Anytime Belief Revision

Mary-Anne Williams
Information Systems
Department of Management
University of Newcastle
Australia
maryanne@frey. newcastle. edu. au

Abstract

Belief Revision is a ubiquitous process underlying many forms of intelligent behaviour. The AGM paradigm is a powerful framework for modeling and implementing belief revision systems based on the principle of Minimal Change; it provides a rich and rigorous foundation for computer-based belief revision architectures. Maxi-adjustment is a belief revision strategy for theory bases that can be implemented using a standard theorem prover, and one that has been used successfully for several applications. In this paper we provide an *anytime decision procedure* for maxi-adjustments, and study its complexity. Furthermore, we outline a set of guidelines that serve as a protomethodology for building belief revision systems employing a maxi-adjustment. The algorithm is under development in the belief revision module of the CIN Project.

1 Introduction

Belief Revision underlies many forms of intelligent behaviour. An intelligent agent must be adept at revising its beliefs in a rational way. The AGM paradigm, so named after its founders Alchourron, Gardenfors and Makinson [1985], is a powerful theoretical framework for modeling and implementing belief revision systems; it provides a rich and rigorous foundation for principled computer-based architectures that endow agents with the ability to change their beliefs in a coherent and rational fashion.

Gardenfors and Makinson [1988] provided a constructive means for defining revision functions based on an epistemic entrenchment ordering of a reasoning agent's beliefs. Furthermore, they showed that there is a one-to-one relationship between revision functions and epistemic entrenchment orderings.

Iterated revision can be achieved by transmuting epistemic entrenchment orderings where the emphasis is not exclusively on acceptance and removal of beliefs from a theory, but also on raising and lowering of the degree of acceptance of beliefs. Raising the degree of acceptance of a belief corresponds to a revision, whilst lowering it corresponds to a contraction.

Maxi-adjustment is a specific strategy for implementing belief revision systems. It strives for maximal inertia of information under change, and was devised by Williams [1996]. It has been shown to be successful in applications where the systems designer or knowledge engineer is able to specify dependencies among beliefs [MacNish and Williams, 1996]. In essence, when incoming information is inconsistent with the agent's knowledge, a maxi-adjustment retracts only those minimally entrenched beliefs that are inconsistent with the new information.

In this paper we provide an anytime decision procedure for maxi-adjustments. Furthermore, we outline a set of guidelines that serve as a protomethodology for building belief revision systems. The algorithm is under development in the belief revision module of the CIN Project; a project that is seeking to develop an intelligent information management toolkit [Antoniou and Williams, 1996].

Section 2 outlines belief revision in the AGM paradigm. Section 3 discusses several important modeling problems that arise when AGM change functions are used in practice, it thence outlines how maxi-adjustments overcome them. Section 4 describes subsumption removal; an optional feature of maxi-adjustment that can be used to enhance its performance. In section 5 we give an anytime algorithm for maxi-adjustments, and in section 6 we discuss its complexity. In section 7 we make several methodological remarks concerning the design and development of belief revisions systems that employ the maxi-adjustment strategy.

2 Belief Revision

An intelligent information system must possess the ability to revise its knowledge base when it receives new information. The AGM paradigm has become one of the standard frameworks for modeling change.

Let us begin with some technical preliminaries: $\mathcal{L}$ will denote a (possibly first-order) language which contains a complete set of Boolean connectives. We will denote sentences in $\mathcal{L}$ by lower case Greek letters. We assume $\mathcal{L}$ is governed by a logic that is identified with its consequence relation $\vdash$ which is assumed to satisfy the following conditions [Gärdenfors, 1988]: (a) if α is a truth-functional tautology, then $\vdash \alpha$, (b) if $\vdash (\alpha \rightarrow \beta)$ and $\vdash \alpha$, then $\vdash \beta$ (*modus ponens*), (c) $\vdash$ is consistent, i.e. $\not\vdash \bot$, where $\bot$ denotes the inconsistent theory, (d) $\vdash$ satisfies the deduction theorem, and (e) $\vdash$ is compact.

The set of all logical consequences of a set $T \subseteq \mathcal{L}$, i.e. $\{\alpha : T \vdash \alpha\}$, is denoted by $\text{Cn}(T)$. A *theory* of $\mathcal{L}$ is any subset of C closed under Cn. We let L^{∞}, pronounced 'elbow', denote the set of *contingent* sentences.

