakening of a logical equivalent formula, with either a quantifier, a disjunction or a conjunction.

Assume that α, β and $\gamma(p)$, where p is some constant, are formulas that verify the theorem. Since $W \cup \{\mathrm{Wk}(\gamma(p), W)\}$ is consistent by hypothesis, then $W \cup \{\mathrm{Wk}(\exists(x)\gamma(x), W)\}$ is also consistent, by definition of WkE. Accordingly, given that $W \cup \{\mathrm{Wk}(\alpha, W)\}$ is consistent, and, therefore, $W \cup \{\mathrm{Wk}(\alpha, W) \vee \mathrm{Wk}(\beta, W)\}$ is consistent, we prove that $W \cup \{\mathrm{Wk}(\alpha \vee \beta, W)\}$ is also consistent. Finally, let $W' = W \cup \{\mathrm{Wk}(\alpha, W)\}$, which, as we have seen, is consistent. Since, by hypothesis, $W' \cup \{\mathrm{Wk}(\beta, W')\}$ is consistent, i.e., $W \cup \{\mathrm{Wk}(\alpha, W), \mathrm{Wk}(\beta, W')\}$ is consistent, we have that $W \cup \{\mathrm{Wk}(\alpha, W) \wedge \mathrm{Wk}(\beta, W')\}$ is consistent, from where it follows trivially that $W \cup \{\mathrm{Wk}(\alpha \wedge \beta, W)\}$ is consistent, which finishes our proof. □

Theorem 3.2 shows that the result of weakening a formula is something not stronger than the original formula, i.e., we are not introducing new consequences.

Theorem 3.2 *Let W be a set of formulas, and ϕ any formula. Then $\phi \vdash \mathrm{Wk}(\phi, W)$.*

Proof. If $\phi \vdash \bot$ then $\phi \vdash \psi$ for every formula ψ, and in particular for $\psi = \mathrm{Wk}(\phi, W)$. If ϕ is consistent with W, then $\mathrm{Wk}(\phi, W) = \phi$ and, obviously, $\phi \vdash \phi = \mathrm{Wk}(\phi, W)$. Otherwise, as above, we will prove by induction on the structure of the formula ϕ that the weakening function produces a formula not stronger than the original.

The structure of this proof is similar to the previous one: if ϕ is a literal or a universally quantified formula, then $\mathrm{Wk}(\phi, W) = \top$, and $\phi \vdash \top$; if ϕ is of the form $\neg\alpha$ or $\alpha \Rightarrow \beta$, the weakening of ϕ reduces to the weakening of a logical equivalent formula, with either a quantifier, a disjunction or a conjunction.

By eliminating the existential quantifier, we have that $\exists(x)\gamma(x) \vdash \gamma(p)$ for some Skolem constant p. By hypothesis, $\gamma(p) \vdash \mathrm{Wk}(\gamma(p), W) = \mathrm{Wk}(\exists(x)\gamma(x), W)$, and, therefore, $\exists(x)\gamma(x) \vdash \mathrm{Wk}(\exists(x)\gamma(x), W)$.

Assume that α and β are formulas that verify the theorem. Given that, by hypothesis, $\alpha \vdash \mathrm{Wk}(\alpha, W)$, then $\alpha \vdash \mathrm{Wk}(\alpha, W) \vee \mathrm{Wk}(\beta, W)$, and, likewise, since $\beta \vdash \mathrm{Wk}(\beta, W)$ then $\beta \vdash \mathrm{Wk}(\alpha, W) \vee \mathrm{Wk}(\beta, W)$. Joining the two, we have that $\alpha \vee \beta \vdash \mathrm{Wk}(\alpha, W) \vee \mathrm{Wk}(\beta, W)$, i.e., $\alpha \vee \beta \vdash \mathrm{Wk}(\alpha \vee \beta, W)$. To finish the proof, let's see that conjunction preserves the theorem: from $\alpha \vdash \mathrm{Wk}(\alpha, W)$ and $\beta \vdash \mathrm{Wk}(\beta, W \cup \{\mathrm{Wk}(\alpha, W)\})$, it follows trivially that $\alpha \wedge \beta \vdash \mathrm{Wk}(\alpha \wedge \beta, W)$. □

3.3 Universal rules

One of the aspects that is still under study is the way of weakening universal rules. In the last section, we said that universal rules are weakened to $\top$, which is obviously too drastic a solution.

This aspect can be improved in two directions. When considering a monotonic logic, a universal rule can be weakened following the general ideas presented in last section. For instance, if we have $\forall(x)\ a(x) \wedge b(x)$, and revise this with $\neg a(p)$, the universal rule must be abandoned, but it can be weakened to $\forall(x)\ b(x)$.

