Investigations Into a Theory of Knowledge Base Revision: Preliminary Report

Mukesh Dalal[*†]
Rutgers University
Department of Computer Science
New Brunswick, NJ 08903

Abstract

A fundamental problem in knowledge representation is how to revise knowledge when new, contradictory information is obtained. This paper formulates some desirable principles of knowledge revision, and investigates a new theory of knowledge revision that realizes these principles. This theory of revision can be explained at the knowledge level, in purely model-theoretic terms. A syntactic characterization of the proposed approach is also presented. We illustrate its application through examples and compare it with several other approaches.

1 Introduction

At the core of very many AI applications built in the past decade is a knowledge base — a system that maintains knowledge about the domain of interest. Knowledge bases need to be revised when new information is obtained. In many instances, this revision contradicts previous knowledge, so some previous beliefs must be abandoned in order to maintain consistency. As argued in [Ginsberg, 1986], such situations arise in diverse areas such as diagnosis, design, database updates, planning, and natural language understanding. In this paper, we investigate a new theory of knowledge revision.

In [Levesque, 1984a], Levesque presents formal foundations of a functional approach to knowledge representation, where knowledge bases (KBs) are characterized in terms of what they can be asked or told about some domain:

$$Tell: \; KB \times \mathcal{L} \rightarrow KB$$

$$Ask: \; KB \times \mathcal{L} \rightarrow \{yes,\, no,\, unknown\}$$

where $\mathcal{L}$ is some language to talk about the domain. Since *Tell* can be used to tell only information which is consistent with the knowledge base, it is not the appropriate operation for knowledge revision [Levesque, 1984a, page 182]. For this purpose, we add an additional operation:

$$Revise: \; KB \times \mathcal{L} \rightarrow KB$$

Levesque argues that one should define the operations on a KB at the knowledge level [Newell, 1981], independently of the particular symbols/sentences used to build up the KB. In this spirit, we define revision purely in terms of the models of the KB. We also give an equivalent symbol level description by presenting a syntactic method for revising knowledge bases.

We show the relation of our work to research in Philosophy on the formal aspects of the logic of theory change [Makinson, 1985] which has recently attracted attention in the AI community. For any revision scheme, it is desirable that it preserve as much as possible the beliefs held prior to revision. We provide one possible formalization of the notion of *knowledge retained* by a revision scheme. We also briefly discuss some applications of revision.

2 Principles of Knowledge Revision

For the purpose of this paper, we abstractly represent the knowledge in a knowledge base by a finite set of formulae in a propositional logic language $\mathcal{L}$; this set describes the possible states of the world – its models. The revision is presented as a formula in $\mathcal{L}$.

Let $\psi \circ \mu$ denote the revised knowledge obtained by revising the old knowledge ψ by the new information μ, i.e., $\psi \circ \mu = Revise(\psi, \mu)$. The problem of knowledge revision is: given ψ and μ, define $\psi \circ \mu$.

In [Dalal, 1988] we motivate certain principles that should be followed when characterizing the revised knowledge $\psi \circ \mu$. These are:

1. Adequacy of Representation: *The revised knowledge should have the same representation as the old knowledge.* Especially in a functional view of knowledge bases, this is essential since the same operations need to be performed on both. By defining the range of *Revise* as KB, $\psi \circ \mu$ is implicitly required to satisfy this principle.

2. Irrelevance of Syntax: *The revised knowledge base should not depend on the syntax (or representation) of either the old knowledge or the new information.* Thus, if ψ is logically equivalent to ($\approx$) ψ' and $\mu \approx \mu'$, then $\psi \circ \mu \approx \psi' \circ \mu'$. This is essential in order to provide a model-theoretic semantics of the revision process. In view of this, we omit the distinction between a set of formulae (with an implicit conjunction) and a single formula, since one can be expressed in terms of the other such that the two are logically equivalent.