Within the AGM paradigm a body of information is represented as a theory, and informational changes are regarded as transformations on theories. The principal AGM functions are contraction and revision. They can be described using the well known AGM rationality postulates, and individual contraction and revision functions can be uniquely determined by any of the several standard constructions, e.g. the epistemic entrenchment ordering construction. Both the postulates and the constructions attempt to encapsulate the principle of Minimal Change. The *magnitude of change* may not be based on set inclusion measures; sometimes the most rational response to avoiding inconsistency is to forfeit more than the minimal number of beliefs, e.g. it may be better to retract several weakly held beliefs than to surrender a single strongly held belief.

A *revision,* T<J, attempts to change a theory T to incorporate a so that the resultant theory is consistent provided a itself is consistent. A *contraction,* T^{-}_{α} involves the removal of a set of sentences from T so that a nontautological sentence is no longer implied. A *withdrawal* function [Makinson, 1987] is a generalised contraction function in that it satisfies all but the most notorious postulate for contraction, namely *recovery;* the property T = Cn(Tα U a). It has been argued in the literature that recovery is not always appropriate for a limited reasoner, however it is one of the most important postulates for capturing the notion of *minimal change* when information is given up.

3 Implementing Belief Revision

For the purpose of developing an implementation of AGM change functions, Gardenfors and Makinson's [1988] work was a significant breakthrough. They showed that an epistemic entrenchment ordering (certain total preorder on the sentences in the language) can uniquely determine how the system will react to the pressures of impinging information. In order to develop computational models based on the entrenchment construction two obvious problems must be overcome: first an epistemic entrenchment ordering has to be propagated by the change function, and second a finite representation for epistemic entrenchment orderings is needed.

We use partial entrenchment rankings[1] as our representation of well-ranked epistemic entrenchment orderings, and we model iterated belief revision by propagating these rankings using a maxi-adjustment; a procedure described in Williams [1996].

3.1 Partial Entrenchment Rankings

Finite partial entrenchment rankings will be sufficient for present purposes. They represent *finite* epistemic entrenchment orderings of £ where the elements of a finite (a not necessarily closed) set of sentences are mapped to the natural numbers.

Definition: *A* finite partial entrenchment ranking *is a function* B *from a finite subset of sentences in* L *into the* *natural numbers* $\mathcal{N}$ *such that the following conditions are satisfied for all* $\alpha \in \mathrm{dom}(\mathbf{B})$*:*

(PER1) *If* $\not\vdash \alpha$ *then* $\{\beta \in \mathrm{dom}(\mathbf{B}) : \mathbf{B}(\alpha) < \mathbf{B}(\beta)\} \not\vdash \alpha$.

(PER2) *If* $\vdash \neg\alpha$ *then* $\mathbf{B}(\alpha) = 0$.

(PER3) $\vdash \alpha$ *if and only if* $\mathbf{B}(\alpha) = \max(\mathbf{B}(\mathrm{dom}(\mathbf{B})))$.

The higher the integer assigned to a sentence by a partial entrenchment ranking the more firmly it is held. **(PER2)** says that inconsistent information is assigned zero, and **(PER3)** says that the tautologies alone are assigned the highest rank. The most piquant property of a ranking, is given in **(PER2)**, namely a nontautological sentence α cannot be entailed by sentences ranked strictly higher than α itself.

We refer to $\mathbf{B}(\alpha)$ as the degree of acceptance of α. The intended interpretation of a partial entrenchment ranking is that sentences assigned a degree of acceptance greater than zero form the agent's *explicit* beliefs of the system, and their logical closure form his *implicit* beliefs.

Definition: *Define the* explicit information content *of a ranking* $\mathbf{B}$ *to be* $\{\alpha \in \mathrm{dom}(\mathbf{B}) : \mathbf{B}(\alpha) > 0\}$, *and denote it by* $\mathbf{exp}(\mathbf{B})$. *Define the* implicit information content *of* $\mathbf{B}$ *to be* $\mathrm{Cn}(\mathbf{exp}(\mathbf{B}))$, *and denote it by* $\mathbf{content}(\mathbf{B})$.