In another direction, i.e., when considering a non-monotonic logic, the most natural way of weakening a universal rule is to turn it into the "corresponding" rule with exceptions (as illustrated by the example at the end of section 3.1). Of course, defining the exact meaning of "corresponding" rule with exceptions depends on the particular non-monotonic logic being considered, but we can state this informally as turning a universal like "All As are Bs" into the default "Typically, As are Bs". After this is done, we have to see how these rules with exceptions are to be integrated with the result of the weakening function, and ensure that we don't get some unexpected results. Finally, the theorems proved in the last section will not only have to be rephrased,[5] but also be more deeply changed. We want to guarantee for instance, that we do not permissively revise a default theory with one extension, to a default theory that has no extensions.

4 Discussion

As every belief revision theory, permissive belief revision is well suited for situations where the available knowledge changes frequently. These changes may be because the world we are modeling changed or because we acquired more knowledge about it.

Traditional belief revision theories [Nebel, 1990; Fuhrmann, 1991] may produce different results when revising logically equivalent theories with the same formula. For example, the fact that both a and b are true may be represented either by $\{a \wedge b\}$ or by $\{a, b\}$. These two representations will provide different results when revised with $\neg a$. Our theory will give the same result, b in both situations. Although this example might suggest that our approach is syntax-independent, a very desirable property, this does not hold. For instance, $\mathrm{Wk}(a \vee b \Rightarrow c, \{a, \neg c\}) = \top$, but $\mathrm{Wk}((a \Rightarrow c) \wedge (b \Rightarrow c), \{a, \neg c\}) = b \Rightarrow c$.

Using the permissive revision approach, the use of universal rules for representing typical situations does not make strong commitments. Therefore, we can use universal rules to represent typical situations and let the revision theory change them to defaults when the need for doing so arises. This is useful because reasoning with universal rules is simpler from a logical point of view, and when the first exception is found the universal is turned into a default rule.

Permissive revision does not simply give up beliefs. Instead, it *weakens* them into beliefs that are consistent with the rest of our knowledge but that still express as much information as they can. Permissive belief revision allows to keep more beliefs than what is usual when using a foundations theory.[6] This has the advantage of being more compliant with what humans do [Harman, 1986],[7] while keeping all the advantages of using a foundations theory, namely the fact that we still have an explanation for each belief that we hold.

[5] The result of weakening a formula will no longer consist of just a classical formula, but probably of a set containing one classical formula and some rule(s) with exceptions.

[6] The distinction between coherence theories and foundations theories is presented in [Harman, 1986]. And [Doyle, 1992] presents a good comparison between them.

[7] This was also documented by psychological experiments in [Hoenkamp, 1987].

Part of this work was implemented on top of SNePSwD [Cravo and Martins, 1993], the implementation of a truth maintenance system [Cravo, 1995][8] based on a foundations belief revision theory [Cravo, 1993a] and a non-monotonic logic SWMC [Cravo, 1993b]. The development and implementation of an early version of this work is described in [Cachopo, 1997].

We should point out that there are some situations where the ability to adapt new beliefs to the existing knowledge, instead of simply discarding them, can be particularly useful:

- When incrementally building a knowledge base, it is likely that whoever is doing it doesn't have complete knowledge from the beginning. So, the knowledge will almost certainly have to be changed. If these changes are made using permissive belief revision, they will be less sensitive to representation errors, because the existing knowledge will be adapted to the knowledge base, instead of simply being discarded.
- The task of merging several knowledge bases can be simplified, because when there is conflicting knowledge, it can be adapted in a way that still keeps as much knowledge as possible.

5 Future work

One of the aspects that we are still studying is the formalization of what it means to weaken a universal rule both in a monotonic and a non-monotonic logic, as discussed in section 3.3.

It is also important to study the properties of the permissive revision, namely which postulates it satisfies. Because we can use a non-monotonic logic, the AGM postulates will not be satisfied. The postulates that the permissive belief revision operation satisfies have not yet been studied.

The *structured theories* non-monotonic formalism developed in [Ryan, 1991] is used to define a new belief revision operation, which produces results that resemble ours. However, Ryan's approach is mainly semantic, while ours is syntactic. Although the two are not equivalent, something might be gained in understanding the differences between them. Understanding what weakening a formula is, in terms of models, could give us new insight in some of the problems we have now, such as the weakening of the universally quantified formulas.

We intend to continue the implementation of the weakening of each type of formula on SNePSwD. A very simple program that weakens formulas was implemented, but it includes neither the belief revision theory nor the underlying logic.

The belief revision theory described in [Cravo, 1993a] uses preferences to choose which rule(s) to abandon during a contraction or a revision operation. We think it is possible to use the same mechanism to select the rules to be weakened, but this still needs more study.

[8] The guarantee that each belief has (at least) one justification allows the use of a truth maintenance system to efficiently make the changes proposed by the belief revision theory.

6 Acknowledgements

This work was partially supported by project PRAXIS XXI 2/2.1/TIT/1568/95.

