Preference-based belief revision for rule-based agents

Natasha Alechina
School of Computer Science, University of Nottingham, UK
(nza@cs.nott.ac.uk)

Mark Jago
Department of Philosophy, University of Nottingham, UK[*]
(mark.jago@nottingham.ac.uk)

Brian Logan
School of Computer Science, University of Nottingham, UK
(bsl@cs.nott.ac.uk)

Abstract. Agents which perform inferences on the basis of unreliable information need an ability to revise their beliefs if they discover an inconsistency. Such a belief revision algorithm ideally should be rational, should respect any preference ordering over the agent's beliefs (removing less preferred beliefs where possible) and should be fast. However, while standard approaches to rational belief revision for classical reasoners allow preferences to be taken into account, they typically have quite high complexity. In this paper, we consider belief revision for agents which reason in a simpler logic than full first-order logic, namely rule-based reasoners. We show that it is possible to define a contraction operation for rule-based reasoners, which we call McAllester contraction, 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 polynomial time. We prove a representation theorem for McAllester contraction with respect to the basic AGM postulates (minus recovery), and two additional postulates. We then show that our contraction operation removes a set of beliefs which is least preferred, with respect to a natural interpretation of preference. Finally, we show how McAllester contraction can be used to define a revision operation which is also polynomial time, and prove a representation theorem for the revision operation.

1. Introduction

One reason agents interact with other agents or their environment is to acquire new information. In general, it is impossible to ensure that such information will be consistent with the agent's current beliefs, and when an inconsistency is discovered the agent must revise its beliefs to restore consistency. We can identify a number of desiderata for such a revision operation. First it should be theoretically well-motivated, in the sense of producing revisions which conform to a generally accepted set of postulates characterising rational belief revision. Secondly, when choosing among possible revisions, the agent should forgo beliefs which are less preferred, e.g., those which are less certain, credible or useful. Finally, the revision operation itself should be

[*] As of July 2008, Department of Philosophy, Macquarie University, Sydney, Australia

computationally efficient, since further interaction with the environment or other agents is precluded until revision is complete. However, while standard approaches to rational belief revision for classical reasoners allow preferences to be taken into account, they typically have quite high computational complexity, making them unsuitable for use when an agent is interacting with other agents or the environment in real time.

In this paper, we present an approach to rational belief revision which takes the agent's preferences regarding beliefs into account (reflecting e.g., how certain, credible or useful it considers a particular belief to be), and is at the same time efficient. We focus on 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, observations of the agent's environment, downloaded information from various web sources, 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 it is impossible to derive a fact and its negation from a consistent knowledge base using the agent's rules, 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.

This paper extends work reported by Alechina et al. (2006b; 2006a). In (Alechina et al., 2006b), a belief contraction algorithm for rule-based agents was described which runs in time linear in the size of the agent's knowledge base in the propositional case. It was also shown that the operation defined by the algorithm satisfies all but one of the basic Alchourrón, Gärdenfors and Makinson (AGM) postulates. Alechina et al. (2006a) showed how this algorithm can be incorporated in the AgentSpeak agent programming language. In this paper, we give extensional definitions of the contraction and revision operations from (Alechina et al., 2006b) and prove representation theorems for each operation.

The rest of the paper is organised as follows. In section 2, we briefly survey 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 present a logic which characterises the 'deductive abilities' of a forward-chaining rule-based agent. In section 4 we define the 'McAllester contraction' operation, and prove a representation theorem for it. In section 5, we show that McAllester contraction can be computed in time polynomial 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 McAllester contraction removes a set of literals which is least preferred. In section 7 we show how to modify the definition of contraction so that it removes a minimal set of least preferred literals. In section 8 we show how McAllester contraction can be used to define a revision operation which is also polynomial time, and prove a representation theorem for the revision operation. We briefly survey related work in section 9, and conclude.

2. Approaches to 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 (Alchourrón et al., 1985) and reason-maintenance style belief revision (Doyle, 1977).

Classical AGM-style belief revision describes an idealised reasoner with a potentially infinite set of beliefs closed under logical consequence. Revision is based on the ideas of coherence and informational economy, i.e., 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. When new information becomes available, a reasoner must modify its belief set to incorporate it. The AGM theory defines three operators on a belief set K: 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 and closing under consequence. 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 (Alchourrón et al., 1985). 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 basic postulates for revision are:

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

(K$\dot{+}$2) $A \in K \dot{+} A$

(K$\dot{+}$3) $K \dot{+} A \subseteq K + A$

(K$\dot{+}$4) If $\{A\} \cup K$ is consistent, then $K + A = K \dot{+} A$[1]

(K$\dot{+}$5) $K \dot{+} A$ is inconsistent if, and only if, A is inconsistent.

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

If the agent is a reasoner in classical logic, a revision operator $\dot{+}$ can be defined in terms of contraction: $K \dot{+} A \stackrel{df}{=} (K \dot{-} \neg A) + A$ (this is known as the Levi identity).

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 A must be given up, sufficient foundational beliefs have to be withdrawn to render A underivable. Moreover, if all the justifications for A are withdrawn, then A should no longer be held. A more detailed comparison of the two approaches can be found in, for example (Doyle, 1992). More recently, a number of approaches have been proposed which try to combine elements of AGM-style and reason maintenance-style revision; see for example cautious revision by Tennant (Tennant, 2006) and work by Dixon and Foo (Dixon and Foo, 1993).

