Belief revision for rule-based agents

Natasha Alechina, Mark Jago and Brian Logan

ABSTRACT. Agents which perform inferences on the basis of possibly unreliable information need an ability to revise their beliefs if they discover an inconsistency. However, most belief revision algorithms for classical reasoners have quite high complexity, which makes them unsuitable for resource-bounded reasoners. We consider belief revision for agents which reason in a simpler logic that the full first-order logic, namely rule-based reasoners. We show that it is possible to define a contraction operation for rule-based reasoners which satisfies all the basic Alchourrón, Gärdenfors and Makinson (AGM) postulates for contraction (apart from the recovery postulate) and at the same time can be computed in linear time. We prove a representation theorem for this operation with respect to the basic AGM postulates (minus recovery), and two additional postulates.

1 Introduction

In this paper, we present an efficient approach to rational belief revision for rule-based agents. Rule-based agents have a knowledge base consisting of *rules* (Horn clauses) and *facts* (ground literals). The facts used by an agent to draw inferences may come from a variety of sources (user input, communication with other agents, downloaded information from various web sources, observations of the agent's environment etc.) and change over time, both as a result of the inference process itself and as a result of the addition and deletion of facts from the agent's knowledge base. In general, it is impossible to ensure that the agent's knowledge base is always consistent. Even if the rules used by the agent make it impossible to derive a fact and its negation from a consistent knowledge base, there is always the possibility of, e.g., derived information being inconsistent with communicated or observed information. This makes a belief revision strategy necessary: the agent needs to have a way of removing enough facts from its knowledge base to make sure that a contradiction is no longer derivable. If an agent is expected to interact with other agents and its environment in real time, this procedure should also be efficient.

This paper is based on the previous research reported in [AJL06, ABH$^+$06]. In [AJL06], we described a belief contraction algorithm for rule-based agents, and showed that it runs in time linear in the size of the agent's knowledge base. We also showed that the operation defined by the algorithm satisfies all but one of the basic Alchourrón, Gärdenfors and Makinson (AGM) postulates. In [ABH$^+$06], we showed how this algorithm can be incorporated in an agent programming language AgentSpeak. In this paper, we define our contraction operation not by means of an algorithm but extensionally, and prove a representation theorem for it.

The rest of the paper is organised as follows. In section 2, we briefly introduce some background information on the principal theories of belief change, and explain why rational belief change operations are generally assumed to apply only to idealised agents (i.e., are at least NP-hard). In section 3 we introduce rule-based agents, and describe a logic which characterises the 'deductive abilities' of a typical forward-chaining rule-based agent. In section 4 we define a contraction operation which we call McAllester contraction, and prove a representation theorem for it. In section 5, we show that McAllester contraction can be computed in time linear in the size of the agent's belief set. We introduce a preference order based on the notion of quality of justifications for beliefs in section 6, and show that it can be re-computed in polynomial time when new beliefs are derived. We briefly survey related work in section 7, and conclude.

2 Background on belief revision

Two main approaches to belief revision have been proposed in the literature: AGM (Alchourrón, Gärdenfors and Makinson) style belief revision as characterised by the AGM postulates [AGM85] and reason-maintenance style belief revision [Doy77]. AGM style belief revision is based on the ideas of coherence and informational economy. It requires that the changes to the agent's belief state caused by a revision be as small as possible. In particular, if an agent has to give up a belief in A, it does not have to give up believing in things for which A was the sole justification, so long as they are consistent with its remaining beliefs. Classical AGM-style belief revision describes an idealised reasoner, with a potentially infinite set of beliefs closed under logical consequence. Reason-maintenance style belief revision, on the other hand, is concerned with tracking dependencies between beliefs. Each belief has a set of justifications, and the reasons for holding a belief can be traced back through these justifications to a set of foundational beliefs. When a belief must be given up, sufficient foundational beliefs have to be withdrawn to render the belief underivable. Moreover, if all the justifications for a belief are withdrawn, then that belief itself should no longer be held. Most implementations of reason-maintenance style belief revision are incomplete in the logical sense, but tractable. A more detailed comparison of the two approaches can be found in, for example, [Doy92].