3. Maintenance of Consistency: *If ψ and μ are both consistent, than so is $\psi \circ \mu$.* If not for this, *Revise* and *Tell* could be identical.

4. Primacy of New Information: *The revised knowledge of the system should conform to the new information.* Thus, $\psi \circ \mu \models \mu$. This implies a complete reliance on the truth of the new information.

*Net: dalal@aramis.rutgers.edu

†The author was supported by the CAIP center at Rutgers during this research. Currently on leave to IBM T.J. Watson Research Laboratory, Yorktown Heights, N.Y. for the summer.

5. Persistence of Prior Knowledge: *As much old knowledge as possible should be retained in the revised knowledge.* Thus $\psi \circ \mu$ should be obtained by some form of *minimal change* in ψ. Note that there may be multiple notions of "minimality", but that if $\psi \cup \{\mu\}$ is consistent then $\psi \circ \mu \approx \psi \cup \mu$.

6. Fairness: *If there are many candidates for the revised KB that satisfy the above principles then one of them should not be arbitrarily chosen.* Arbitrariness is clearly undesirable; yet we wish to avoid non-reproducibility (hence non-determinism), and by the specification of Revise, we can rely only on the contents of the KB to choose. One possible solution is to define the revised KB as the "intersection" of all these candidates. Note that this involves compromising the principle of persistence of prior knowledge only to the extent of reflecting this ambiguity.

Note that we do not claim that the above principles are the best for every application. We only make explicit certain principles to characterize the intuitive notion of knowledge base revision.

3 Semantics of Revision

Let Λ be the set of atoms of the underlying language $\mathcal{L}$. An interpretation $\mathcal{A}$ is a truth assignment to the atoms in Λ. An interpretation $\mathcal{A}$ is a model of a formula ψ if ψ evaluates to *true* in $\mathcal{A}$. $\mathcal{A}$ is a model of a set of formulae if it is a model of every formula in the set. Let $mod(\psi)$ denote the set of all models of ψ, where ψ could be a single formula or a set of formulae.

Consider the knowledge base $Revise(\psi, \mu)$. The possible states of the world consistent with ψ are the models of ψ, i.e., $mod(\psi)$. If μ is inconsistent with ψ, μ does not hold in any of them. We can make changes in the models of ψ such that μ holds in (some or all of) these changed interpretations. What type of changes can we make? How do we quantify these changes so that we can formalize the notion of *minimal change*? We answer these questions in this section.

Consider changes first. The smallest change in an interpretation is a change in the truth value of a single atom. Since we do not wish to be biased in favor of any single atom, all changes in truth values of all possible single atoms will be our smallest unit of change in an interpretation.

Definition: If w is an interpretation over a set of atoms Λ, then define

$$g(w) \;=\; \{w' \mid w' \text{ and } w \text{ differ in the truth-value of at most one atom in } \Lambda\}$$

Note that $w \in g(w)$. We can extend the definition of g to sets and formulae:

Definition: If $\mathcal{A}$ is a set of interpretations, define

$$g(\mathcal{A}) \;=\; \bigcup_{w \in \mathcal{A}} g(w)$$

If ψ is a formula or a set of formulae then $G(\psi)$ is defined in terms of its models as[1]

$$mod(G(\psi)) \;=\; g(mod(\psi))$$

[1]Note that while g is a function on interpretation(s), G is a function on a formula or a set of formulae.

Note that $\psi \models G(\psi)$ and that $G(\psi)$ is consistent iff ψ is consistent. g can be read as a generalization operator which takes a set of interpretations and generalizes them to a larger set. G is also a generalization operator which can be thought of as taking a formula or a set of formulae and returning a subset of its logical closure.