Most implementations of reason-maintenance style belief revision are logically incomplete, but tractable. For example, McAllester's boolean constraint propagation algorithm (McAllester, 1990) does not find all the classical logical consequences of boolean formulas, sacrificing completeness for efficiency. The AGM theory, on the other hand, 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 within the AGM approach, an approach called belief base revision has been proposed by Makinson (1985), Nebel (1989), Williams (1992), Hansson (1993) and Rott (1998), amongst others. 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

[1] We replaced '$\neg A \notin K$' with '$\{A\} \cup K$ is consistent' here, since the two formulations are classically equivalent.

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) (Nebel, 1994). The reason for the high complexity is the need to check for classical consistency while performing the operations. Similarly, dependency-network contraction was shown to be NP-complete in (Tennant, 2003).

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 by Nebel (1992)). This is essentially what we do 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

$$\texttt{PartOf}(\texttt{Bordeaux}, \texttt{France})$$

and

$$\texttt{PartOf}(\texttt{x}, \texttt{Bordeaux})$$

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

$$A_1, \ldots, A_n \rightarrow B$$

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. We assume variables are universally quantified, and that rules do not contain functional symbols. An example of a rule is:[2]

$$\texttt{Region}(\texttt{x}, \texttt{y}), \texttt{PartOf}(\texttt{y}, \texttt{z}) \rightarrow \texttt{Region}(\texttt{x}, \texttt{z}).$$

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

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

[2] The rules are adapted from McGuinness et al. (1994)'s wine ontology.

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, $\mathtt{Pauillac}$ to y, and $\mathtt{Bordeaux}$ to z, then

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

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

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.

The agent's beliefs are closed under logical consequence in a logic W which has a single inference rule, generalised modus ponens (GMP):

$$\frac{\delta(A_1), \ldots, \delta(A_n), \quad A_1, \ldots, A_n \rightarrow B}{\delta(B)}$$

where δ is a substitution function which replaces all the free variables of $A_1, \ldots, A_n \rightarrow B$ with constants. We will use $\vdash$ to denote derivability in W and Cn closure under consequence in W. Note that W is much weaker than classical logic. 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 the law of excluded middle). Also, B and $\neg B$ do not entail arbitrary formulas.

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

R1 $\mathtt{Region(x,y)}, \mathtt{PartOf(y,z)} \rightarrow \mathtt{Region(x,z)}$

R2 $\mathtt{Region(x,France)} \rightarrow \neg\mathtt{Region(x,Australia)}$

and facts:

F1 $\mathtt{Region(c,Pauillac)}$

F2 $\mathtt{PartOf(Pauillac,Bordeaux)}$

F3 $\mathtt{PartOf(Bordeaux,France)}$

F4 $\mathtt{PartOf(Tasmania,Australia)}$

From these rules and facts, the agent can derive

F5 $\mathtt{Region(c,Bordeaux)}$ (from F1, F2, R1)

F6 $\mathtt{Region(c,France)}$ (from F5, F3, R1)

F7 $\neg\mathtt{Region(c,Australia)}$ (from F6, R2)

To illustrate the need for such an agent to revise its beliefs, assume further that the agent is told that Chateau Lafite Rothschild Pauillac is a region of Tasmania:

F8 $\mathtt{Region}(\mathtt{c}, \mathtt{Tasmania})$

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

F9 $\mathtt{Region}(\mathtt{c}, \mathtt{Australia})$ (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 $\mathtt{Region}(\mathtt{c}, \mathtt{Australia})$ or by $\neg\mathtt{Region}(\mathtt{c}, \mathtt{Australia})$.

4. McAllester contraction

In this section, we define a contraction operation for rule-based reasoners. In defining our contraction operation we have chosen to allow contraction only by literals, and not by rules. For many rule-based agents it is reasonable to suppose that the agent's rules are 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, perhaps unreliable sources, and as such are a possible source of inconsistencies. To model sources of belief, we assume that the agent prefers some beliefs to others. For example, it may trust communicated information (such as **F8** in the example above) less than the information in its original knowledge base (such as **F1**). Those preferences are used to decide which beliefs to remove to restore consistency. To make the notion of preference precise, we assume that there is a preference order $\preceq$, which is a total order on the set of ground literals.[3] 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, for example, a lexicographical order to decide which one precedes the other in $\preceq$. For any finite set of ground literals Γ, we denote by $w(\Gamma)$ ('the weakest element of Γ') the element of Γ which is minimal with respect to $\preceq$.

[3] A total or linear order $\preceq$ on a set X is a binary relation satisfying, for all $A, B, C \in X$:

(reflexivity) $A \preceq A$

(antisymmetry) if $A \preceq B$ and $B \preceq A$, then $A = B$

(transitivity) if $A \preceq B$ and $B \preceq C$, then $A \preceq C$

(comparability) $A \preceq B$ or $B \preceq A$.

Let us denote the set of the 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. $A_1, \ldots, A_n \rightarrow B \in K$, $\delta(A_1), \ldots, \delta(A_n) \in K$, 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. $A_1, \ldots, A_n \rightarrow C \in K, \delta(A_1), \ldots, \delta(A_n) \in K$, $\delta(A) = w(\delta(A_1), \ldots, \delta(A_n))$, and $\delta(C) \gg_K \delta(B)$.

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.[4]

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

$$K \dot{-} A \stackrel{df}{=} Cn(K \backslash \Gamma)$$

where $\Gamma \subseteq \{C : C \gg_K A\}$ *and* $K \backslash \Gamma \not\vdash A$.

