Relating Belief Revision and Circumscription*

Paolo Liberatore and Marco Schaerf
Dipartimento di Informatica e Sistemistica
Università di Roma "La Sapienza"
via Salaria 113, I-00198 Roma, Italy
email: {liberato,schaerf}@assi.dis.uniroma1.it

Abstract

Nonmonotonic formalisms and belief revision operators have been introduced as useful tools to describe and reason about evolving scenarios. Both approaches have been proven effective in a number of different situations. However, little is known about their relationship. Previous work by Winslett has shown some correlations between a specific operator and circumscription. In this paper we greatly extend Winslett's work by establishing new relations between circumscription and a large number of belief revision operators. This highlights similarities and differences between these formalisms. Furthermore, these connections provide us with the possibility of importing results in one field into the other one.

1 Introduction

During the last years, many formalisms have been proposed in the AI literature to model commonsense reasoning. Particular emphasis has been put in the formal modeling of a distinct feature of commonsense reasoning, that is, its nonmonotonic nature. The AI goal of providing a logic model of human agents' capability of reasoning in the presence of incomplete or contradictory information has proven to be a very hard one. Nevertheless, many important formalisms have been put forward in the literature.

Two main approaches have been proposed to handle the nonmonotonic aspects of commonsense reasoning. The first one deals with this problem, by defining a new logic equipped with a nonmonotonic consequence operator. Important examples of this approach are *default logic* proposed in [Reiter, 1980] and *circumscription* introduced in [McCarthy, 1980]. The second one relies on preserving a classical (monotonic) inference operator, but introduces a *revision* operator

*Work partially supported by the ESPRIT Basic Research Action N. 6810 COMPULOG II, N. 6471 MEDLAR and the Progetto Finalizzato Informatica & Calcolo Parallelo of the CNR (Italian Research Council)

that accommodates a new piece of information into an existing body of knowledge. Specific revision operators have been introduced, among the others, in [Ginsberg, 1986] and in [Dalal, 1988]. A general framework for revision has been proposed by Alchourron, Gärdenfors and Makinson in [Alchourrón *et al.*, 1985, Gärdenfors, 1988]. A close variant of revision is *update*. The general framework for update has been studied in [Katsuno and Mendelzon, 1989, Katsuno and Mendelzon, 1991] and specific operators have been proposed in [Winslett, 1990] and [Forbus, 1989].

In this paper we investigate the relationship between circumscription and many operators for belief revision and update. A first study of these relations has been done in [Winslett, 1989], where she relates her operator to circumscription. We expand her results showing similar connections between several other belief revision operators and circumscription. To this end, we also introduce a variant of circumscription based on cardinality, rather than set-containment.

The established correlations highlight the relations between the two fields. Moreover, as side benefits, they provide us with the opportunity to import results in one field into the other one.

A distinct approach to model the nonmonotonic aspect of commonsense reasoning is via a logic of actions. Even though this aspect is out of the scope of this paper, we want to point out the results presented in [Kharta and Lifschitz, 1994] where it is shown how to express a very general logic of action using circumscription.

The paper is organized as follows: In Section 2 we recall some key definitions and results for belief revision and circumscription, introduce a variant of circumscription (NCIRC) and explain the notation used throughout the following sections. In Section 3 we show the main relations between revision operators and circumscription, while in Section 4 we show relations and reductions between the various operators. In Section 5 we focus on syntactically-restricted knowledge bases. Section 6 discusses the impact of our results with particular attention to the computational complexity analysis. Finally, in Section 7 we draw some conclusions.

2 Preliminaries

In this section we (very briefly) present the background and terminology needed to understand the results presented later in the paper. For the sake of simplicity, throughout this paper we restrict our attention to a (finite) propositional language.

The *alphabet* of a propositional formula is the set of all propositional atoms occurring in it. Formulae are built over a finite alphabet of propositional letters using the usual connectives $\neg$ (not), $\vee$ (or) and $\wedge$ (and). Additional connectives are used as shorthands, $\alpha \rightarrow \beta$ denotes $\neg\alpha \vee \beta$, $\alpha = \beta$ is a shorthand for $(\alpha \wedge \beta) \vee (\neg\alpha \wedge \neg\beta)$ and $\alpha \neq \beta$ denotes $\neg(\alpha = \beta)$.