Now we have a way to systematically characterize changes in models. We also have a quantitative measure of this change: $g^i(\mathcal{A})$ is a smaller change in $\mathcal{A}$ than $g^j(\mathcal{A})$ iff $i < j$. The definition of $g^i(\mathcal{A})$ is the obvious one: (1) $\mathcal{A}$ if $i = 0$, (2) $g^{i-1}(g(\mathcal{A}))$ otherwise. If ψ is inconsistent with μ, we change the models of ψ by applying the operator g. If we obtain at least one interpretation that makes μ true then we are done: we can define $\psi \circ \mu$ to be $G(\psi) \cup \{\mu\}$. If not, we apply g again and keep on this way.

Let k be the least value of i for which μ holds in some interpretation in the set $g^i(mod(\psi))$. It is clear that this is also the least value of i for which the set of formulae $G^i(\psi) \cup \{\mu\}$ is consistent.

Definition: $\psi \circ \mu \;=\; G^k(\psi) \cup \{\mu\}$, where k is the least value of i for which $G^i(\psi) \cup \{\mu\}$ is consistent.

At first sight it might appear that we are doing an overkill by generalizing ψ with respect to all the ground atoms in it, since the cause of inconsistency might be located in only a few of them. In [Dalal, 1988] we show that revised knowledge is the same (modulo logical equivalence) even if ψ is generalized with respect to only the conflicting atoms. Thus, if it is easier to find the set of conflicting atoms, then it would be advantageous to generalize ψ with respect to the conflicting atoms only.

Example: Let $\psi = \{a, \neg b\}$ and $\mu = b$. Then $mod(\psi) = \{\{a\}\}$[2] and $mod(\mu) = \{\{a, b\}, \{b\}\}$. Since $\psi \cup \{\mu\}$ is inconsistent, we generalize ψ:

$$mod(G(\psi)) = g(mod(\psi)) = \{\phi, \{a\}, \{a, b\}\}$$

Since $G(\psi)$ is consistent with μ, $k = 1$. Thus $Revise(\psi, \mu)$ is $G(\psi) \cup \{\mu\}$, whose only model is $\{a, b\}$. Since we are not interested in exact syntactic representation of a formula, $\psi \circ \mu$ can be expressed as any set of formulae, whose only model is $\{a, b\}$.

4 Syntactic Characterization of Revision

We present a technique to compute $G(\psi)$ by syntactic transformation of ψ, without using models of ψ or invoking any model-theoretic constructions. Since $\psi \circ \mu$ is defined in terms of G we would effectively have a syntactic transformation technique to compute $\psi \circ \mu$. For the purpose of this section, we represent a set of formulae by a conjunction of all the formulae in the set. We use the following lemma and definition from [Weber, 1987]:

Lemma: Let ψ be a formula and α be an atom. There exists formulae ψ_α^+ and ψ_α^- such that (1) ψ_α^+ and ψ_α^- do not contain α, and (2) $\psi \approx (\alpha \wedge \psi_\alpha^+) \vee (\neg\alpha \wedge \psi_\alpha^-)$.

We replace each α in ψ by *true* (or *false*) to obtain ψ_α^+ (or ψ_α^-). The resulting expressions can be simplified by evaluating subexpressions consisting of *false*'s or *true*'s, until all of these constants are eliminated.

[2]We restrict the set of atoms to $\Lambda(\psi \cup \mu)$, and we represent an interpretation by the set of atoms which are assigned *true*.

Definition: Let ψ, α, ψ_α^+ and ψ_α^- be as above; then

$$res_\alpha(\psi) = \psi_\alpha^+ \vee \psi_\alpha^-$$

is called the *resolvent* of ψ with respect to α.

Theorem[3]: Let ψ be a formula and $\{\alpha_1, \ldots, \alpha_p\}$ be the set of atoms occurring in ψ, then

$$G(\psi) = res_{\alpha_1}(\psi) \vee \ldots \vee res_{\alpha_p}(\psi)$$

Thus we have a method to compute $G(\psi)$ given any formula ψ. This method can be used to compute the revised knowledge $\psi \circ \mu$ following the definition in the last section. A caveat in this characterization is that it requires checks of logical consistency in order to establish the minimum k for which $G^k(\psi)$ needs to be computed. This problem is in general NP-Complete for propositional logic. One such check is required for every step of generalization.