In the remainder of this section we describe the AGM approach in more detail. As we mentioned earlier, the AGM theory assumes that a reasoner has a *belief set*, which is deductively closed. When new information becomes available, a reasoner must modify its belief set to incorporate it. The AGM theory defines three operators on belief sets: expansion, contraction and revision. *Expansion*, denoted $K + A$, simply adds a new belief A to K and the resulting set is closed under logical consequence. *Contraction*, denoted by $K \dot{-} A$, removes a belief A from the belief set and modifies K so that it no longer entails A. *Revision*, denoted $K \dot{+} A$, is equivalent to expansion if A is consistent with the current belief set, otherwise it minimally modifies K to make it consistent with A, before adding A.

Contraction and revision cannot be defined uniquely, since in general there is no unique maximal set $K' \subset K$ which does not imply A. Instead, the set of 'rational' contraction and revision operators is characterised by the AGM postulates [AGM85]. In this paper, we will consider only the contraction operation, which is essential for restoring consistency to the belief set.

The basic AGM postulates for contraction are:

(K$\dot{-}$1) $K \dot{-} A = Cn(K \dot{-} A)$ (closure)

(K$\dot{-}$2) $K \dot{-} A \subseteq K$ (inclusion)

(K$\dot{-}$3) If $A \notin K$, then $K \dot{-} A = K$ (vacuity)

(K$\dot{-}$4) If not $\vdash A$, then $A \notin K \dot{-} A$ (success)

(K$\dot{-}$5) If $A \in K$, then $K \subseteq (K \dot{-} A) + A$ (recovery)

(K$\dot{-}$6) If $Cn(A) = Cn(B)$, then $K \dot{-} A = K \dot{-} B$ (equivalence)

where $Cn(K)$ denotes closure of K under logical consequence.

The AGM theory assumes that the set of beliefs is closed under logical consequence, and is therefore generally considered to apply only to idealised agents. To model practical (implementable) agents, an approach called belief base revision has been proposed (see for example [Mak85, Neb89, Wil92, Han93, Rot98]). A belief base is a finite representation of a belief set. Revision and contraction operations can be defined on belief bases instead of on logically closed belief sets. However the complexity of these operations ranges from NP-complete (full meet revision) to low in the polynomial hierarchy (computable using a polynomial number of calls to an NP oracle which checks satisfiability of a set of formulas) [Neb94]. The reason for the high complexity is the need to check for classical consistency while performing the operations. One way to define an operation with a feasible complexity is to weaken the language and the logic of the agent so that the consistency check is no longer an expensive operation (as suggested in [Neb92]).

In this paper, we present an approach to belief revision and contraction for rule-based agents which is a synthesis of AGM and reason-maintenance style belief revision. Our approach is tractable while at the same time being complete and rational with respect to the agent's logic.

3 Rules and corresponding logic

We assume that the agent's beliefs are represented in predicate logic, more precisely, in the form of literals and Horn clause rules. We fix a set of predicate symbols $\mathcal{P}$, a set of variables $\mathcal{X}$ and a set of constants $\mathcal{D}$. A literal A is a predicate symbol of n arguments followed by n variables or constants and possibly preceded by a negation symbol '$\neg$'. For example, if `PartOf` is a binary predicate and `Bordeaux` and `France` are constants, then $\mathtt{PartOf}(\mathtt{Bordeaux}, \mathtt{France})$ and $\mathtt{PartOf}(\mathtt{x}, \mathtt{Bordeaux})$ are both literals. When every argument of the predicate symbol in a literal is an element of $\mathcal{D}$, we call the literal a *ground literal*. We consider an agent with a finite set $\mathcal{R}$ of *rules*, which are of the form

$$A_1, \ldots, A_n \to B$$

where $A_1, \ldots, A_n$ $(n \geq 1)$, B are literals. B is called the *consequent*, and each A_i a *premise*, of the rule. An example of a rule is:[1]

$$\mathtt{Region}(\mathtt{x}, \mathtt{y}), \mathtt{PartOf}(\mathtt{y}, \mathtt{z}) \to \mathtt{Region}(\mathtt{x}, \mathtt{z})$$

We do not allow functional symbols in rules. Variables are assumed to be universally quantified.