The motivation behind this definition of contraction is simple: to contract by A, we need to 'destroy', by removing some premise, each rule instance which can be used to derive A, and we choose to remove those beliefs which are least preferred. [5] Note that since some rule instances may share premises, 'destroying' one rule instance may mean that another rule instance becomes destroyed, too, and it is not always necessary to remove all of $\{C : C \gg_K A\}$ to make A underivable.

McAllester contractions can be characterised by the following set of postulates.

[4] It is also different from the notion of dependence in logic programming (Kowalski, 1979) where the head of a rule depends on any literal in the body (not just the weakest).

[5] This is similar in spirit to safe contraction (Alchourrón and Makinson, 1985) and kernel contraction (Hansson, 1994) but does not involve considering all the kernels/inclusion-minimal subsets of K which imply A. It can also be shown to remove, in some cases, a different set of literals than any kernel contraction would remove. The following example demonstrates this. Let $K = \{A(0), A(1), A(2), r\}$, where $r = A(x), A(0) \rightarrow A(y)$, and the preference order is $A(2) \preceq A(1) \preceq A(0)$. Suppose we contract by $A(2)$. There are three rule instances of r which derive $A(2)$, obtained by substitutions $(x/0, y/2)$, $(x/1, y/2)$ and $(x/2, y/2)$. McAllester contraction may remove the weakest premise from each, namely $A(0)$, $A(1)$ and $A(2)$. However there are only two $A(2)$-kernels, namely $\{A(2)\}$ and $\{A(0), r\}$, so there is no function which removes one element from every kernel which would remove three literals. Hence there is no kernel contraction corresponding to this particular McAllester contraction.

THEOREM 1. *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 $A_1, \ldots, A_n \rightarrow B$*, if* $A_1, \ldots, A_n \rightarrow B \in K$*, then* $A_1, \ldots, A_n \rightarrow B \in K \dot{-} B$ *(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 we explicitly deductively close $K\backslash\Gamma$. (K$\dot{-}$2) is satisfied because we never add literals to K. (K$\dot{-}$3) is satisfied because we remove only A and the literals on which A depends; if A is not in K, it is also not derivable from K, hence there are no such literals. (K$\dot{-}$4) is satisfied because we remove A and destroy all means of deriving A from K. (K$\dot{-}$6) is trivially true, since $Cn(A) = Cn(B)$ for 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 removes only literals, and (K$\dot{-}$F) is satisfied because it removes only the literals on which A depends.

Now assume we have an operation $\dot{-}$ which satisfies the postulates. We need to show that it removes a subset Γ of $\{C : C \gg_K A\}$ such that $K\backslash\Gamma$ does not derive A, and is deductively closed. (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* those literals C such that $C \gg_K A$ are removed. (K$\dot{-}$1) and (K$\dot{-}$4) guarantee that enough of these literals are removed to make A underivable. Finally, (K$\dot{-}$1) guarantees that the set is closed under consequence. □

Note that, for a McAllester contraction, (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 those for which $C \gg_K A$ holds, remain in $K \dot{-} A$; and from $A \notin K$ and (K$\dot{-}$1), we conclude that there are no such C in K. (K$\dot{-}$6) is also not required for the representation theorem and is included to allow comparison with the classic AGM postulates.

McAllester contraction does not satisfy the recovery postulate (K$\dot{-}$5). The reason (K$\dot{-}$5) is not satisfied is simple. Suppose we have a single rule $A(x) \rightarrow B(x)$, K contains $A(c)$ and $B(c)$, and that $A(c) \gg_K B(c)$. 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 $A(c)$. However, there are significant objections in belief revision literature to the recovery postulate (see for example Tennant (2006)) and there have been a number of proposals for rational belief revision on finite bases which do not satisfy the recovery postulate (see e.g., (Makinson, 1987; Hansson, 1991)).

5. Complexity

In this section, we show that McAllester contraction can be implemented so as to run in time polynomial in the size of its belief set (the set of literals in working memory and the set of the agent's rules).

Assume that the agent maintains a directed graph of beliefs and *justifications* for beliefs, corresponding to fired rule instances. Each justification consists of a belief and a *support list* containing premises of the rule used to derive this belief: for example, the justification of a belief $B(c)$ derived using the rule $A_1(x), \ldots, A_n(x) \rightarrow B(x)$ from $A_1(c), \ldots, A_n(c)$ is $(B(c), [A_1(c), \ldots, A_n(c)])$. 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 (non-derived) beliefs have a justification with an empty support list, for example, $(\texttt{Region}(\texttt{c}, \texttt{Pauillac}), [\ ])$. In the graph, each justification has one outgoing edge to the belief it justifies, and an incoming edge from each belief in its support list. We assume that each support list s has a designated least preferred member $w(s)$, which is accessible in constant time.

Algorithm 1 implements McAllester contraction by A. The algorithm consists of two main loops. The first loop removes all justifications which have an incoming edge from A. The second loop iterates through all justifications for A, and for each justification, either removes it (if it has an empty support), or recurses to contract by the weakest member of the justification's support list, $w(s)$. Note that the algorithm computes a deductively closed set by ensuring that justifications for each of the removed literals are destroyed. The algorithm runs in time $O(kr + n)$, where r is the number of rule instances, k the maximal number of premises in a rule, and n the number of literals in

Algorithm 1 McAllester contraction by A

```
for all j = (B, s) with an edge from A do
    remove j (and all edges to and from j) from the graph
end for
for all j = (A, s) with an edge to A do
    if s == [] then
        remove j (and the edge from j to A)
    else
        contract by w(s)
    end if
end for
remove A
```

K. 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 instance). The last step in the contraction algorithm (removing a belief) is executed at most n times. Note that in the propositional case, when the number of rules equals the number of rule instances, the algorithm runs in time linear in the size of the agent's belief set (Alechina et al., 2006b).