Example(continued): Consider the example of the last section. It is convenient to express ψ as $a \wedge \neg b$. ψ and μ conflict in the truth-value of the atom b. Thus, we need to resolve ψ with only b:

$$(\psi)_b^+ = false \qquad (\psi)_b^- = a$$

$$res_b(\psi) = (\psi)_b^+ \vee (\psi)_b^- = a$$

$$\text{Thus,} \quad \psi \circ \mu = G(\psi) \cup \{\mu\} = a \wedge b$$

5 Retained Knowledge

Let $\circ$ be any revision scheme. The revised knowledge $\psi \circ \mu$ is expected to represent the composition of the old knowledge ψ and the new information μ. In this section we will formalize the notion of how much knowledge represented by ψ and μ is retained in $\psi \circ \mu$.

Definition: For formulae ψ and μ and a revision scheme $\circ$, if there exist formulae σ_ψ, σ_μ and $\sigma_{\psi\circ\mu}$ such that $\psi \models \sigma_\psi, \mu \models \sigma_\mu$, and

$$\psi \circ \mu \approx \sigma_\psi \wedge \sigma_\mu \wedge \sigma_{\psi\circ\mu}$$

then σ_ψ is *old-knowledge retained* by $\circ$, σ_μ is *new-knowledge retained* by $\circ$, and $\sigma_{\psi\circ\mu}$ is *extra-knowledge added* by $\circ$.

Theorem: If $\circ$ is a revision scheme then for every ψ and μ the following statements are equivalent:

1. $\psi \wedge \mu \models \psi \circ \mu$ and $\psi \circ \mu \models \mu$.
2. there exists σ_ψ such that $\psi \models \sigma_\psi$ and $\psi \circ \mu \approx \sigma_\psi \wedge \mu$.

Since Principles 4 and 5 (section 2) entail condition 1, this theorem demonstrates that any acceptable revision scheme retains complete new-knowledge and adds no extra-knowledge. Such schemes differ only in the amount of old-knowledge retained. A scheme that retains maximum old-knowledge is more desirable.

Definition: Let $\circ$ and $\circ'$ be any two revision schemes that retain complete new knowledge and add no extra knowledge. $\circ$ is said to *retain at least as much knowledge as* $\circ'$, i.e., $\circ' \preceq \circ$, if for all ψ and μ there exists σ_ψ and σ'_ψ such that

$$\psi \models \sigma_\psi \; ; \; \psi \models \sigma'_\psi$$

$$\psi \circ \mu \approx \sigma_\psi \wedge \mu \; ; \; \psi \circ' \mu \approx \sigma'_\psi \wedge \mu$$

and $\sigma_\psi \models \sigma'_\psi$. $\circ$ is said to *retain more knowledge* than $\circ'$ iff $\circ' \preceq \circ$ and $\circ \not\preceq \circ'$.

[3]Unless otherwise mentioned, proofs appear in [Dalal, 1988].

The following theorem gives a more direct way to determine whether one revision scheme retains more old-knowledge than the other.

Theorem: $\circ$ retains more old-knowledge than $\circ'$ iff for every ψ and μ, $\psi \circ \mu \models \psi \circ' \mu$, while for some ψ and μ, $\psi \circ' \mu \not\models \psi \circ \mu$.

6 Related Work and Applications

6.1 A Logic of Theory Change

[Makinson, 1985] provides an excellent survey on the work by Gärdenfors, Alchourrón and Makinson (GAM) on the formal aspects of a logic of theory change. A theory is defined as a set of propositions (formulae) closed under logical consequence, i.e., A is a theory iff $Cn(A) = A$, where Cn is a consequence operation. Three operations are defined on a set of propositions A — *expansion*, where a new proposition x is set-theoretically added to A; *contraction* ($A \dot{-} x$), where a proposition x which is in the theory $Cn(A)$ is rejected; and *revision* ($A \dot{+} x$), where a proposition x inconsistent with the theory $Cn(A)$ is added to it under the requirement that the revised theory be consistent. This operation of revision is very similar to the notion of revision introduced in this paper.