Given a rule $A_1, \ldots, A_n \to B$, we define an *instance* of the rule as

$$\delta(A_1, \ldots, A_n \to B)$$

where δ is some substitution function from the set of variables of the rule into $\mathcal{D}$. For example, if δ assigns $\mathtt{c} = \mathtt{ChateauLafiteRothschildPauillac}$ to x, `Pauillac` to y, and `Bordeaux` to z, then

$$\delta(\mathtt{Region}(\mathtt{x}, \mathtt{y}), \mathtt{PartOf}(\mathtt{y}, \mathtt{z}) \to \mathtt{Region}(\mathtt{x}, \mathtt{z})) =$$

$$= \mathtt{Region}(\mathtt{c}, \mathtt{Pauillac}), \mathtt{PartOf}(\mathtt{Pauillac}, \mathtt{Bordeaux}) \to \mathtt{Region}(\mathtt{c}, \mathtt{Bordeaux})$$

[1] The rules are adapted from the wine ontology in [MAR+94].

The set of agent's beliefs contains only rules and ground literals. We consider the agent's beliefs when the agent's rules have run to quiescence, i.e., after all the agent's rules have been applied to all the literals in the agent's memory (note that this set is finite if the original set of rules and ground literals is finite). This means that we assume that the agent's beliefs are closed under logical consequence in a logic W which has a single inference rule, generalised modus ponens (GMP):

$$\frac{\delta(A_1), \ldots, \delta(A_n), \quad \forall \bar{x}(A_1 \wedge \ldots \wedge A_n \rightarrow B)}{\delta(B)}$$

where $\bar{x}$ are all the free variables of $A_1 \wedge \ldots \wedge A_n \rightarrow B$, and δ is a substitution function which replaces $\bar{x}$ with constants.

Note that this logic is much weaker than classical. The only new formulas which are derivable from a set of rules and ground literals are new ground literals. Another limitation is that from $A \rightarrow B$ and $\neg A \rightarrow B$ the agent cannot derive B (it cannot reason using excluded middle). Also, B and $\neg B$ do not entail arbitrary formulas.

As an example, assume that an agent has the rules:

R1 `Region(x, y), PartOf(y, z)` → `Region(x, z)`

R2 `Region(x, France)` → ¬`Region(x, Australia)`

and facts:

F1 `Region(c, Pauillac)`

F2 `PartOf(Pauillac, Bordeaux)`

F3 `PartOf(Bordeaux, France)`

F4 `PartOf(Tasmania, Australia)`

From these rules and facts, the agent can derive

F5 `Region(c, Bordeaux)` (from F1, F2, R1)

F6 `Region(c, France)` (from F5, F3, R1)

F7 ¬`Region(c, Australia)` (from F6, R2)

Assume further that the agent is told that Chateau Lafite Rothschild Pauillac is made in Tasmania:

F8 `Region(c, Tasmania)`

This new statement does not directly contradict the agent's beliefs, in the sense that the belief base does not contain a literal $\neg\texttt{Region}(\texttt{c}, \texttt{Tasmania})$, however it does lead to inconsistency, since it derives

F9 $\texttt{Region}(\texttt{c}, \texttt{Australia})$ (from F8, F4, R1)

which is inconsistent with **F7**. The agent now needs to revise its beliefs to restore consistency. It needs to contract either by $\texttt{Region}(\texttt{c}, \texttt{Australia})$ or by $\neg\texttt{Region}(\texttt{c}, \texttt{Australia})$.

In what follows, we assume that the agent prefers some beliefs to others. For example, it may trust communicated information less than the information in its original knowledge base. Those preferences will be used to decide which beliefs to remove to restore consistency. To make the notion of preference formal, we assume that there is a preference order $\preceq$, which is a total order on the set of ground literals. For example, each literal A may be assigned a numerical degree of preference $p(A)$, and $A \preceq B$ if $p(A) < p(B)$; if $p(A) = p(B)$ for two different literals A and B, we can use a lexicographical order to decide which one precedes the other in $\preceq$. For each finite set of ground literals Γ, we define $w(\Gamma)$ to be the element of Γ which is minimal with respect to $\prec$.

4 McAllester contraction

In this section, we define a contraction operation for rule-based reasoners. We have chosen to allow only contraction by literals, and not by rules. For many rule-based agents it is reasonable to suppose that the agent's rules are themselves not open to revision: for example, if the rules constitute certain knowledge about the domain, e.g., ontological rules, or if they constitute its program and so cannot safely be revised. On the other hand, facts or literals may be acquired from multiple, often not very reliable sources, and are a possible source of inconsistencies.