Typically, we are not only interested in the degree of acceptance of explicit information, but also in the degree of acceptance of sentences they entail. A partial entrenchment ranking represents a system's incomplete preferences from which a complete entrenchment ranking can be generated. There could, well, be an infinite number of entrenchment rankings which are compatible with a given partial specification. The following function '**degree**' derives the minimum possible degree of acceptance for implicit information as specified by a partial entrenchment ranking.

Definition: *Let* $\alpha \in \mathcal{L}^{\bowtie}$. *If* $\mathbf{B}$ *is a finite partial entrenchment ranking, then define* $\mathbf{degree}(\mathbf{B}, \alpha)$ *to be*

$$\begin{cases} \text{largest } j \text{ such that } \{\beta \in \mathbf{exp}(\mathbf{B}) : \mathbf{B}(\beta) \geq j\} \vdash \alpha \\ \qquad\qquad \text{if } \alpha \in \mathbf{content}(\mathbf{B}) \\ 0 \qquad\quad \text{otherwise} \end{cases}$$

Example: If $\mathbf{B}(\alpha \rightarrow \beta) = 3$, $\mathbf{B}(\alpha) = 2$, $\mathbf{B}(\gamma) = 1$, then we can compute $\mathbf{degree}(\mathbf{B}, \beta) = 2$, $\mathbf{degree}(\mathbf{B}, \alpha \wedge \beta) = 2$, $\mathbf{degree}(\mathbf{B}, \alpha \vee \gamma) = 2$, $\mathbf{degree}(\mathbf{B}, \alpha \wedge \gamma) = 1$, $\mathbf{degree}(\mathbf{B}, \delta) = 0$, $\mathbf{degree}(\mathbf{B}, \neg\delta) = 0$.

3.2 Maxi-Adjustments

The (α, i)–maxi-adjustment of the partial entrenchment ranking $\mathbf{B}$, is denoted $\mathbf{B}^{*}(\alpha, i)$, it modifies $\mathbf{B}$ by assigning α the new ranking i whilst maintaining the properties of a partial entrenchment ranking. It is formally defined in [Williams, 1996].

The input i can be specified as an integer, as a number in a specified range, e.g. [0,1], as $\mathbf{B}(\beta)$ for some β which would mean give the new information α the same rank as β, or in linguistic terms depending on the application at hand.

Maxi-adjustments are motivated by Spohn's notion of a *reason*. According to Spohn [1983] β is a *reason* for α if and only if raising the epistemic rank of β would raise the epistemic rank of α.

Maxi-adjustment is based on the idea that if you believe a reason your coffee cup is leaking is that it has a hole, then whereupon closer inspection you discover its not leaking at all, it would seem a rational response

[1]Partial entrenchments were defined in [Williams 1995], and essentially identical representations can be found in [Dubios et al 1994, Rott 1992, Williams 1992], and elsewhere.

to retract *your cup has a leak* and *your cup has a hole*. Contrariwise, if you are unaware that holes are reasons for leaks, then you could quite happily continue to believe *your cup has a hole* even when you discover that it is not leaking. In other words, common sense suggests that information should be retracted only if there is good reason to do so.

To evaluate reasons, it seems eminently sensible to require that raising the rank of α should disturb the agent's background ranking as little as possible. It is easy to show that whenever α and β are in $\mathbf{exp}(\mathbf{B})$, if we simply enforce the properties of a partial entrenchment ranking then β is a reason for α if and only if $\mathbf{degree}(\mathbf{B}, \beta \rightarrow \alpha) > \mathbf{B}(\alpha)$. Maxi-adjustments aim to use a closed world assumption with respect to reasons: if it cannot be derived that β is a reason for α, then it is assumed that β is not a reason for α. The system designer specifies reasons by ensuring that $\mathbf{degree}(\mathbf{B}, \beta \rightarrow \alpha)$ is strictly larger than $\mathbf{degree}(\mathbf{B}, \alpha)$. All nontautological reasons are defeasible, and there may be reasons for reasons.