An *interpretation* of a formula is a truth assignment to the atoms of its alphabet. A *model* M of a formula F is an interpretation that satisfies F (written $M \models F$). Interpretations and models of propositional formulae will be denoted as sets of atoms (those which are mapped into 1). A theory T is a set of formulae. An interpretation is a model of a theory if it is a model of every formula of the theory. Given a theory T and a formula F we say that T entails F, written $T \models F$, if F is true in every model of T. Given a propositional formula or a theory T, we denote with $\mathcal{M}(T)$ the set of its models. We say that T is consistent, written $T \not\models \bot$, if $\mathcal{M}(T)$ is non-empty.

2.1 Belief Revision and Update

Belief revision is concerned with the modeling of accommodating a new piece of information (the revising formula) into an existing body of knowledge (the knowledge base), where the two might contradict each other. A slightly different perspective is taken by knowledge update. An analysis of the relative merits of revision and update is out of the scope of this paper, or an interesting discussion on the differences between belief revision and update we refer the reader to the work [Katsuno and Mendelzon, 1991]. We assume that both the revising formula and the knowledge base can be either a single formula or a theory.

We now recall the different approaches to revision and update, classifying them into formula-based and model-based ones. A more thorough exposition can be found in [Eiter and Gottlob, 1992]. We use the following conventions: the expression $card(S)$ denotes the cardinality of a set S, and symmetric difference between two sets S_1, S_2 is denoted by $S_1 \Delta S_2$. If S is a set of sets, $\cap S$ denotes the set formed intersecting all sets of S, and analogously $\cup S$ for union; $min_{\subseteq} S$ denotes the subset of S containing only the minimal (w.r.t. set inclusion) sets in S, while $max_{\subseteq} S$ denotes its maximal sets.

Formula-based approaches operate on the formulae syntactically appearing in the knowledge base K. Let $C(K, A)$ be the set of the subsets of K which are consistent with the revising formula A:

$$C(K, A) = \{K' \subseteq K \quad | \quad K' \cup \{A\} \not\models \bot\}$$

and let $W(K, A)$ be the set of the maximal sets of $C(K, A)$:

$$W(K, A) = max_{\subseteq} C(K, A)$$

The set $W(K, A)$ contains all the plausible subsets of K that we may retain when inserting A.

Ginsberg. In [Ginsberg, 1986] the revised knowledge base is defined as a set of theories: $K *_G A \doteq \{K' \cup \{A\} \mid K' \in W(K, A)\}$. That is, the result of revising K is the set of all maximal subsets of K consistent with A, plus A. Logical consequence in the revised knowledge base is defined as logical consequence in each of the theories, i.e. $K *_G A \models Q$ iff for all $K' \in W(K, A)$, $K' \cup \{A\} \models Q$. In other words, Ginsberg considers all sets in $W(K, A)$ equally plausible and inference is defined skeptically, i.e. Q must be a consequence of each set.

Model-based approaches instead operate by selecting the models of A on the basis of some notion of proximity to the models of K. Model-based approaches assume K to be a single formula, if K is a set of formulae it is implicitly interpreted as the conjunction of all the elements. Many notions of proximity have been defined in the literature. We distinguish them between pointwise proximity and global proximity.

We first recall approaches in which proximity between models of A and models of K is computed pointwise w.r.t. each model of K. That is, they select models of K one-by-one and for each one choose the closest model of A. These approaches are considered as more suitable for knowledge update [Katsuno and Mendelzon, 1991]. Let M be a model, we define $\mu(M, A)$ as the set containing the minimal differences (w.r.t. set inclusion) between each model of A and the given M; more formally, $\mu(M, A) \doteq min_{\subseteq}\{M \Delta N \mid N \in \mathcal{M}(A)\}$.

Winslett. The work [Winslett, 1990] defines the models of the updated knowledge base as $\mathcal{M}(K *_W A) \doteq \{N \in \mathcal{M}(A) \mid \exists M \in \mathcal{M}(K) : M \Delta N \in \mu(M, A)\}$. In other words, for each model of K it chooses the closest (w.r.t. set-containment) model of A.

Borgida. This operator $*_B$, defined in [Borgida, 1985], coincides with Winslett's one, except in the case when A is consistent with K, in which case Borgida's revised theory is simply $K \cup \{A\}$.

Forbus. This approach [Forbus, 1989] takes into account cardinality: Let $k_{M,A}$ be the minimum cardinality of sets in $\mu(M, A)$. The models of Forbus' updated theory are $\mathcal{M}(K *_F A) \doteq \{N \in \mathcal{M}(A) \mid \exists M \in \mathcal{M}(K) : card(M \Delta N) = k_{M,A}\}$. Note that by means of cardinality, Forbus can compare (and discard) models which are incomparable in Winslett's approach.