Let us denote the set of agent's beliefs (rules and ground literals) by K. For two ground literals $\delta(A)$ and $\delta(B)$, let us say that $\delta(B)$ depends on $\delta(A)$ in K, in symbols $\delta(A) \gg_K \delta(B)$, if either:

1. $\delta(A) = \delta(B)$; or,
2. $\forall \bar{x}(A_1, \ldots, A_n \rightarrow B) \in K$, $\delta(A_1), \ldots, \delta(A_n) \in K$, and $\delta(A)$ is the least preferred premise of the rule instance $\delta(A_1, \ldots, A_n \rightarrow B)$, formally: $\delta(A) = w(\delta(A_1), \ldots, \delta(A_n))$, or
3. $\forall \bar{x}(A_1, \ldots, A_n \rightarrow C) \in K$, $\delta(A_1), \ldots, \delta(A_n) \in K$, $\delta(A) = w(\delta(A_1), \ldots, \delta(A_n))$, and $C \gg_K B$.

This notion of dependence is different from entailment. In order for $A \gg_K B$ to hold, A and B have to be in K, B should be derivable from K, A should be a literal which is actively involved in the derivation of B, and, in addition, it has to be involved as *the weakest* premise of some rule used in the derivation.

DEFINITION 1. The *McAllester* contraction of K by a literal A, $K \dot{-} A$, is defined as

$$K \dot{-} A =_{df} K \backslash \{C : C \gg_K A\}$$

The motivation behind this definition of contraction is simple: to contract by A, we need to remove at least one premise of every rule instance which can be used to derive A; and we choose to remove those beliefs which are least preferred. The set of literals which is removed as a result of this contraction does not necessarily have minimal cardinality, but it is minimal in the following order on subsets of K:

$$\Gamma \leq \Gamma' =_{df} \forall\ A \in \Gamma\ \exists\ B \in \Gamma'(A \preceq B)$$

McAllester contraction can be characterised in terms of a set of postulates.

THEOREM 2. *Each McAllester contraction satisfies the postulates (K$\dot{-}$1)–(K$\dot{-}$F) below, and conversely, if a contraction operation satisfies the postulates, then it is a McAllester contraction.*

(K$\dot{-}$1) $K \dot{-} A = Cn(K \dot{-} A)$, *where* $Cn(K)$ *denotes the closure of* K *with respect to GMP (closure)*

(K$\dot{-}$2) $K \dot{-} A \subseteq K$ *(inclusion)*

(K$\dot{-}$3) If $A \notin K$, *then* $K \dot{-} A = K$ *(vacuity)*

(K$\dot{-}$4) $A \notin K \dot{-} A$ *(success)*

(K$\dot{-}$6) If $Cn(A) = Cn(B)$, *then* $K \dot{-} A = K \dot{-} B$ *(equivalence)*

(K$\dot{-}$R) For each rule $A_1, \ldots, A_n \to B$, *if* $A_1, \ldots, A_n \to B \in K$, *then* $A_1, \ldots, A_n \to B \in K \dot{-} B$ *(rule persistence)*

(K$\dot{-}$F) If $C \in K$ *and* $C \notin K \dot{-} A$ *then* $C \gg_K A$ *(minimality)*

Proof. First we show that McAllester contractions satisfy the postulates. (K$\dot{-}$1) is satisfied because every time we remove a literal, we also destroy all means of deriving it in K, by removing the weakest premise of every rule which may be used to derive the literal. The resulting smaller set is still closed under GMP application, because no literals which were not originally in K are derivable, and no literals removed from K can be re-derived. (K$\dot{-}$2) is satisfied because we remove literals from K. (K$\dot{-}$3) is satisfied because we only remove A and the literals which A depends on; if A is not in K, it is also not derivable from K, hence there are no such literals. (K$\dot{-}$4) is satisfied because we remove A and destroy all means of deriving A from K (by removing one premise for each rule used in the derivation of A). (K$\dot{-}$6) is trivially true, since in $Cn(A) = Cn(B)$ for two literals A, B with respect to GMP if, and only if, $A = B$ (note that we are concerned with the consequences of a single literal with respect to GMP, not in $Cn(\{A\} \cup K)$). Finally, (K$\dot{-}$R) is satisfied because McAllester contraction only removes literals, and (K$\dot{-}$F) is satisfied because it only removes the literals A depends on.