Like most revision procedures on theory bases maxi-adjustments are syntactically dependent on $\mathbf{exp}(\mathbf{B})$. Despite this syntax dependence it can be shown that (i) $\mathbf{B}^\star(\alpha, i)$ is a partial entrenchment ranking in which α is assigned the degree of acceptance i, (ii) if i is greater than zero then $\mathbf{content}(\mathbf{B}^\star(\alpha, i))$ is an AGM revision $(\mathbf{content}(\mathbf{B}))^*_\alpha$, (iii) $\mathbf{content}(\mathbf{B}^\star(\alpha, 0))$ is an AGM withdrawal function, and (iv) $\mathrm{Cn}(\mathbf{B}^\star(\alpha, 0) \cup \neg\alpha) \cap \mathbf{content}(\mathbf{B})$ is an AGM contraction $(\mathbf{content}(\mathbf{B}))^-_\alpha$.

The change functions defined above by maxi-adjustment may not be the same as those obtained via Gärdenfors and Makinson's construction using the epistemic entrenchment ordering derived in the obvious way from the relative ordering given by **B** using the function **degree**. Ordinary adjustment [Williams 1995] is a procedure that complies exactly with the standard entrenchment construction. There exist several variants of maxi-adjustment, all are minor variations of the main algorithm described in the next section. Our implementation [Williams and Williams, 1997] offers anytime algorithms for all of the variants. By performing a so-called hybrid-adjustment (an adjustment followed by a maxi-adjustment) we can guarantee that at least as much information is retained as would be using the standard entrenchment procedure. In practice hybrid-adjustments and maxi-adjustments retain a great deal more information than the standard recipe.

Other properties of maxi-adjustment include: (i) every ranking is *reachable* in a finite language via a finite number of maxi-adjustments, (ii) maxi-adjustments only use the relative ranking of sentences, (iii) maxi-adjustments *preserve finiteness*; $|\mathbf{exp}(\mathbf{B}^\star(\alpha, i))| \leq |\mathbf{exp}(\mathbf{B})| + 1$, (iv) maxi-adjustment does not reassign explicit beliefs holding ranks greater than $\max(\{i, \mathbf{B}(\alpha), \mathbf{B}(\neg\alpha)\})$, i.e. it preserves information that is more inportant than the incoming information, (v) moving a sentence up j ranks in one step results in the same ranking as 'jacking it up' j ranks one rank at a time, and (vi) similarly, moving a sentence down j ranks in one step results in the same ranking as moving it down j ranks one rank at a time.

Finally, based on the well known connections between belief revision, possibilistic reasoning and nonmonotonic reasoning, maxi-adjustments can be applied to possibilistic knowledge bases [Dubois *et al*, 1994], and nonmonotonic knowledge bases [Gärdenfors and Makinson 1994, Lehmann 1992].

4 Subsumption Removal

The idea of using reasons to determine what information to retract during the contraction of a sentence α works well for all beliefs except for some ranked at $\mathbf{degree}(\mathbf{B}, \alpha)$. For example, if $\beta \rightarrow \alpha$, β and α are equally ranked, then neither $\beta \rightarrow \alpha$ nor β is a Spohnian reason for α, yet belief in both cannot be maintained if α is to be contracted. In this case the Gärdenfors and Makinson construction would remove both $\beta \rightarrow \alpha$ and β. A standard maxi-adjustment is designed to do the same, however it also provides an *optional* procedure that uses *subsumption removal* which if selected removes all the beliefs subsumed by the information to be contracted. For the case in question above, this would result in the retraction of $\beta \rightarrow \alpha$ in preference to β because $\beta \rightarrow \alpha$ is subsumed by α.

Subsumption removal is not always appropriate, has a computational cost, and should be used with caution. When contracting (or lowering) a belief α using a maxi-adjustment we justify removing beliefs subsumed by α on the following grounds:

(i) If **B** is interpreted to be a partial specification of an agent's epistemic entrenchment, then it will, in general, be compatible with several epistemic entrenchments (which could, in principle, be generated using a different mechanism for **degree**). If $\mathbf{B}(\beta \rightarrow \alpha) = \mathbf{degree}(\mathbf{B}, \alpha)$ then there may exist a compatible epistemic entrenchment ordering that would retain β using the standard Gärdenfors and Makinson construction, but no such compatible ordering could exist that would retain $\beta \rightarrow \alpha$[2]. If $\alpha \vee \beta$ is explicit and $\mathbf{B}(\alpha \vee \beta) = \mathbf{degree}(\mathbf{B}, \alpha)$ then we may prefer not to use subsumption removal because it could be interpreted to mean that the user definitely wants β removed when α is - this situation can be trivially detected, and built into the algorithm.