We now recall approaches where proximity between models of A and models of K is defined considering globally *all* models of K. In other words, these approaches consider at the same time all pairs of models $M \in \mathcal{M}(K)$ and $N \in \mathcal{M}(A)$ and find all the closest pairs. Let $\delta(K, A) \doteq min_{\subseteq} \bigcup_{M \in \mathcal{M}(K)} \mu(M, A)$.

Satoh. In [Satoh, 1988], the models of the revised knowledge base are defined as $\mathcal{M}(K *_S A) \doteq \{N \in \mathcal{M}(A) \mid \exists M \in \mathcal{M}(K) : N \Delta M \in \delta(K, A)\}$. That is, Satoh selects all closest pairs (by set-containment of the difference set) and then projects on the models of A.

Dalal. This approach is similar to Forbus', but global. Let $k_{K,A}$ be the minimum *cardinality* of sets in $\delta(K, A)$; in [Dalal, 1988] the models of a revised theory are defined as $\mathcal{M}(K *_D A) \doteq \{N \in \mathcal{M}(A) \mid \exists M \in \mathcal{M}(K) : card(N \Delta M) = k_{K,A}\}$. That is, Dalal selects all closets pairs (by cardinality of the difference set) and then projects on the models of A.

The complexity of deciding $K * A \models Q$ (where $*$ is one of $\{*_G, *_W, *_B, *_F, *_S, *_D\}$, K, A and Q are the input) was studied in [Eiter and Gottlob, 1992]: in Dalal's approach, the problem is $\Delta_2^p[\log n]$-complete, while in all other approaches it is

Π_2^p-complete.

2.2 Circumscription

Circumscription has been originally introduced in [McCarthy, 1980]. Further extensions have been proposed by several authors. Here we stick to the semantic formulation of circumscription and restrict our interest to a propositional language. Following [Lifschitz, 1985], we define:

Definition 1 *Let T be a propositional formula, $X = \{x_1, \ldots, x_n\}$ its alphabet, P and Z disjoint sets of letters partitioning X (*i.e. $P \cup Z = X$*) and $M \in \mathcal{M}(T)$. M is called a P-minimal model of T if there is no model N of T such that $(N \cap P) \neq (M \cap P)$ and $(N \cap P) \subset (M \cap P)$.*

Definition 2 *The circumscription of T w.r.t. the two sets of letters P and Z, denoted as $CIRC(T; P, Z)$, is the set of all P-minimal models of T,* i.e. *$M \models CIRC(T; P, Z)$ iff M is a P-minimal model of T.*

Informally, P is the set of letters we want to minimize, while letters in Z are allowed to vary. Notice that we are using a version of circumscription where fixed predicates are not allowed. Due to the results of [de Kleer and Konolige, 1989] on eliminating fixed predicates, this restriction does not lead to any loss of expressiveness.

2.3 Cardinality-based Circumscription

The minimality criterion of circumscription is based on set-containment. We now introduce, for propositional languages, a version of circumscription based on cardinality.

Definition 3 *Let T be a propositional formula, $X = \{x_1, \ldots, x_n\}$ its alphabet, P and Z disjoint sets of letters partitioning X (*i.e. $P \cup Z = X$*) and $M \in \mathcal{M}(T)$. M is called a P-cardinality-minimal model of T if there is no model N of T such that $|N \cap P| < |M \cap P|$.*

Definition 4 *The cardinality-based circumscription of T w.r.t. the two sets of letters P and Z, denoted as $NCIRC(T; P, Z)$, is the set of all P-cardinality-minimal models of T,* i.e. *$M \models NCIRC(T; P, Z)$ iff M is a P-cardinality-minimal model of T.*

In other words, I am preferring models with the least number of true letters of the set P, rather than models with a least set of true letters.

2.4 Notations

In order to make formulae more compact and easier to understand, we introduce a number of notations that we use in the rest of the paper.

In the following sections, we make use of variable renaming. To make this clear, we explicitly mention over which alphabet a formula is built upon. More precisely, let $X = \{x_1, \ldots, x_n\}$ be a set of letters, we denote as $T(X)$ a formula built over X. Given a new alphabet $Y = \{y_1, \ldots, y_n\}$ one-to-one with X, with $T(Y)$ we denote the formula with the same structure of $T(X)$ but every occurrence of x_i is replaced by y_i for all $i \geq 1$ and $i \leq n$. For example, let

$$T(X) = (x_1 \wedge (\neg x_3 \vee x_2))$$

then the formula $T(Y)$ is $(y_1 \wedge (\neg y_3 \vee y_2))$.