Now assume we have an operation $\dot{-}$ which satisfies the postulates. We need to show that it removes exactly the set of literals $\{C : C \gg_K A\}$. (K$\dot{-}$2) and (K$\dot{-}$R) guarantee that $K \dot{-} A$ is obtained from K by removing some literals. (K$\dot{-}$F) guarantees that *only* the literals C with $C \gg_K A$ are removed. (K$\dot{-}$1) and (K$\dot{-}$4) guarantee that *all* such C are removed. ∎

Note that it satisfies all but one of the basic AGM postulates. Note also that (K$\dot{-}$3) follows from the other postulates. Namely, by (K$\dot{-}$2), $K \dot{-} A \subseteq K$; by (K$\dot{-}$R), all the rules remain in $K \dot{-} A$; by (K$\dot{-}$F), all the literals C apart from $C \gg_K A$, remain in $K \dot{-} A$; and from $A \notin K$ and (K$\dot{-}$1), we conclude that there are no such C in K.

The reason (K$\dot{-}$5) is not satisfied is simple. Suppose we have a single rule $A \to B$ and K contains $A(c)$ and $B(c)$, and that $A(c) \gg_K B(c)$ holds. After contraction by $B(c)$, both $A(c)$ and $B(c)$ are removed. When we expand by $B(c)$, this becomes the only fact in K, since there is no way to re-derive $A(c)$.

5 Complexity

In this section, we show that the contraction operation above can be implemented to run in linear time.

Assume that the agent's belief state is a directed graph, where the nodes are either beliefs or *justifications*, which correspond to fired rule instances. A justification consists of a belief and a *support list* containing premises of the rule used to derive this belief: for example, $(A, [B, C])$, where A is a derived belief and the rule used to derive it is $B, C \to A$. In the example in section 3, there is a single justification for $\texttt{Region}(\texttt{c}, \texttt{Bordeaux})$, which is

$$(\texttt{Region}(\texttt{c}, \texttt{Bordeaux}), [\texttt{Region}(\texttt{c}, \texttt{Pauillac}), \texttt{PartOf}(\texttt{Pauillac}, \texttt{Bordeaux})])$$

with the support list

$$[\texttt{Region}(\texttt{c}, \texttt{Pauillac}), \texttt{PartOf}(\texttt{Pauillac}, \texttt{Bordeaux})]$$

Foundational beliefs which were not derived, have a justification of the form $(D, [])$, for example $(\texttt{Region}(\texttt{c}, \texttt{Pauillac}), [])$. In the graph, each justification has one outgoing edge to the belief it is a justification for, and an incoming edge from each belief in its support list. We assume that each support list s has a designated *least preferred* member $w(s)$, accessible in constant time.

The following algorithm implements McAllester contraction by A:

```
For each of A's outgoing edges to a justification (C, s),
    remove (C,s) from the graph.

For each of A's incoming edges from a justification (A, s),
    if s is empty:
        remove (A, s);
    else:
        contract by w(s);

Remove A.
```

For *reason-maintenance* type contraction, we also remove beliefs which have no incoming edges in the first step of the algorithm.

The algorithm runs in time $O(kr + n)$, where r is the number of rules, k the maximal number of premises in a rule, and n the number of literals in K. Indeed, the upper bound on the number of steps required to remove justifications corresponding to rule instances is $r(k + 1)$ (one constant time operation for each premise and one for the conclusion of the rule). Removing all justifications corresponding to foundational beliefs costs n steps. The last step in the contraction algorithm (removing a belief) is executed at most n times.

6 Preferences

In this section, we give an example of a preference order on beliefs based on assigning degrees of preference to beliefs. Given a function p assigning numerical preferences to beliefs, we can define $A \preceq B$ as $p(A) \leq p(B)$.

We define preferences using a notion of *quality* associated with justifications. We assume that the quality of a justification is represented by non-negative integers in the range $0, \ldots, m$, where the value of 0 means lowest quality and m means highest quality. We assume that an agent associates an *a priori* quality with each non-inferential justification for its foundational beliefs. For example, a justification $(\texttt{Region}(\texttt{c}, \texttt{Pauillac}), [])$ will have some number assigned to it; so will $(\texttt{Region}(\texttt{c}, \texttt{Tasmania}), [])$. For example, the first justification may have a higher quality since it is part of the original knowledge base, while the second may have lower quality due to being communicated by an unreliable source.