(ii) Since $\alpha \vdash (\beta \rightarrow \alpha)$ the sentence $\beta \rightarrow \alpha$ does not add any *epistemic power* to the ranking if α itself is an explicit belief. In other words, the same epistemic entrenchment ordering would be generated using the function **degree** irregardless of whether $\beta \rightarrow \alpha$ is explicit. If subsumption removal is applied then whenever α is an implicit but not explicit belief the maxi-adjustment algorithm makes it explicit. In other words, reducing the ranking of α will not remove α even if it wasn't previously explicit, except of course when the new rank is zero.

5 Anytime Maxi-Adjustments

An anytime algorithm is one that if interrupted has constructed a partial solution that approximates the actual solution, and the longer it runs the better the approximation. The proposed algorithm for maxi-adjustment is

[2] In other words, if $\alpha = \beta \rightarrow \alpha$, then $\alpha \vee \beta > \alpha$ is possible, but $\alpha \vee (\beta \rightarrow \alpha) > \alpha$ would violate the postulates of entrenchment.

anytime because it constructs the desired resultant partial entrenchment ranking in a top down iterative fashion refining the partly constructed ranking until it converges on the actual ranking. Furthermore, it captures important beliefs before less important beliefs, by rebuilding the ranking from the highest rank to the lowest.

Put simply, the algorithm consists of two main phases: First find the largest cut of the ranking that is consistent with the sentence to be moved, and second salvage as much of the remainder of the ranking by giving preference to higher ranked beliefs.

An important and endearing feature of the algorithm is that after completing its first phase the content of the ranking constructed so far is guaranteed to satisfy the AGM postulates for withdrawal and revision. Furthermore, there is enough information to construct a contraction function as well.

The second phase refines the ranking, in the sense that it recaptures as many beliefs as possible by removing only the minimal subsets at each rank that together with the higher ranked beliefs to be kept entail the information to be contracted.

For simplicity of exposition, and without loss of generality, nonbeliefs and tautologies are not explicit in our rankings, in addition we focus on the principle case of movement of contingent beliefs within the ranking.

The algorithm requires the services of a standard theorem prover to implement references to entailment in the function degree, and in the generation of minimal subsets that entail sentences to be moved down the ranking via the movedown procedure. Before the procedure can decide whether to raise or to lower the degree of acceptance of the sentence α it computes its current degree of acceptance. If its degree is to decrease, then movedown is performed. If its degree is to increase then -α is decreased first (if necessary) using movedown, and then a is moved up the ranking using moveup. The procedure maxi-adjustment(a,i, B,newB) is the (α,i)-maxi-adjustment of B resulting in the formation of the new ranking newB. The highest integer given to any sentence in the domain of B is assigned to the global variable maxjdegree.

If the **maxi-adjustment** procedure is interrupted then **newB** constructed so far is returned and if $i > 0$ then we check that α has been added to rank i, if not we do so.

PROCEDURE: maxi-adjustment(α, i, B, newB)
Input: A partial entrenchment ranking **B**, a sentence α, and an integer i which represents the new desired rank of α.

if $\alpha \in$ dom(**B**) **then degree_α := B(α)**
else degree_α := degree(B, α)
end if
max_degree = max(**B**(dom(**B**)))
case of degree_α:
greater than i:
movedown(α, degree_α, i, B, newB).
lower than i:
if degree_α = 0
then degree_not_α := **degree(B, $\neg\alpha$)**
else degree_not_α = 0
end if
if degree_not_α > 0 then
movedown(α, degree_not_α, 0, B, newB)
end if
B := newB
moveup(α, i, B, newB)
otherwise:
newB := B
end case

Output: The new partial entrenchment ranking **newB**.

FUNCTION: degree(B, α)
Input: A partial entrenchment ranking **B**, and a sentence α.

degree := **max_degree**+1
do degree := degree -1
until $\{\beta$: degree $\geq$ **B**$(\beta)\} \vdash \alpha$ **or** degree = 0
assign degree(α, B) := degree

Output: The degree of α in the ranking **B**.