In order to make the formulae more compact and readable, we overload the boolean connectives to apply to sets of letters. For example, given three disjoint sets of letters W, S and R with the same number of elements k, we use the notation $\neg S$ as a shorthand for the formula $\bigwedge\{\neg s_i | s_i \in S\}$, $S = R$ to denote $\bigwedge\{s_i = r_i | 1 \leq i \leq k\}$, $S \neq R$ to denote $\bigwedge\{s_i \neq r_i | 1 \leq i \leq k\}$ and $W = (S \neq R)$ for $\bigwedge\{w_i = (s_i \neq r_i) | 1 \leq i \leq k\}$.

3 General Cases

In this section we establish relations between circumscription and the various belief revision operators. Due to the lack of space we cannot present complete proofs for all the results, but we provide a sketch of some of the proofs.

3.1 Dalal's revision

The links between the cardinality-based circumscription and Dalal's revision [Dalal, 1988] are very simple. This is due to the similarity of these operations: $NCIRC$ takes the models with a minimum number of positive atoms of the set P, whereas Dalal's revision selects the models of A with a minimum number of differences with models of K.

To translate $NCIRC(T; P, Z)$ into a Dalal's revision it is enough to revise the knowledge base with all literals in P negated. More precisely, we have

$$NCIRC(T; P, Z) = (\neg P) *_D T$$

In fact, the cardinality-minimal models of T are exactly the models of T closer to the knowledge base $\neg P$. The above relation is simple because revision seems somewhat more powerful than $NCIRC$. In fact, it has $NCIRC$ as a sub-case, where K is a set of literals. However, it can be shown that Dalal's revision can be translated into cardinality-based circumscription.

Given three disjoint sets of letters X, Y and W, each one containing n letters, we denote with $\Gamma(X, Y, W)$ the formula $K(Y) \wedge A(X) \wedge (W = (X \neq Y))$. Γ

admits a model M iff $M_X = (M \cap X)$ is a model of A and $M_Y = (M \cap Y)$ is a model of T. Which letters of W will belong to $M_W = (M \cap W)$ is uniquely determined by M_X and M_Y. In fact, $w_i \in M_W$ if and only if $x_i \in M_X$ and $y_i \notin M_Y$ or $x_i \notin M_X$ and $y_i \in M_Y$.

If we force M_W to contain a minimal number of letters, only the models of Γ where the differences between the assignments to X and Y are as few as possible, will be retained. Thus we obtain

$$K(X) *_D A(X) = NCIRC(\Gamma(X,Y,W); W, X \cup Y)$$

where we minimize the letters in W, but not those in $X \cup Y$. More precisely we have:

Theorem 1 *For any model M of $K(X)\ *_D\ A(X)$ there exists a model N of $NCIRC(\Gamma(X,Y,W); W, X \cup Y)$ such that $M = N \cap X$. Furthermore, for any model N of $NCIRC(\Gamma(X,Y,W); W, X \cup Y)$ $N \cap X$ is a model of $K(X)\ *_D\ A(X)$.*

Proof (sketch). We first prove that for any model $M_X \subseteq X$ of $K(X) *_D A(X)$ there exists two sets $M_Y \subseteq Y$ and $M_W \subseteq W$ such that $M = M_X \cup M_Y \cup M_W$ is a model of $NCIRC(K(Y) \wedge A(X) \wedge (W = (X \neq Y)); W, X \cup Y)$. Since M_X is a model of $K(X) *_D A(X)$, it follows that $M_X \models A$ and that there exist a model $N_X \in \mathcal{M}(K(X))$ such that $card(M_X \Delta N_X) = k$, where k is the minimum distance between models of K and models of A. We define $M_Y = \{y_i | x_i \in N_X\}$ and $M_W = \{w_i | ((x_i \in M_X) and (y_i \notin M_Y)) or ((x_i \notin M_X) and (y_i \in M_Y))\}$. Obviously, it holds $M \models K(Y) \wedge A(X) \wedge (W = (X \neq Y))$. Furthermore, we can prove by contradiction that M is a cardinality-minimal model.