We take the preference of a literal A, $p(A)$, to be that of its highest quality justification:

$$p(A) = max\{qual(j_0), \ldots, qual(j_n)\},$$

where $j_0, \ldots, j_n$ are all the justifications for A, and define the quality of an inferential justification to be that of the least preferred belief in its support:

$$qual(j) = min\{p(A) : A \in \text{ support of } j\}.$$

Literals with no supports (as opposed to an empty support) are viewed as having an empty support of lowest quality. This is similar to ideas in argumentation theory: an argument is only as good as its weakest link, yet a conclusion is at least as good as the best argument for it. This approach is also related to Williams 'partial entrenchment ranking' [Wil95] which assumes that the entrenchment of any sentence is the maximal quality of a set of sentences implying it, where the quality of a set is equal to the minimal entrenchment of its members.

To perform a preferred contraction, we preface the contraction algorithm given in section 5 with a step which computes the preference of each literal in K, and for each justification, finds the position of a least preferred member of its support list. An algorithm for computing preferences is given in [AJL06]; it runs in $O(n \log n + kr)$.

A McAllester contraction based on the order of preference induced by the numerical values of preferences, minimises the preference value of the literals removed as a result of contraction. The following notion makes this precise. Define the *worth* of a set of literals as $worth(\Gamma) = max\{p(A) : A \in \Gamma\}$. We can prove that McAllester contraction removes the set of literals with the least worth:

PROPOSITION 3. *If contraction of the set of literals in K by A resulted in removal of the set of literals Γ, then for any other set of literals Γ' such that $K \backslash \Gamma'$ does not imply A, $worth(\Gamma) \leq worth(\Gamma')$.*

The proof is given in [AJL06].

7 Related work

Our contraction algorithm is similar to the algorithm proposed by McAllester in [McA90] for boolean constraint propagation. McAllester also uses a notion of the 'certainty' of a node, which is similar to our definition of preference. However his reason maintenance system was designed to work with arbitrary boolean formulas, and was not logically complete. As far as we know, the relationship between his belief change operations and AGM contraction has not been investigated.

Our approach to defining the preference order on beliefs is similar to the approach developed in [Dix93, DW93, Wil95] by Williams, Dixon and Wobcke. However, since they work with full classical logic, and calculating entrenchment of a sentence involves considering all possible derivations of this sentence, the complexity of their contraction and revision operations is at least exponential.

Complexity of a Horn clause knowledge base belief revision was studied in [EG96], but with respect to the consequence relation in classical logic. Perhaps the work most similar to ours is that of Bezzazi et al [BJKP98], where belief revision and update operators for forward chaining reasoners were defined and analysed from the point of view of satisfying rationality postulates. The operators are applied to programs, which are finite sets of rules and literals, and are presented as 'syntactic' operators, which do not satisfy the closure under consequence and equivalence postulates. Rather, the authors were interested in preserving the 'minimal change' spirit of revision operators, which resulted in algorithms with high (exponential) complexity. Our approach is different in that we do not consider contraction by rules, and we are concerned less with minimality than with keeping more preferred beliefs.

8 Conclusions

In this paper we have shown how rule-based agents can be modelled as reasoners in a very weak logic with a single inference rule of generalised modus ponens. Their belief sets are deductively closed with respect to this rule; the closure of a finite set of sentences in this logic is still a finite set, thus reducing the distinction between belief bases and theories for rule-based reasoners. Furthermore, we show that it is possible to define a contraction operation for rule-based reasoners which revise only by facts, that satisfies all the basic Alchourrón, Gärdenfors and Makinson (AGM) postulates for contraction (apart from the recovery postulate) and at the same time can be computed in linear time. We prove a representation theorem for this operation with respect to the basic AGM postulates (minus recovery), and two additional postulates.

BIBLIOGRAPHY

[ABH+06] Natasha Alechina, Rafael Bordini, Jomi Hübner, Mark Jago, and Brian Logan. Automating belief revision for AgentSpeak. In Matteo Baldoni and Ulle Endriss, editors, *Declarative Agent Languages and Technologies IV, 4th International Workshop, DALT 2006, Selected, Revised and Invited Papers*, volume 4327 of *Lecture Notes in Artificial Intelligence*, pages 61–77. Springer, 2006.

[AGM85] C. E. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: Partial meet functions for contraction and revision. *Journal of Symbolic Logic*, 50:510–530, 1985.

[AJL06] Natasha Alechina, Mark Jago, and Brian Logan. Resource-bounded belief revision and contraction. In Matteo Baldoni, Ulle Endriss, Andrea Omicini, and Paolo Torroni, editors, *Declarative Agent Languages and Technologies III, Third International Workshop, DALT 2005, Utrecht, The Netherlands, July 25, 2005, Selected and Revised Papers*, volume 3904 of *Lecture Notes in Computer Science*, pages 141–154. Springer, 2006.