PROCEDURE: movedown(α, j, i, B, newB)
Input: A partial entrenchment ranking **B**, a sentence α, an integer j representing the current degree of α in **B**, and an integer i representing the desired new degree of α.

for k = **max_degree** down to $j+1$ **do**
newB$(\beta) := k$
for all sentences β such that **B**$(\beta) = k$.
for $k = j$ down to i **do**
- generate minimal subsets of sentences at rank k that together with dom(**newB**) entail α
- **newB**$(\beta) := k$ for all β not in any such minimal subset.

Output: The new partial entrenchment ranking **newB**.

PROCEDURE: moveup(α, i, B, newB)
Input: A partial entrenchment ranking **B**, a sentence α, a integer i representing the desired new degree of α.

for k = **max_degree** down to $i+1$ **do**
newB$(\beta) := k$
for all sentences β such that **B**$(\beta) = k$.
newB$(\alpha) := i$
for $k = i-1$ down to 1 **do**
newB(β) := **degree(newB**$, \alpha \to \beta$**)**.

Output: The new partial entrenchment ranking **newB**.

Subsumption removal is an optional feature of the maxi-adjustment procedure. There are several ways to use it: (i) eliminate subsumed beliefs before the minimal subsets are generated, or (ii) eliminate subsumed beliefs from the minimal subsets after they have been found. The first approach will, in general, result in more beliefs being removed than the second, and will also reduce the number of minimal subset that must be generated. To incorporate the first approach (the second is equally obvious) we insert the following instruction between the first and second **for** loop constructs in the **movedown** procedure, so that subsumed sentences are removed before the minimal subsets are generated: **remove** sentences in dom(**B**) **if** **B**(β) = **degree**$(\alpha$, **B**$)$ and $\alpha \vdash \beta$. Then after the second **for** construct we add: **if** $i > 0$ **then newB**$(\alpha) := i$

6 Complexity

The procedure maxi-adjustment(B, a, i,newB) returns a revised ranking newB for any allocation of computational time. The longer the algorithm runs the closer the ranking newB approximates B*(α, *i).*

First-order logics satisfy the conditions required of the underlying logic given in section 2, hence the AGM framework and the proposed algorithm supports changes to rich knowledge bases, in principle. However, it is well known that satisfiability in first-order languages is undecidable, consequently nontrivial belief revision algorithms that are guaranteed to terminate cannot be constructed. The maxi-adjustment decision procedure described herein is an anytime algorithm, so it can be used to generate an infinite sequence of better and better approximations to revision and contraction functions.

As noted in the previous section the maxi-adjustment procedure essentially consists of two main phases.

The first phase is computationally easier than the second and if it is completed then we are guaranteed to satisfy the AGM postulates for revision and withdrawal, and notably we have enough information to construct a contraction function, if desired. It is interesting to note that for query evaluations, such as is *0* € B*(α, i), only the first phase need be carried out.

The function degree is the workhorse of the first phase. Computing degree(B,α) is NP-hard. If a polynomial fragment of propositional logic is used, then computing degree is polynomial. Our anytime algorithm uses a top-down strategy that, if interrupted will always err on the side of overestimating the degree of sentences which in turn will never lead to inconsistency. A purely bottom up procedure would not exhibit this behaviour. However a hybrid strategy that combined a top-down and a bottom-up binary search would be more efficient on average, than a purely top-down linear search for the degree of a, and if B has *n* natural partitions then it requires [$\log_2$ *n]* satisfiability checks [Lang, 1997]. An interpolation strategy that used information about ranking's history, or information available from the application at hand would also improve the performance of the function degree. Hybrid and informed techniques have been investigated in Lang [1997]. If we adopt a hybrid algorithm for degree in our anytime algorithm then should the program be interrupted we simply use the most recent upper bound to calculate the resultant ranking.