We now show that for any model M of $NCIRC(K(Y) \wedge A(X) \wedge (W = (X \neq Y)); W, X \cup Y)$, the set $M_X = M \cap X$ is a model of $K(X) *_D A(X)$. It immediately follows that $M_X \models A(X)$, if M_X is one of the models of A closer to models of K the thesis follows, so assume to the contrary that there exists a $N_X \subseteq X$, different from M_X, such that $N_X \models A(X)$, the distance of M_X from the closest model of $K(X)$ is k_M, the distance of N_X from the closest model of $K(X)$ is k_N and $k_N < k_M$. Let $V_X \subseteq X$ be one of the models of $K(X)$ closer to N, $N_Y = \{y_i | x_i \in V_X\}$, $N_W = \{w_i | ((x_i \in N_X) and (y_i \in N_Y)) or ((x_i \notin N_X) and (y_i \notin N_Y))\}$ and $N = N_X \cup N_Y \cup N_W$. Obviously N is a model of $K(Y) \wedge A(X) \wedge (W = (X \neq Y))$ moreover, the cardinality of $N \cap W$ is k_N, the cardinality $M \cap W$ is k_M. Since $k_n < k_M$ it follows that M is not a cardinality-minimal model of $K(Y) \wedge A(X) \wedge (W = (X \neq Y))$, hence contradiction arises. $\square$

3.2 Satoh's revision

The same reductions between Dalal's revision and cardinality-based circumscription hold between Satoh's revision [Satoh, 1988] and usual (set-containment-based) circumscription.

$$CIRC(T; P, Z) = (\neg P) *_S T$$

To reduce Satoh's revision into circumscription, we use the same relation adopted to reduce Dalal's revision into $NCIRC$

$$\begin{aligned} T(X, Y, Z) &= K(Y) \wedge A(X) \wedge (W = (X \neq Y)) \\ K(X) *_S A(X) &= CIRC(T(X, Y, Z)); W, X \cup Y) \end{aligned}$$

The models of T can be decomposed in a model of K, a model of A, and the difference between them. Circumscription minimizes W, hence makes them as close as possible, where closeness is w.r.t. set containment.

3.3 Winslett's update

Winslett's update method modifies models of K one-by-one, replacing each one with the closest one within the models of A. Local proximity methods are better related to circumscription where all letters are minimized. Circumscription without varying letters is immediately expressed as

$$CIRC(T(X); X, \emptyset) = \neg X *_W T(X)$$

In order to correlate Winslett's update with circumscription, we must be sure that to each distinct model of K correspond incomparable models in the circumscriptive theory. This is obtained by the following reduction:

$$T(X, Y, Z) = K(Y) \wedge (Y \neq Z) \wedge A(X) \wedge (W = (X \neq Y))$$

$$K(X) *_W A(X) = CIRC(T(X, Y, Z)); Y \cup Z \cup W, X)$$

The sub-formula $Y \neq Z$ guarantees that every two models M and N of K, with different assignments to Y, are incomparable because it cannot be the case that $(M \cap Y) \subseteq (N \cap Y)$ and $(M \cap Z) \subseteq (N \cap Z)$ at the same time. The above reduction can be rephrased so as to eliminate varying letters (X). In fact, we obtain:

$$\begin{aligned} T(X, Y, W, Z) &= K(Y) \wedge (Y \neq Z) \wedge A(Y \neq W) \\ K(X) *_W A(X) &= CIRC(T(X, Y, W, Z)); Y \cup Z \cup W, \emptyset) \\ &\quad \wedge [X = (W \neq Y)] \end{aligned}$$

where $A(Y \neq W)$ denotes A where all occurrences of x_i are replaced by $y_i \neq w_i$ for all $i \geq 1$ and $i \leq n$.

3.4 Borgida's revision

Borgida's revision operator [Borgida, 1985] is very similar to Winslett's one, the only difference being that the result of the first one has to be $K \wedge A$ when not contradictory. Since $*_B$ and $*_W$ coincide when K is a set of literals, the reduction

$$CIRC(T(X); X, \emptyset) = \neg X *_B T(X)$$

holds for $*_B$.

In the other direction, one can find a direct transformation from Borgida's revision into circumscription, very much like Winslett's one.