Gärdenfors developed some general postulates that seem desirable for contraction and revision. His postulates for revision can be expressed as:

($G1$) $A \dot{+} x$ is always a theory;
($G2$) $x \in A \dot{+} x$;
($G3$) If $\neg x \notin Cn(A)$ then $A \dot{+} x = Cn(A \cup \{x\})$;
($G4$) If $\neg x \notin Cn(\phi)$ then $A \dot{+} x$ is consistent;
($G5$) If $Cn(x) = Cn(y)$ then $A \dot{+} x = A \dot{+} y$;
($G6$) $A \dot{+} (x \wedge y) \subseteq Cn((A \dot{+} x) \cup \{y\})$ for any theory A;
($G7$) $Cn((A \dot{+} x) \cup \{y\}) \subseteq A \dot{+} (x \wedge y)$ for any theory A, provided that $\neg y \notin A \dot{+} x$.

For a theory A, contraction is then defined using the identity: $A \dot{-} x = (A \dot{+} \neg x) \cap A$, and conversely (for any set of propositions A) : $A \dot{+} x = Cn((A \dot{-} \neg x) \cup \{x\})$.

How does our approach compare with that of GAM?

A superficial difference between the approach of GAM and the one presented here is that theirs is *defined* in terms of the set of formulas expressing the KB. If the KB is taken however to be the logical closure of these formulas (as suggested by a knowledge-level approach) this difference disappears. In fact, in the expanded version of this paper we characterize the revision schemes $\dot{+}$ satisfying G1-G7 in model-theoretic terms. On the other hand, the GAM approach is more general since it applies to any logic for which a notion of logical closure Cn is defined, while ours currently applies to only standard propositional logic.

We do however have

Theorem: The revision scheme $\circ$ satisfies the Gärdenfors postulates G1-G7.

It is obvious that $\circ$ satisfies axioms G1-G5, and simple model-theoretic arguments establish conditions G6-G7[4].

[4]There is also a proof involving the notion of "partial meet function" introduced in [Alchourron *et al.*, 1985].

There are of course many other revision schemes that satisfy these conditions, including defining $A \dot{+} x$ to be just $Cn(x)$ whenever $\neg x \in Cn(A)$; and $Cn(A \cup \{x\})$ otherwise. Our scheme is more *conservative* than at least some of those satisfying G1-G7: it preserves more old-knowledge than the previous admittedly trivial revision, as will be shown in the example of section 6.4.

Secondly, not all revision schemes satisfying axioms G1-G7 satisfy our postulate of *fairness*: One form of fairness would be to require that the result of $\dot{+}$ not depend on the accidents of naming propositions; i.e., if f is an isomorphism on Λ, then we would expect $f(A \dot{+} x) = f(A) \dot{+} f(x)$. Even some of the revisions considered in [Makinson, 1985] are unfair in this sense: they pick arbitrary maximally consistent subsets of A which do not contain x.

Finally, GAM do not suggest any algorithm to implement their constructions, although they do have a theorem characterizing the acceptable revisions in terms of maximally consistent subsets of A which do not entail x. Such a definition would seem to be much more difficult to implement than that presented in Section 4; but then $\circ$ requires tests of consistency, so we cannot make any great claims to efficiency.

Observe also that the definition of $\circ$ shows that, contrary to the intuitions voiced in [Makinson, 1985], contraction is not necessarily more primitive/basic than revision: defining $\circ$ does not involve contraction[5].