The first phase determines the most important core of the ranking to survive the change, the second phase refines it by maintaining as many other beliefs as possible based on the original ranking. The second stage of the maxi-adjustment is in A_2, and hence solvable with a polynomial number of calls to an NP oracle (c.f. [Nebel 1991, Eiter and Gottlob 1992]). The worse case arises when all explicit beliefs are equally ranked. The computational cost decreases as the number of ranks increase. So the more discerning the agent the easier it is for him to modify his beliefs using a maxi-adjustment. This property concords with our intuition, i.e. it seems psychologically plausible, but not all revision strategies exhibit it. For example, ordinary adjustment which is based on the standard entrenchment construction does not.

7 Methodological Remarks

7.1 Contraposition

The use of maxi-adjustment presupposes that the knowledge engineer is able to identify reasons. The inability to identify reasons simply means that more information than is perhaps intended is retained in practice. If *0* is a reason for *α* then *0* —► *α* should be placed higher in the ranking than a. Using material implication in this way has an important ramification, namely *contraposition of reasons,* i.e. whenever A is a (whole) reason for B, then B is a (whole) reason for ¬«α. For example, if one of the reasons my hang-guilder is ascending is that the up-lift is sufficient to overcome the effects of gravity, then one of the reasons it is not ascending is that the up-lift is insufficient. Consequently, reasons *do not capture causality* in a broader context.

Another effect of contraposition is that if *0* is a reason for a, then when an agent revises by adding ->*α* he will, if only implicitly, ¬baccept . Maxi-adjustment can also be used to model applications in which contraposition of reasons is not an appropriate assumption. This is achieved by breaking down the changes to the ranking into more primitive operations, and composing a *transaction* on the ranking!

7.2 Some Guidelines

Knowledge Engineers and System Designers are accustomed to the syntax sensitivity present in prevalent information modelling methods, such as (restricted) logical languages, entity-relationship models, conceptual graphs, etc. Methodologies have been developed for these traditional techniques: they guide the development process to a faithful, hopefully optimal, representation of the application at hand.

Maxi-adjustment is also syntax dependent, and whilst a methodology for using maxi-adjustments is not yet available the following application independent guidelines have helped in the development of several belief revision applications.

(1) Important information should be explicit.

(2) Information should be in its simplest logical form.

(3) The number of ranks should be maximised if incoming information is expected to be inconsistent with highly entrenched information.

(4) Conjunction can be used to bind information items together, e.g. if the application calls for *α* to be removed whenever *0* is and vice versa, then *α ^ B* can be used. If the conjuncts themselves are not explicit, or not derivable from other explicit beliefs then they will stand and fall together.

(5) Represent sentences that do not need to be bound as independent sentences, e.g. if *α* and *0* are not related then using *α* and *0* is preferable to the compound sentence a A *0.*

(6) Irredundant rankings are preferable.

(7) If the set *{0\, 02>..., 0_n]* constitutes a reason for *a* (i.e. their simultaneous satisfaction would mean that Q must hold), then the sentence *0\ A02* A... *$A0_n$* —► a is placed higher in the ranking than a.

(8) Subsumption at the same rank should be avoided, and should only be used to satisfy guideline (1).

Guidelines 4, 5, and 6 are related to data normalisation in database design; a process used to transform a database into a representation that minimises update anomalies. As is commonly the case with sets of guidelines, there are exceptions and some guidelines may be in conflict with one another for a particular application. For example, following (1) may lead to a redundant ranking which clearly offends (6).

Methodologies will have to be developed to support the effective use of belief revision in real-world applications. Their development will be facilitated through experience with implemented prototype systems out in the field. Our aim is, not only, to develop a robust belief revision system but also to assist the user in making design choices. This is partially achieved in our system [Williams and Williams, 1997] by making as many of the consequences of a users ranking representation visible. For example beliefs are highlighted before being moved, changes can be stepped through, reasons can be queried for in advance (since they are determined by the current ranking), and rankings can be unwound and saved using *commit* and *rollback* mechanisms. Rankings can also be modified hypothetically during development and testing.

The anytime algorithm for maxi-adjustment is under development in the CIN Project: a project that is seeking to develop an Intelligent Information Management Toolkit. It offers a suite of sophisticated methods for default reasoning and belief revision. The system is currently founded on an objected-oriented design. The core of the belief revision system is implemented in C++ using a state-of-the-art tableau theorem prover, and a Java based Graphical User Interface that provides facilities for *dropping* and *dragging* sentences up and down a ranking. Several rankings can be manipulated simultaneously.