It is straightforward to modify algorithm 1 to perform reason-maintenance style contraction. Reason-maintenance contraction by A involves removing not only those justifications whose supports contain A, but also all beliefs which have no justifications whose supports do not contain A. In this case, in addition to removing the justifications where A is in the support list, we check to see if this leaves some literal with no incoming edges (no justifications for it). If so, we remove the literal and recurse forwards, following edges to the justifications in whose support list it appears. This adds another traversal of the graph, but the overall complexity remains $O(kr + n)$.

6. Preferences

In this section, we give an example of a preference order on beliefs based on assigning degrees of preference to beliefs. Recall from section 4 that, given a function p assigning numerical preferences to beliefs, we can define $A \preceq B$ as $p(A) < p(B)$; if $p(A) = p(B)$ for two distinct literals A and B, we use, for example, lexicographic order to decide whether $A \preceq B$ or vice versa.

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. In our example, the justification $(\texttt{Region}(\texttt{c}, \texttt{Pauillac}), [\,])$ will have some number assigned to it; so will $(\texttt{Region}(\texttt{c}, \texttt{Tasmania}), [\,])$. The first justification may have a higher quality since it is part of the original knowledge base, while the second may have lower quality if it has been 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. It is also related to Williams's *partial entrenchment ranking* (1995), 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. While this approach is intuitively appealing, nothing hangs on it and Theorem 1 holds for any total order over literals. For example, the preference of a derived literal could be a property of the rule used to derive it or given by some function of the preferences of its antecedents.

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 the least preferred member of its support list. An algorithm for computing preferences which runs in $O(n \log n + kr)$ is given in (Alechina et al., 2006b).

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

Proof. If $A \notin K$, the statement is immediate since $\Gamma = \varnothing$. Assume that $A \in K$. In this case, $A \in \Gamma$ and $A \in \Gamma'$ (otherwise $K \setminus \Gamma$ and $K \setminus \Gamma'$ would still derive A). It is also easy to see that A is the maximal element of

Γ, because a literal B is in Γ if and only if either (1) $p(B) = qual(j_i)$ for some justification j_i for A, and since $p(A) = \max(qual(j_0), ..., qual(j_n))$, $p(B) \leq p(A)$; or (2) B is the least preferred element of a support set for some literal A depends on, in which case its preference is less or equal to the preference of the literal it is justification for, which in turn is less or equal to $p(A)$. So, since A is an element of both Γ and Γ', and A has the maximal preference in Γ, then $worth(\Gamma) \leq worth(\Gamma')$. □

7. Minimal change and least worth

The set of literals removed as a result of McAllester contraction of K by A is not necessarily a minimal set of literals which has to be removed from K to make A underivable. Consider the following example:

$$K = \{(B, C \rightarrow A), \quad (C, D \rightarrow A), \quad A, B, C, D\}$$

and assume $B \preceq C \preceq D$. Contraction by A may result in removing A, B and C, but it would have been sufficient to remove just A and C.

Proposition 1 tells us however that given the preference order defined in Section 6, the set of literals removed as a result of McAllester contraction is minimal in the following order on subsets of K:[6]

$$\Gamma \leq \Gamma' \stackrel{df}{=} \forall\, A \in \Gamma\, \exists\, B \in \Gamma' \;\; (p(A) \leq p(B))$$

This means that the Γ removed by McAllester contraction by A may contain some less preferred belief B which did not need to be removed to prevent the rederivation of A, i.e., $K \setminus (\Gamma \setminus B)$ would not have derived A. However, any other set of literals Γ' which has to be removed in order to make A underivable will have to contain at least one belief which is at least as preferred as the most preferred belief removed by the McAllester contraction.

It is possible to extend the algorithm for McAllester contraction so that the set of literals it removes both has least worth and is minimal, in the sense that no proper subset of this set would be sufficient for contraction. The complexity of the algorithm remains polynomial.

DEFINITION 2. *The* minimal *McAllester contraction of K by a literal A, $K \dot{-}_m A$, is defined as*

$$K \dot{-}_m A \stackrel{df}{=} Cn(K \setminus \Gamma)$$

where Γ is a subset of $\{C : C \gg_K A\}$ which is minimal in the sense that for every $\Gamma' \subset \Gamma$, $K \setminus \Gamma' \vdash A$.

[6] This definition of the order on subsets is similar to Rott's $\lhd^-$ (1998).

Algorithm 2 describes a polynomial-time procedure to compute the minimal McAllester contraction by A.[7]

Algorithm 2 Minimal McAllester contraction by A

$K \dot{-}_m A = K \dot{-} A$
$\Gamma = K \setminus (K \dot{-} A)$
for all C in Γ **do**
  if $A \notin Cn((K \dot{-}_m A) \cup \{C\})$ **then**
    $K \dot{-}_m A = (K \dot{-}_m A) + C$
  end if
end for

A representation theorem for a minimal McAllester contraction is very similar to Theorem 1, except that the minimality postulate (K$\dot{-}$F) in Theorem 1 must be replaced by the following *strong minimality postulate*:

(K$\dot{-}$F-strong) If $C \in K$ and $C \notin K \dot{-}_m A$ then $C \gg_K A$ and $(K \dot{-}_m A) \cup C \vdash A$ (strong minimality)

8. Revision

In the previous sections we described contraction. Now let us consider revision, which is adding a new belief in a manner that does not result in an inconsistent set of beliefs. Recall that a set of beliefs K is inconsistent if for some literal A, both A and $\neg A$ are in K.