The fact that the result must be $K \wedge A$ can be taken into account by selecting the models of this formula as minimal.

$$K(X) *_B A(X) = CIRC(T(X, Y, R, Z, W); R \cup Z \cup W, X \cup Y)$$

where T is defined as follows:

$$\begin{aligned} T(X, Y, R, Z, W) \quad = \quad & [K(X) \vee ((Y = R) \wedge (R \neq Z))] \wedge \\ & \wedge K(Y) \wedge A(X) \wedge [W = (X \neq Y)] \end{aligned}$$

This reduction coincides with Winslett's one, exception made for the subformula $K(X) \vee ((Y = R) \wedge (R \neq Z))$. Given a model X of P, if there is a model of K with the same value, the formula T has a model $(X, Y = X, R = \emptyset, Z = \emptyset)$ that is surely minimal.

If such a model does not exist, the only models of $T(X, Y, R, Z)$ are those having an Y such that $K(Y)$ is true, R equal to Y and Z the complement of Y. Hence, we have only the models of Winslett's transformation, so the result coincides with the result of applying Winslett's revision. This is exactly the definition of revision given by Borgida.

3.5 Ginsberg's revision

Ginsberg's revision is quite similar to Satoh's principle of minimization. The main difference between them is that the latter minimizes distance given as set of literals, while the first one maximizes the number of true formulae of K.

Two simple relations correlating Ginsberg's revision and circumscription are the following ones:

$$\begin{aligned} CIRC(T; P, Z) \quad &= \quad \neg P *_G T \\ \{f_1, \ldots, f_m\} *_G A(X) \quad &= \quad CIRC(T(X, Y); Y, X) \end{aligned}$$

where $T(X, Y)$ is defined as

$$T(X, Y) = A(X) \wedge (y_1 \neq f_1) \wedge \cdots \wedge (y_m \neq f_m)$$

The first formula follows from the known fact that if K is a set of literals, then $K *_G A = K *_S A$ (see [Eiter and Gottlob, 1992]).

Regarding the second reduction, remind that Ginsberg's revision finds the maximal subsets K' of K s.t. K' and A are not contradictory, whereas the circumscription of a formula takes only the models with a maximal set of false variables. Hence, to revise a set of formulae, we have to include a formula f_i if and only if a variable y_i (not contained in the original formulae) is false. The transformation follows.

3.6 Forbus' update

While all the above reductions are quite simple and straightforward, what we found for Forbus' update operator $*_F$ [Forbus, 1989] is much more complex. Since the resulting formula is somehow cumbersome and difficult to read, we prefer to give a sketch of the steps needed in the reduction. Similarly to the other local-proximity model-based formalisms, we establish relations between Forbus' update and circumscription (or $NCIRC$) where all letters are minimized.

We first observe how circumscription and $NCIRC$ can be expressed using Forbus' update. Reduction of $NCIRC$ to Forbus' operator is trivial:

$$NCIRC(T(X); X, \emptyset) = \neg X *_F T(X)$$

It is also possible to reduce circumscription to Forbus' update. Very briefly, this can be obtained by adding a suitable number of new variables ($O(n^3)$) and imposing that to each (set-containment) minimal model of $T(X)$ corresponds a cardinality-minimal model over the extended alphabet.

The reduction of Forbus' update to circumscription is very similar to Borgida's one. We have only to take in account that Forbus' update is based upon a minimizzation of the cardinality of the distances between models.

Consider the formula:

$$K(X) *_F A(X) = CIRC(T(X, Y, W, V); V, X \cup Y \cup W)$$

where $T(X, Y, W, V)$ expresses the fact that that X is a model of A, that Y is a model of T, and V is a valutation of the cardinality of the distance between them.

Now let

$$\begin{aligned} T(X, Y, W, V) = K(Y) \wedge A(X) \wedge (Y \neq Z) \wedge (W = (X \neq Y)) \\ \wedge EQ(W, V) \wedge BEGIN(V) \end{aligned}$$

Given a model of T and a model of A, this formula has exactly one model, namely the model in which X represents the model of A and Y the model of T. The variables W represents the simmetrical difference between X and Y. The formula $EQ(W, V)$ is the polynomial-size formula that is true if and only if W and V have exactly the same number of positive literals. Finally, $BEGIN(V)$ states that the positive literals of V are its first ones:

$$BEGIN(V) = (v_n \to v_{n-1}) \wedge \ldots \wedge (v_2 \to v_1)$$

Such a formula imposes that V and W have the same number of true literals, and that the set V has all the true atoms "at the beginning", so if W have three positive literals then $V = \{v_1, v_2, v_3\}$.