8 Discussion

Iterated belief revision can be achieved by transmuting a partial entrenchment ranking. In this paper we described an anytime decision procedure for iterated belief revision.

We discussed the complexity of our anytime algorithm. In essence, it possesses two main phases: (i) determine the degree of acceptance for the information to be moved down the ranking, and (ii) remove minimal sets of sentences at each rank that entail the information to be moved down the ranking. The first phase is computationally simpler than the second, and if the first phase is completed before the algorithm is interrupted then we are guaranteed to be able to identify a theory base whose closure satisfies the AGM postulates for revision and withdrawal. Furthermore, the ranking so far contains enough information to construct a theory satisfying the postulates for contraction.

The proposed anytime algorithm has been used for several applications, and as a result of the experienced gained in using our system we were able to provide a domain independent protomethodology for developing belief revision applications based on maxi-adjustment.

References

Alchourr6n, C, Gardenfors, P., and Makinson, D. [1985], *On the logic of theory change: Partial meet functions for contraction and revision,* Journal of Symbolic Logic 50: 510 - 530.

Antoniou, G. and Williams, M.A. [1996], *Reasoning with Incomplete and Changing Information,* in N. Terashima and E. Altman (eds), *Advanced Information Technology Tools,* Chapman and Hill, 395 -401.

Dubois, D., Lang, J., and Prade, H. [1994], *Possibilistic Logic,* Handbook of Logic in Artificial Intelligence and Logic Programming, Gabbay, D., Hogger, C, and Robinson, J. (eds), Claredon Press, Oxford.

Eiter, T and Gottlob, G. [1992], *The Complexity of Propositional Knowledge Base Revision,* Artificial Intelligence 57: 227 - 270.

Gardenfors, P. [1988], *Knowledge in Flux,* A Bradford Book, The MIT Press.

Gardenfors, P., and Makinson, D. [1988], *Revisions of Knowledge Systems using Epistemic Entrenchment,* Proceedings of TARK, 83 - 96.

Gardenfors, P. and Makinson, D. [1994], *Nonmonotonic Inference Based on Expectations,* Artificial Intelligence 65: 197-245.

Lang, J. [1997], *Possibilistic Logic: Algorithms and Complexity* in J. Kohlas and S. Moral (eds), *Handbook of Algorithms for Uncertainty and Defeasible Reasoning,* Kluwer Academic Publishers.

Lehmann, D. [1992], *What should a Knowledge Base Entail,* in Principles of Knowledge Representation and Reasoning, Morgan Kaufmann, 1992.

MacNish, C.K. and Williams M.A. [1996], *From Belief Revision to Design Revision: Applying Theory Change to Changing Requirements,* in Reasoning about Incomplete and Changing Information, Lecture Note Series, Springer Verlag, 1997.

Makinson, D. [1987], *On the Status of the Postulate of Recovery in the Logic of Theory Change,* Journal of Philosophical Logic, 16: 383 - 394.

Nebel, B., [1991] *Belief Revision and Default Reasoning,* in Principles of Knowledge Representation and Reasoning, Morgan Kaufmann, 417 - 428.

Rott, H. [1992] *Preferential Belief Change using Generalised Epistemic Entrenchment,* Journal of Logic, Language and Information, 1: 45 - 78.

Spohn, W., [1983] *Deterministic and Probabilistic Reasons,* Erkenntis 19: 371 - 396.

Williams M.A. [1992], *Two functions for theory base change,* in the Proceedings of the Australian Joint Conference on Artificial Intelligence, p259 - 265.

Williams, M.A. [1995] *Iterated Theory Base Change: A Computational Model,* IJCAI-95, Morgan Kaufmann, 1541 - 1550.

M.A. Williams, [1996] *A Practical Appraoch to Belief Revision: Reason-Based Change,* in L. Aiello, and S. Shapiro (eds), Principles of Knowledge Representation and Reasoning: Proceedings of the Fifth International Conference, Morgan Kaufmann, San Mateo, CA, 412-421, 1996.

Williams, M.A., and Williams, D.M., *A Belief Revision System for the World Wide Web,* IJCAI Workshop on Artificial Intelligence and the Internet, 1997, http://u2.newcastle.edu.au/webworld/papers/ai-internet.html