As noted in section 2, if the agent is a reasoner in classical logic, revision is definable in terms of contraction and vice versa. Given a contraction operator $\dot{-}$ which satisfies postulates (K$\dot{-}$1)–(K$\dot{-}$4) and (K$\dot{-}$6), a revision operator $\dot{+}$ which satisfies (K$\dot{+}$1)–(K$\dot{+}$6) can be defined via the Levi identity as

$$K \dot{+} A \stackrel{df}{=} (K \dot{-} \neg A) + A.$$

Conversely, if a revision operator satisfies (K$\dot{+}$1)–(K$\dot{+}$6), then contraction defined via the Harper identity as

$$K \dot{-} A \stackrel{df}{=} (K \dot{+} \neg A) \cap K$$

satisfies postulates (K$\dot{-}$1)–(K$\dot{-}$6): see Gärdenfors (1988).

However, for an agent which is not a classical reasoner, contraction and revision are not necessarily inter-definable in this way. In particular, they are

[7] To simplify the exposition we give a naive algorithm here; more efficient algorithms are possible.

not inter-definable when the consistency of $K + A$ is not equivalent to $K \nvdash \neg A$. For a rule-based agent which reasons in the logic W, applying the Levi identity to McAllester contraction results in a revision operation which does not satisfy (K$\dot{+}$5). One of the reasons for this is that contracting the agent's belief set by $\neg A$ does not make this set consistent with A, so $(K \dot{-} \neg A) + A$ may be inconsistent. Instead, we define revision by A as $(K + A) \dot{-} \bot$, that is, as expansion by A followed by elimination of all contradictions.[8]

Algorithm 3 Revision by A

Add A to K and run rules to quiescence
Let $L = \{(B_1, \neg B_1), \ldots, (B_n, \neg B_n)\}$ be the list of all contradictions in K, ordered by preference order on $w(\{B_i, \neg B_i\})$
for $(B_i, \neg B_i)$ in L **do**
  contract by $w(\{B_i, \neg B_i\})$
end for

Algorithm 3 implements this revision operation. Note that we need to specify the order in which the contradictions are eliminated, for McAllister contraction is order-dependent. That is:

$$(K \dot{-} B_1) \dot{-} B_2 \neq (K \dot{-} B_2) \dot{-} B_1.$$

To see why this is so, consider the following example. Let:

$$K = \{A_1, A_2 \rightarrow B_1,\ \ A_2 \rightarrow B_2,\ \ A_1,\ A_2,\ B_1,\ B_2\}$$

and $w(\{A_1, A_2\}) = A_1$. Then $(K \dot{-} B_1) = \{A_2, B_2\}$ and $(K \dot{-} B_1) \dot{-} B_2 = \varnothing$. On the other hand, $(K \dot{-} B_2) = \{A_1, B_1\}$ and $(K \dot{-} B_2) \dot{-} B_1 = \{A_1\}$.

The declarative definition of the revision operation computed by algorithm 3 is as follows. Let $(B_1, \neg B_1), \ldots, (B_n, \neg B_n)$ be the list of all contradictions in $K + A = Cn(K \cup \{A\})$, ordered by preference order on $w(B_i, \neg B_i)$, and let $\sim B_i = w(\{B_i, \neg B_i\})$. Then

$$K \dot{+} A \stackrel{df}{=} (K + A) \dot{-} \sim B_1 \dot{-} \sim B_2 \dot{-} \ldots \dot{-} \sim B_n.$$

Call this revision *ordered contraction by contradictions* (OCC).

THEOREM 2. *Each OCC revision satisfies the postulates (K$\dot{+}$1)–(K$\dot{+}$OCC) below and, conversely, if a revision operation satisfies the postulates, then it is an OCC revision.*

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

(K$\dot{+}$3) $K \dot{+} A \subseteq K + A$

[8] This is called semi-revision by Hansson (1997).

(K$\dot{+}$4) If $\{A\} \cup K$ *is consistent, then* $K + A = K \dot{+} A$

(K$\dot{+}$5) $K \dot{+} A$ *is inconsistent if, and only if,* A *is inconsistent.*

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

(K$\dot{+}$R) For each rule $A_1, \ldots, A_n \rightarrow B$*,* $A_1, \ldots, A_n \rightarrow B \in K$ *iff* $A_1, \ldots, A_n \rightarrow B \in K \dot{+} B$

(K$\dot{+}$OCC) If $C \in K + A$ *and* $C \notin K \dot{+} A$*, then for some* i*,*

$$C \gg_{(K+A)\dot{-}\sim B_1 \dot{-} \ldots \dot{-} \sim B_{i-1}} \sim B_i$$

where $(B_1, \neg B_1), \ldots, (B_n, \neg B_n)$ *are all the contradictions in* $K + A$*,* $\sim B_i = w(B_i, \neg B_i)$*,* $\sim B_1 \preceq \ldots \preceq \sim B_n$*, and* $i \in \{1, \ldots, n\}$*.*

Proof. (K$\dot{+}$1) is satisfied because after we add A, we run the rules to quiescence. (K$\dot{+}$3) is satisfied because the construction of $K \dot{+} A$ starts with A being added to K, which is then closed under consequence (which is $K + A$), and after that literals can only be removed from K. (K$\dot{+}$4) holds because, if adding A does not cause an inconsistency, then $K \dot{+} A = K + A$ by the definition of OCC revision. (K$\dot{+}$5) holds trivially because A and $K \dot{+} A$ are never inconsistent. Finally, recall that in the agent's logic, $Cn(A) = Cn(B)$ only if $A = B$, so (K$\dot{+}$6) holds trivially. (K$\dot{+}$R) holds because the set of the agent's rules does not change as a result of closing under consequence (since only literals are derivable by GMP) or contracting by contradictions (due to (K$\dot{-}$R)). (K$\dot{+}$OCC) holds because the only beliefs removed from $(K + A) \dot{-} \sim B_1 \dot{-} \ldots \dot{-} \sim B_{i-1}$ when contracting by $\sim B_i$ are the weakest premises of rules which can be used to derive $\sim B_i$ in $(K + A) \dot{-} \sim B_1 \dot{-} \ldots \dot{-} \sim B_{i-1}$.