6.2 Counterfactuals

A counterfactual is a statement like "if p, then q", where the premise p is either known or expected to be false. It is represented as $p > q$ and is defined to be *true* in a world[6] iff q is true in every *most similar (possible) world* in which the premise p holds. In an excellent paper [Ginsberg, 1986], Ginsberg presents a formal description of counterfactual implication and discusses the issues involved in implementing it.

In our framework, $p > q$ in a world ψ is defined to be *true* iff "$\psi \circ p \models q$". There is only one most similar possible world — $\psi \circ p$.

Given a world F, [Ginsberg, 1986] defines a partial order among the subsets of F based on set inclusion. The set of possible worlds for p in F is defined to be:

$$W(p,F) = \{T \subseteq F \mid T \not\models \neg p, \neg B(T) \text{ and } \forall U,\ T \subset U \subseteq F \Rightarrow U \models \neg p \text{ or } B(U)\}$$

The predicate B is called the *badworld* predicate. Its purpose is to rule out certain worlds, say, which are completely meaningless. $p > q$ is defined to be *true* in a world F iff for every $T \in W(p,F)$, $T \cup \{p\} \models q$.

Because of certain examples involving counterfactual statements, Ginsberg opts for a definition of $p > q$ which depends on the syntactic form of p and q. As such, his definition clearly differs from our semantic definition. In the full version of this paper, we plan to show the relationship of Ginsberg's definition to the work of AGM, and hence further relate it to our own definition of $\circ$.

6.3 Diagnosis from First Principles

Assume one is first given a description of some system (say, a physical device) and then an observation of the system's behavior. If the two are inconsistent then one is confronted with a diagnostic problem, namely, to determine those system components whose abnormal behavior can account for this discrepancy.

Suppose ψ is the system's description, where there are propositions asserting the *normality* of all components; and suppose μ is an observation that is inconsistent with ψ. By *protecting* all but the normality propositions, it is possible to view $\psi \circ \mu$ as representing the revised description of the system[7]. This revised description will implicitly contain information about all abnormal components – the ones for which normality propositions do not hold.

In [Reiter, 1987], Reiter proposes a theory of diagnosis from first principles (references to other work on diagnosis can be found in Reiter's paper) which starts form the same initial ψ and μ. He then suggests an algorithm which produces the set of abnormal components explicitly. Space limitations only permit us to state that our scheme would find only those diagnoses which involve the least number of abnormal components: thus if one diagnosis blamed component b, and the other components c,d and e, then using $\circ$ only the former would be reported, while [Reiter, 1987] would report both.

6.4 Updates in Logical Databases

A database can be considered as a set of formulae which models our knowledge about the real world. One can add new information to the database and query it about its current knowledge. Given new information, the update problem is to define and compute the revised state of the database. Notable approaches to solving this problem have been suggested by Fagin, Ullman and Vardi [Fagin *et al.*, 1983], Borgida [Borgida, 1985], Winslett [Winslett, 1986] and Weber [Weber, 1987]. We suggest that the update should be considered as the revision operator $\circ$.

Example: Let $\psi = \{a \wedge b\}$ and $\mu = \neg a \vee \neg b$. Since $\psi \cup \{\mu\}$ is inconsistent, we generalize ψ with respect to both a and b.

$$G(\psi) = res_a(\psi) \vee res_b(\psi) = a \vee b$$

Since $G(\psi)$ is consistent with μ, we are done: the revised knowledge is $G(\psi) \cup \{\mu\} \approx \{(a \wedge \neg b) \vee (\neg a \wedge b)\}$. Given the model-theoretic nature of our revision mechanism, the result of the update will be the same whether the knowledge base is presented as above, or as $\{a, b\}$ or even $\{a, b, a \wedge b\}$.

In contrast, all four of [Fagin *et al.*, 1983], [Winslett, 1986], [Ginsberg, 1986] and [Weber, 1987] obtain $\{\neg a \vee \neg b\}$ as the revised database. They are thus less "conservative", losing all the knowledge in the hypotheses set ψ. Moreover, [Ginsberg, 1986] and [Fagin *et al.*, 1983] would report a different answer (the one produced by our mechanism) if the database was presented as $\{a, b\}$, but not as $\{a \wedge b\}$! It seems counter-intuitive that updates should produce different results even in such relatively minor variations in the

[5] $G^k(\psi)$ is *not* $\psi \dot{-} \mu$!