3.7 AGM operators

In previous sections we showed how we can reduce specific belief revision operators to circumscription and vice versa. Here we present a general methodology to transform any belief revision operator. The most general form of belief revision is given by the well-known postulates for revision (AGM postulates [Alchourrón *et al.*, 1985]). AGM postulates give eight basic properties that any belief revision operator should satisfy. Operators ($*_{AGM}$) satisfying the AGM postulates can be expressed as:

$$\mathcal{M}(K *_{AGM} A) = min(\mathcal{M}(A), \leq_K)$$

where $\leq_K$ is a transitive, reflexive and total relation based on K.

Note that not all the presented operators satisfy all AGM postulates. Updates and revisions defined by Borgida, Forbus, Satoh and Dalal are better generalized as reflexive and transitive orderings over pairs of models. The result of $K * A$ is $K \wedge A$ if consistent and $\mathcal{M}(K * A) = \{J | \exists I \in \mathcal{M}(K).\ \langle I, J \rangle \in min(\mathcal{M}(K) \times \mathcal{M}(A), \leq)\}$ otherwise.

In both cases, we must choose the minimal models of a formula w.r.t. a given ordering $\leq$ (or $\leq_K$). Any ordering over interpretations can be represented via a propositional formula $LEQ(\cdot, \cdot)$ (resp. $LEQ_K(\cdot, \cdot)$), such that $LEQ(X, Y)$ (resp. $LEQ_K(X, Y)$) is true iff $X \leq Y$ (resp. $X \leq_K Y$). Using this formula, AGM revision operators can be reduced to circumscription via:

$$K(X) *_{AGM} A(X) = w \wedge CIRC(T(X, Y, Z, \{w\}); X \cup Z \cup w, Y)$$

where $T(X, Y, Z, w)$ is defined as follows:

$$\begin{aligned} T(X, Y, Z, w) &= (X \neq Z) \wedge A(X) \wedge A(Y) \\ &\quad \wedge(\neg w = LESS_K(Y, X)) \end{aligned}$$

and $LESS_K(Y, X)$ represents the fact that $Y <_K X$. More precisely, $LESS_K(Y, X) \equiv LEQ_K(Y, X) \wedge \neg LEQ(X, Y)$. Note that K is missing in the circumscription, since it is implicit in LEQ_K. The formula $X \neq Z$ makes two models with different valuations over X incomparable. Models of $[(X \neq Z) \wedge A(X) \wedge A(Y) \wedge (\neg w = LEQ_K(Y, X))]$ assign $false$ to w iff X is not a minimal model of A w.r.t. $\leq_K$. Therefore, for every assignment to X not being a minimal model of A, the models of the circumscription make w false. Conjoining the result with w, we get rid of the non-minimal models.

The other generalization is even more complex. In fact, we must enforce that we choose the models M of A such that there exists a model N of K and for all models M' of A and models N' of K it does not hold that $M' \neq M$, $N' \neq N$ and $LEQ(M'N', MN)$.

We want to point out that these transformations are not necessarily polynomial. In fact, we do not know what is the size of the formula $LEQ(X, Y)$ w.r.t. the size of X and Y. It might very well be exponential.

4 Relations among belief revision operators

In Section 3 we found relations between circumscription and belief revision operators. Here we focus on relations among the various revision operators.

In particular, we show that Satoh's and Ginsberg's operators can be reduced one to the other and that Winslett's one can be reduced to both. Note that these operators belong to three different classes of operators, namely formula-based (Ginsberg), model-based with global proximity (Satoh) and model-based with local proximity (Winslett). Therefore, our results make evident the similarities between all these operators, pointing out, at the same time, their differences.

Ginsberg's operator can be reduced to Satoh's operator via:

$$
\begin{aligned}
A'(X,Y) &= A \wedge (y_1 \rightarrow f_1) \wedge \cdots \wedge (y_m \rightarrow f_m) \\
\{f_1,\ldots,f_m\} *_G A &= \{y_1,\ldots,y_m\} *_S A'(X,Y)
\end{aligned}
$$

The reverse reduction is:

$$
\begin{aligned}
A''(X,Y) &= K(Y) \wedge A(X) \wedge (W \rightarrow (X = Y)) \\
K(X) *_S A(X) &= \{w_1,\ldots,w_n\} *_G A''(X,Y)
\end{aligned}
$$

More complex, but still polynomial, is the reduction of Winslett's operator into Ginsberg's one:

$$K(X) *_W A(X) = (W \cup Y \cup Z) *_G F(X,Y,Z,W)$$

where $F(X,Y,Z,W)$ is:

$$F(X,Y,Z,W) = K(Y) \wedge A(X) \wedge (\neg Y \vee \neg Z) \wedge (W \rightarrow (X = Y))$$

Through the above reductions it is also possible to reduce $*_W$ to $*_S$.