Now assume that $\dot{+}$ is an operation satisfying the postulates above. By (K$\dot{+}$3), we know that $K \dot{+} A = (K + A) \setminus \Gamma$ for some (possibly empty) set Γ. By (K$\dot{+}$R), Γ is a set of literals. If $K + A$ is consistent, by (K$\dot{+}$4), $K \dot{+} A = K + A$, in other words, it is the result of contracting $K + A$ by an empty set of literals, so it is an OCC revision.

Suppose $K + A$ is not consistent, namely it contains contradictions

$$(B_1, \neg B_1), \ldots, (B_n, \neg B_n)$$

for some $n \geq 1$. As before, let $w(B_i, \neg B_i)$ be denoted by $\sim B_i$, and assume that $\sim B_1 \preceq \ldots \preceq \sim B_n$. We need to prove that $K \dot{+} A$ is really the result of a series of McAllester contractions by $\sim B_1, \ldots, \sim B_n$. By (K$\dot{+}$6), we know that a literal C is in Γ (is removed from $K + A$ as a result of the revision operation) only if there exists a McAllester contraction $\dot{-}$ such that for some i, $C \in (K + A) \dot{-} \sim B_1 \dot{-} \ldots \dot{-} \sim B_{i-1}$ and $C \gg_{(K+A)\dot{-}\sim B_1 \dot{-} \ldots \dot{-} \sim B_{i-1}} \sim B_i$. This makes it is easy to prove by induction on n that we do indeed have a series of contraction operations applied to $K + A$. Namely, we partition

Γ into sets of literals Γ_i removed at stage i (if different contractions would remove some literal at different stages, we choose the latest stage to make sure that each successive set is deductively closed):

$\Gamma_1 \subseteq \{C : C \gg_{K+A} {\sim}B_1\}$,

$\Gamma_i \subseteq \{C : C \gg_{(K+A)\dot{-}{\sim}B_1\dot{-}\ldots\dot{-}{\sim}B_{i-1}} {\sim}B_i\}$.

If $n = 1$, then by (K$\dot{+}$1) and (K$\dot{+}$6), $K \dot{+} A = Cn(K \setminus \Gamma_1))$ in other words, $(K+A) \dot{-} {\sim}B_1$. If by the inductive hypothesis, $(K+A)\setminus(\Gamma_1 \cup \ldots \cup \Gamma_{n-1}) = (K + A) \dot{-} {\sim}B_1 \dot{-} \ldots \dot{-} {\sim}B_{n-1}$, then removing Γ_n will give a McAllester contraction by ${\sim}B_n$:

$$K \dot{+} A = (K + A) \setminus \Gamma = (K + A) \dot{-} {\sim}B_1 \dot{-} \ldots \dot{-} {\sim}B_n.$$

Hence, any operation satisfying (K$\dot{+}$3), (K$\dot{+}$4), (K$\dot{+}$R) and (K$\dot{+}$OCC) is an OCC revision. $\square$

Note that (K$\dot{+}$2), $A \in K \dot{+} A$, does not hold. Suppose we add A to K and derive some literal B, but $\neg B$ is already in K and has a higher preference value than B. Then we contract by B, which may well result in contraction by A. Another example of a situation in which (K$\dot{+}$2) is violated is revision by A in the presence of a rule $A \rightarrow \neg A$.

One could question whether (K$\dot{+}$2) is a desirable property. For example, Galliers (1992) has argued that it would not be satisfied by an agent which performs autonomous belief revision. However, if we do want to define a belief revision operation which satisfies (K$\dot{+}$2), we need to ensure that: (1) A is preferred to all other facts in working memory; and (2) A on its own cannot be responsible for an inconsistency. One way to satisfy the first requirement is to use a preference order based on timestamps: more recent information is more preferred. To satisfy the second requirement, we may postulate that the agent's rules are not *perverse*. We call a set of rules $\mathcal{R}$ perverse if there is a literal A such that running $\mathcal{R}$ to quiescence on $K = \{A\} \cup \mathcal{R}$ results in deriving a contradiction $\{B, \neg B\}$ (including the possibility of deriving $\neg A$). This is equivalent to saying that no singleton set of literals is *exceptional* in the sense of Bezzazi et al. (1998).

Analogously to the minimal McAllester contraction defined in section 7 we can define a maximal revision by A operation, $K \dot{+}_m A$, which retains a maximal set of literals consistent with A. To do this we replace the McAllester contraction in the definition of OCC by the minimal McAllester contraction.

9. Related work

The problem of characterising rational belief change has been intensively investigated, and there is a substantial literature. In this section, we give a brief overview of the work most closely related to our own, focusing specifically on resource-bounded revision and contraction.