[6] A world is a set of propositions, which are not necessarily atomic.

[7] A protected formula must hold even after the revision. It is like an integrity constraint in a database.

syntax of the database – even the limited logic of explicit beliefs in [Levesque, 1984b] considers these formulations equivalent!

In a more complete version of this paper, we relate the other update schemes to the Gärdenfors postulates. In [Dalal, 1988] we show that ∘ preserves more old knowledge than them.

7 Conclusions

The major contribution of this paper is a *semantic* definition of revision in propositional knowledge bases, providing a new point in the spectrum of approaches to this old-standing problem. This definition is founded on a number of a priori principles (especially minimality of change and fairness) and is also given a syntactic characterization. The application of the approach in several domains is also discussed. The notion of old-knowledge retained is formalized, and the approach defined in this paper is shown to retain more old-knowledge then some previous proposals.

In addition to the results mentioned earlier, we also propose to investigate the extension of this work in several directions:

- Establish further criteria for fairness and preservation of old knowledge, and evaluate all the proposals against these.

- Extend the language of revisions to first order logic and epistemic languages like Levesque's FOPC.

- Extend the notion of KB to allow differential treatment of certain atoms, or even formulas (e.g., integrity constraints in a data base), so that some beliefs are more easily given up. This of course relaxes the principle of fairness.

Acknowledgements

I wish to express my gratitude to Alex Borgida for his advice and help at every stage of this research. This paper has benefited from the discussions and comments of my colleagues, including T. Imielinski, J. Radhakrishnan, J. Chomicki, K. Williamson, A. Prieditis, T. Marlowe and J. Koeneman. In addition, I wish to thank D. Etherington, J.-F. Lamy, and the two anonymous referees for their extremely useful suggestions.

References

[Alchourron *et al.*, 1985] C. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: partial meet contraction and revision functions. *The Journal of Symbolic Logic*, 50(2):510–530, June 1985.

[Borgida, 1985] A. Borgida. Language features for flexible handling of exceptions in information systems. *ACM — Transactions on Database Systems*, 10(4):565–603, 1985.

[Dalal, 1988] M. Dalal. *Updates in Propositional Databases.* Technical Report DCS-TR-222, Department of Computer Science, Rutgers University, February 1988.

[Fagin *et al.*, 1983] R. Fagin, J.D. Ullman, and M.Y. Vardi. On the semantics of updates in databases. In *Proc. of 2nd ACM Symp. on Principles of Database Systems*, pages 352–365, Atlanta, March 1983.

[Ginsberg, 1986] M.L. Ginsberg. Counterfactuals. *Artificial Intelligence*, 30(1):35–79, 1986.

[Levesque, 1984a] H.J. Levesque. Foundations of a functional approach to knowledge representation. *Artificial Intelligence*, 23(2):155–212, 1984.

[Levesque, 1984b] H.J. Levesque. A logic of explicit and implicit belief. *Proceedings of AAAI-84*, pages 198–202, 1984.

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

[Newell, 1981] A. Newell. The knowledge level. *The AI Magazine*, 2(2):1–20, 1981.

[Reiter, 1987] R. Reiter. A theory of diagnosis from first principles. *Artificial Intelligence*, 32(1):57–95, 1987.

[Weber, 1987] Andreas Weber. Updating propositional formulas. In L. Kerschberg, editor, *Expert Database Systems*, pages 487–500, The Benjamin/Cummings Publishing Company, Inc., 1987.

[Winslett, 1986] M. Winslett. Is belief revision harder than you thought? In *Proceedings of AAAI-86*, pages 421–427, 1986.