[BJKP98] Hassan Bezzazi, Stéphane Janot, Sébastien Konieczny, and Ramón Pino Pérez. Analysing rational properties of change operators based on forward chaining. In Burkhard Freitag, Hendrik Decker, Michael Kifer, and Andrei Voronkov, editors, *Transactions and Change in Logic Databases*, volume 1472 of *Lecture Notes in Computer Science*, pages 317–339. Springer, 1998.

[Dix93] Simon Dixon. A finite base belief revision system. In *Proceedings of Sixteenth Australian Computer Science Conference (ACSC-16): Australian Computer Science Communications*, volume 15, pages 445–451, Brisbane, Australia, January 1993. Queensland University of Technology, Australia.

[Doy77] Jon Doyle. Truth maintenance systems for problem solving. In *Proceedings of the Fifth International Joint Conference on Artificial Intelligence, IJCAI 77*, page 247, 1977.

[Doy92] Jon Doyle. Reason maintenance and belief revision. Foundations vs coherence theories. In P. Gärdenfors, editor, *Belief Revision*, volume 29, pages 29–51. Cambridge University Press, Cambridge, UK, 1992.

[DW93] Simon Dixon and Wayne Wobcke. The implementation of a first-order logic AGM belief revision system. In *Proceedings of 5th IEEE International Conference on Tools with AI*, pages 40–47, Boston, MA, 1993. IEEE Computer Society Press.

[EG96] Thomas Eiter and Georg Gottlob. The complexity of nested counterfactuals and iterated knowledge base revisions. *J. Comput. Syst. Sci.*, 53(3):497–512, 1996.

[FR04] Eduardo Fermé and Hans Rott. Revision by comparison. *Artificial Intelligence*, 157:5 – 47, 2004.

[Han93] Sven Ove Hansson. Theory contraction and base contraction unified. *J. Symb. Log.*, 58(2):602–625, 1993.

[Mak85] D. Makinson. How to give it up: A survey of some formal aspects of the logic of theory change. *Synthese*, 62:347–363, 1985.

[MAR+94] D.L. McGuinness, M.K. Abrahams, L.A. Resnick, P.F. Patel-Schneider, R.H. Thomason, V. Cavalli-Sforza, and C. Conati. Classic knowledge representation system tutorial, 1994. http://www.bell-labs.com/project/classic/papers/ClassTut/ClassTut.html.

[McA90] David A. McAllester. Truth maintenance. In *Proceedings of the Eighth National Conference on Artificial Intelligence (AAAI'90)*, pages 1109–1116. AAAI Press, 1990.

[Neb89] Bernhard Nebel. A knowledge level analysis of belief revision. In R. Brachman, H. J. Levesque, and R. Reiter, editors, *Principles of Knowledge Representation and Reasoning: Proceedings of the First International Conference*, pages 301–311, San Mateo, 1989. Morgan Kaufmann.

[Neb92] Bernhard Nebel. Syntax-based approaches to belief revision. In P. Gärdenfors, editor, *Belief Revision*, volume 29, pages 52–88. Cambridge University Press, Cambridge, UK, 1992.

[Neb94] Bernhard Nebel. Base revision operations and schemes: Representation, semantics and complexity. In Anthony G. Cohn, editor, *Proceedings of the Eleventh European Conference on Artificial Intelligence (ECAI'94)*, pages 341–345, Amsterdam, The Netherlands, August 1994. John Wiley and Sons.

[Rot98] Hans Rott. "Just Because": Taking belief bases seriously. In Samuel R. Buss, Petr Hájaek, and Pavel Pudlák, editors, *Logic Colloquium '98—Proceedings of the 1998 ASL European Summer Meeting*, volume 13 of *Lecture Notes in Logic*, pages 387–408. Association for Symbolic Logic, 1998.

[Wil92] Mary-Anne Williams. Two operators for theory base change. In *Proceedings of the Fifth Australian Joint Conference on Artificial Intelligence*, pages 259–265. World Scientific, 1992.

[Wil95] Mary-Anne Williams. Iterated theory base change: A computational model. In *Proceedings of Fourteenth International Joint Conference on Artificial Intelligence (IJCAI-95)*, pages 1541–1549, San Mateo, 1995. Morgan Kaufmann.