Our contraction algorithm is inspired by the algorithm proposed by McAllester (1980) for eliminating inconsistencies resulting from boolean constraint propagation. He also used a notion of the *certainty* of a node, which is similar to our definition of preference. McAllester's algorithm was implicitly resource-bounded in that it was intended to form part of a practical AI system. However his reason maintenance system was designed to work with arbitrary boolean formulas, and was not logically complete. To the best of our knowledge, the relationship between the retraction, backtracking and refutation operations in (McAllester, 1980) and AGM contraction and revision has not been investigated.

The use of preference or epistemic entrenchment orderings to define revision and contraction was proposed by Gärdenfors (1988) and Gärdenfors and Makinson (1988), although the properties of entrenchment and the way in which entrenchment determines the result of contraction differs from that proposed in this paper. An anytime algorithm for computing a series of approximations of a new entrenchment order corresponding to a changed degree of belief in a given sentence (and hence the resulting new belief set) was proposed by Williams (1997); however even one iteration of this algorithm involves computing all minimal sets of sentences entailing the given sentence, and hence has a higher complexity then our revision algorithm.

Our approach to defining the preference order on beliefs is similar to the approach developed by Dixon (1993), Dixon and Wobcke (1993) and Williams (1995). Again, 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.

Wasserman (2001) has proposed an approach to resource-bounded belief revision which has similar motivations to our work. Her approach is based on a compartmentalised belief base, defined by a prior notion of relevance between beliefs. Given a belief and a degree of relevance, there is a compartment of the belief base which contains the beliefs relevant to that degree, to the specified belief. A revision operation takes a formula and a degree of relevance and revises the beliefs in the corresponding compartment of the agent's belief set. The revision procedure itself is classical belief base revision; however, as the compartments of the belief base are typically much smaller than the entire belief base, revision of a compartment typically requires less computation than a revision of the entire base.

Wasserman's approach does not ensure that the belief set as a whole is consistent after revision; it ensures only that the compartment in question is consistent. The idea is that the agent is mainly interested in ensuring consistency between the beliefs that are relevant (to some degree) to the present task, rather than ensuring the consistency of its belief set as a whole. A similar approach is taken by Chopra et al. (2001). They define a contraction operation which approximates a classical AGM contraction operation. Their contraction operation has complexity $O(|K| \cdot |A| \cdot 2^S)$, where K is the knowledge base, A the formula to be contracted, and S is a set of 'relevant' atoms. As S gets larger, their contraction operation becomes a closer approximation of classical contraction. Our approach differs both from Wasserman's and from Chopra et al.'s in that it can be applied to ensure the consistency of large rule-based belief sets as a whole.

The complexity of Horn clause knowledge base belief revision was studied by Eiter and Gottlob (1996) but with respect to the consequence relation in classical logic. Perhaps the work most similar to ours is that of Bezzazi et al. (1998). They also consider forward-chaining agents whose program consists of a set of literals and a set of rules. They propose several revision operators, some of which revise both the rules and the literals, and some of which revise only the literals, and consider which of the rationality postulates for belief revision or update (proposed by Katsuno and Mendelzon (1991)) they satisfy. However, Bezzazi et al. are mostly concerned with belief revision operations which are *minimal*, which results in algorithms with high (exponential) complexity. The only polynomial-time operation they consider, ranked revision, makes sense only in the setting of default rules; for programs where rules do not have exceptions, ranked revision of a program P by a program P' is $P + P'$ if P and P' are consistent and P' otherwise.

10. Conclusions

In this paper we have shown how rule-based agents can be modelled as reasoners in a 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 difference between belief bases and theories for rule-based reasoners. We defined a contraction operation, McAllester contraction, for rule-based reasoners which revise only by facts, and showed that it 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 polynomial time. We proved a representation theorem for McAllester contraction with respect to the basic AGM postulates (minus recovery), and two additional postulates. We also showed how our contraction operation can be used to define a corresponding revision operation which is also polynomial time, and proved a representation theorem for the revision operation.

References

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

Alchourrón, C. E. and D. Makinson: 1985, 'The logic of theory change: Safe contraction'. *Studia Logica* **44**, 405–422.

Alechina, N., R. Bordini, J. Hübner, M. Jago, and B. Logan: 2006a, 'Automating Belief Revision for AgentSpeak'. In: M. Baldoni and U. Endriss (eds.): *Declarative Agent Languages and Technologies IV, 4th International Workshop, DALT 2006, Selected, Revised and Invited Papers*, Vol. 4327 of *Lecture Notes in Artificial Intelligence*. pp. 61–77.

Alechina, N., M. Jago, and B. Logan: 2006b, 'Resource-Bounded Belief Revision and Contraction.'. In: M. Baldoni, U. Endriss, A. Omicini, and P. Torroni (eds.): *Declarative Agent Languages and Technologies III, Third International Workshop, DALT 2005, Utrecht, The Netherlands, July 25, 2005, Selected and Revised Papers*, Vol. 3904 of *Lecture Notes in Computer Science*. pp. 141–154.

Bezzazi, H., S. Janot, S. Konieczny, and R. P. Pérez: 1998, 'Analysing Rational Properties of Change Operators Based on Forward Chaining'. In: B. Freitag, H. Decker, M. Kifer, and A. Voronkov (eds.): *Transactions and Change in Logic Databases*, Vol. 1472 of *Lecture Notes in Computer Science*. pp. 317–339.

Chopra, S., R. Parikh, and R. Wassermann: 2001, 'Approximate Belief Revision'. *Logic Journal of the IGPL* **9**(6), 755–768.

Dixon, S.: 1993, 'A Finite Base Belief Revision System'. In: *Proceedings of Sixteenth Australian Computer Science Conference (ACSC-16): Australian Computer Science Communications*, Vol. 15. Brisbane, Australia, pp. 445–451.