References

[Alchourrón *et al.*, 1985] Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: partial meet functions for contraction and revision. *The Journal of Symbolic Logic*, 50(2):510–530, 1985.

[Cachopo, 1997] Ana Cardoso Cachopo. Revisão adaptativa de crenças. Master's thesis, Instituto Superior Técnico, Universidade Técnica de Lisboa, Lisbon, Portugal, 1997.

[Cravo and Martins, 1993] Maria R. Cravo and João P. Martins. SNePSwD, a newcomer to the SNePS family. *Journal of Experimental and Theoretical Artificial Intelligence*, 5:135–148, 1993.

[Cravo, 1993a] Maria R. Cravo. A belief revision theory based on SWMC. Technical Report GIA 93/03, Instituto Superior Técnico, Universidade Técnica de Lisboa, Lisbon, Portugal, November 1993.

[Cravo, 1993b] Maria R. Cravo. SWMC: A logic for default reasoning and belief revision (a new version). Technical Report GIA 93/02, Instituto Superior Técnico, Universidade Técnica de Lisboa, Lisbon, Portugal, November 1993.

[Cravo, 1995] Maria R. Cravo. BRS: A belief revision sistem capable of non monotonic reasoning and belief revision. Technical Report GIA 95/01, Instituto Superior Técnico, Universidade Técnica de Lisboa, Lisbon, Portugal, February 1995.

[Doyle, 1992] Jon Doyle. Reason maintenance and belief revision: Foundations vs. coherence theories. In Peter Gärdenfors, editor, *Belief Revision*, number 29 in Cambridge Tracts in Theoretical Computer Science, pages 29–51. Cambridge University Press, 1992.

[Fuhrmann, 1991] André Fuhrmann. Theory contraction through base contraction. *Journal of Philosophical Logic*, 20(2):175–203, 1991.

[Harman, 1986] Gilbert H. Harman. *Change in View: Principles of Reasoning.* MIT Press, Cambridge, MA, 1986.

[Hoenkamp, 1987] Edward Hoenkamp. An analysis of psychological experiments on non-monotonic reasoning. In *Proceedings of IJCAI-87*, pages 115–117, Milan, Italy, 1987. Morgan Kaufmann Publishers, Inc.

[Lehmann, 1995] Daniel Lehmann. Belief revision revised. In *Proceedings of IJCAI-95*, pages 1534–1540, Montreal, Canada, 1995. Morgan Kaufmann Publishers, Inc.

[McCarthy, 1980] John McCarthy. Circumscription: A form of nonmonotonic reasoning. *Artificial Intelligence*, 13(1–2):27–39, 1980.

[McCarthy, 1986] John McCarthy. Applications of circumscription to formalizing common sense reasoning. *Artificial Intelligence*, 28:89–116, 1986.

[Moore, 1983] Robert C. Moore. Semantical considerations on nonmonotonic logic. In *Proceedings of IJCAI-83*, pages 272–279, Karlsruhe, FRG, 1983. Morgan Kaufmann Publishers, Inc.

[Moore, 1985] Robert C. Moore. Semantical considerations on nonmonotonic logic. *Artificial Intelligence*, 25(1):75–94, 1985.

[Moore, 1988] Robert C. Moore. Autoepistemic Logic. In P. Smets, E. H. Mamdani, D. Dubois, and H. Prade, editors, *Non-Standard Logics for Automated Reasoning*, pages 105–136. Academic Press, New York, USA, 1988.

[Nebel, 1989] Bernhard Nebel. A knowledge level analysis of belief revision. In *Proceedings of KR-89*, pages 301–311. Morgan Kaufmann Publishers, Inc., Toronto, Canada, 1989.

[Nebel, 1990] Bernhard Nebel. *Reasoning and revision in hybrid representation systems*, volume 422 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag, Heidelberg, Germany, 1990.

[Poole, 1988] David Poole. A logical framework for default reasoning. *Artificial Intelligence*, 36(1):27–47, 1988.

[Quine and Ullian, 1978] Willard Van Orman Quine and Joseph S. Ullian. *The Web of Belief*. Random House, 2 edition, 1978.

[Reiter, 1980] Raymond Reiter. A logic for default reasoning. *Artificial Intelligence*, 13(1–2):81–132, April 1980.

[Ryan, 1991] Mark D. Ryan. Defaults and revision in structured theories. In *Proceedings of the Sixth IEEE Symposium on Logic in Computer Science (LICS)*, pages 362–373. Morgan Kaufmann Publishers, Inc., 1991.

[Witteveen *et al.*, 1994] Cees Witteveen, Wiebe van der Hoek, and Hans de Nivelle. Revision of non-monotonic theories: Some postulates and an application to logic programming. In *Proceedings of JELIA-94*, number 838 in Lecture Notes in Artificial Intelligence, pages 137–151, York, UK, 1994. Springer-Verlag.