Dixon, S. and N. Foo: 1993, 'Connections between the ATMS and AGM Belief Revision'. In: R. Bajcsy (ed.): *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*. San Mateo, California, pp. 534–539.

Dixon, S. and W. Wobcke: 1993, 'The Implementation of a First-Order Logic AGM Belief Revision System'. In: *Proceedings of 5th IEEE International Conference on Tools with AI*. Boston, MA, pp. 40–47.

Doyle, J.: 1977, 'Truth Maintenance Systems for Problem Solving'. In: *Proceedings of the Fifth International Joint Conference on Artificial Intelligence, IJCAI 77*. p. 247.

Doyle, J.: 1992, 'Reason maintenance and belief revision. Foundations vs coherence theories'. In: P. Gärdenfors (ed.): *Belief Revision*, Vol. 29. Cambridge, UK: Cambridge University Press, pp. 29–51.

Eiter, T. and G. Gottlob: 1996, 'The Complexity of Nested Counterfactuals and Iterated Knowledge Base Revisions.'. *J. Comput. Syst. Sci.* **53**(3), 497–512.

Galliers, J. R.: 1992, 'Autonomous Belief Revision and Communication'. In: P. Gärdenfors (ed.): *Belief Revision*, Cambridge Tracts in Theoretical Computer Science 29. Cambridge University Press, pp. 220–246.

Gärdenfors, P.: 1988, *Knowledge in Flux: Modelling the Dynamics of Epistemic States*. Cambridge, Mass.: The MIT Press.

Gärdenfors, P. and D. Makinson: 1988, 'Revisions of knowledge systems using epistemic entrenchment'. In: M. Y. Vardi (ed.): *Proceedings of the Second Conference on Theretical Aspects of Reasoning About Knowledge*. pp. 83–95.

Hansson, S. O.: 1991, 'Belief Contraction Without Recovery'. *Studia Logica* **50**(2), 251–260.

Hansson, S. O.: 1993, 'Theory Contraction and Base Contraction Unified.'. *J. Symb. Log.* **58**(2), 602–625.

Hansson, S. O.: 1994, 'Kernel Contraction'. *Journal of Symbolic Logic* **59**, 845–859.

Hansson, S. O.: 1997, 'Semi-revision'. *Journal of Applied Non-Classical Logic* **7**(1-2), 151–175.

Katsuno, H. and A. Mendelzon: 1991, 'Propositional knowledge base revision and minimal change'. *Artificial Intelligence* **52**, 263–294.

Kowalski, R.: 1979, *Logic for Problem Solving*. North Holland.

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

Makinson, D.: 1987, 'On the status of the postulate of recovery in the logic of theory change'. *Journal of Philosophical Logic* **16**, 383–394.

McAllester, D. A.: 1980, 'An Outlook on Truth Maintenance'. AI Memo 551, Massachusetts Institute of Technology Artificial Intelligence Laboratory.

McAllester, D. A.: 1990, 'Truth Maintenance'. In: *Proceedings of the Eighth National Conference on Artificial Intelligence (AAAI'90)*. pp. 1109–1116.

McGuinness, D. L., M. K. Abrahams, L. A. Resnick, P. F. Patel-Schneider, R. H. Thomason, V. Cavalli-Sforza, and C. Conati: 1994, 'Classic Knowledge Representation System Tutorial'. http://www.bell-labs.com/project/classic/papers/ClassTut/ClassTut.html.

Nebel, B.: 1989, 'A Knowledge Level Analysis of Belief Revision'. In: R. Brachman, H. J. Levesque, and R. Reiter (eds.): *Principles of Knowledge Representation and Reasoning: Proceedings of the First International Conference*. San Mateo, pp. 301–311.

Nebel, B.: 1992, 'Syntax-based approaches to belief revision'. In: P. Gärdenfors (ed.): *Belief Revision*, Vol. 29. Cambridge, UK: Cambridge University Press, pp. 52–88.

Nebel, B.: 1994, 'Base Revision Operations and Schemes: Representation, Semantics and Complexity'. In: A. G. Cohn (ed.): *Proceedings of the Eleventh European Conference on Artificial Intelligence (ECAI'94)*. Amsterdam, The Netherlands, pp. 341–345.

Rott, H.: 1998, "'Just Because": Taking Belief Bases Seriously'. In: S. R. Buss, P. Hájaek, and P. Pudlák (eds.): *Logic Colloquium '98—Proceedings of the 1998 ASL European Summer Meeting*, Vol. 13 of *Lecture Notes in Logic*. pp. 387–408.

Tennant, N.: 2003, 'Theory-Contraction is NP-Complete'. *Logic Journal of the IGPL* **11**(6), 675–693.

Tennant, N.: 2006, 'New Foundations for a Relational Theory of Theory-revision'. *Journal of Philosophical Logic* **35**(5), 489–528.

Wasserman, R.: 2001, 'Resource-Bounded Belief Revision'. Ph.D. thesis, ILLC, University of Amsterdam.

Williams, M.-A.: 1992, 'Two operators for theory base change'. In: *Proceedings of the Fifth Australian Joint Conference on Artificial Intelligence*. pp. 259–265.

Williams, M.-A.: 1995, 'Iterated Theory Base Change: A Computational Model'. In: *Proceedings of Fourteenth International Joint Conference on Artificial Intelligence (IJCAI-95)*. San Mateo, pp. 1541–1549.

Williams, M.-A.: 1997, 'Anytime Belief Revision'. In: *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence, IJCAI 97*, Vol. 1. pp. 74–81.