5 Syntactically-restricted Knowledge Bases

In this section we focus on knowledge bases of a restricted syntactic form. Among the restricted cases, Horn knowledge bases are of particular interest for several reasons. First of all, since Horn clauses can represent causality relations, they are expressive enough to represent many real situations. Moreover, reasoning with Horn knowledge bases is significantly simpler than reasoning with general ones (see [Dowling and Gallier, 1984]) and also revising them is, in general, simpler than revising general ones (see [Eiter and Gottlob, 1992]).

While reductions from circumscription to belief revision preserve the syntactic form of the original theory, reductions from belief revision to circumscription do not preserve the syntactic form of the formulae. As an example, notice that the relation $X \neq Y$ cannot be expressed as an Horn formula.

As a consequence, it is easy to apply results on restricted cases of belief revision to circumscription, but the other way around is less likely to produce interesting results.

There are several reasons why the revision of Horn theories cannot be expressed as the circumscription of a Horn formula. First of all, results of Eiter and Gottlob show that reasoning with the revision of a Horn knowledge base is coNP-hard for all operators considered, while reasoning with Horn theories under circumscription is a polynomial task. As a consequence, reductions from belief revision to circumscription preserving the syntactic form cannot be done in polynomial time (assuming P $\neq$ NP).

Secondly, the result of revising a Horn knowledge base with a Horn formula might be a non-Horn formula. For example, the result of $\{a, b\} * (\neg a \vee \neg b)$ is $a \neq b$ for all operators, and $a \neq b$ cannot be expressed as an Horn formula. On the other hand, the circumscription of a Horn theory is an Horn theory.

6 Analysis and Discussion

In the previous sections we showed new relations relating belief revision operators and circumscription. These relations point out the close connections between the two fields. Many side benefits can be obtained from the established relations. In this section we want to point out the most important benefits obtained.

6.1 Compact Representation of NCIRC

In two recent papers [Cadoli *et al.*, 1995b, Cadoli *et al.*, 1995a] Cadoli, Donini and the present authors analyze the size of the explicit representation of circumscription and belief revision operators. More precisely, taking as an example belief revision, it is determined the size of the smallest propositional formula K_1 that is equivalent to $K * A$, where $*$ is one of the belief revision operators analyzed.

As it turns out, the size of the explicit representation of the result of revising a knowledge base is, in general, exponential w.r.t. $|K| + |A|$. Differences arise between the various operators. The result of revising a knowledge base using Dalal's revision operator admits a polynomial-sized explicit representation, if we allow new variables in the representation. More precisely, there exists a formula K_1 using the letters of K and A and possibly new ones, whose size is polynomial in $|K|+|A|$, s.t. , for any q using only variables of K and A we have that $K_1 \models q$ if and only if $K *_D A \models q$.

We show that $NCIRC(T; P, Z)$ always admits an explicit representation whose size is polynomial w.r.t. $|T|$, via the proof given for Dalal's belief revision operator. $NCIRC(T; P, Z)$ is the set of models of T with a least number of elements. Given T and P, we can compute the least number k of true letters, hence the explicit representation can be obtained conjoining T with a formula stating that at least k letters of the set P must be true. That is:

$$NCIRC(T; P, Z) = T \wedge ATLEAST(k, P)$$

where $ATLEAST(k, P)$ is a formula of size $O(n^3)$ that we do not show for the sake of brevity.

6.2 Computational Complexity Analysis

A valuable byproduct of the reductions presented in this work is the ability of importing complexity results obtained in one field into the other one. For example, in the general case, inference using the belief revision operators introduced by Satoh, Borgida and Winslett has the same complexity of inference under circumscription. While this result is not novel, it has been proven in [Eiter and Gottlob, 1993, Eiter and Gottlob, 1992], several other interesting results can be obtained. As an example, it is known that deciding whether a clause follows from the circumscription (with all letters minimized) of a theory composed of binary clauses (i.e. clauses with at most two literals) is a coNP-hard problem [Cadoli and Lenzerini, 1994]. We can use this result to prove that inference in the revision of a knowledge base composed of binary clauses is a coNP-hard problem for all operators except Dalal's one.

7 Conclusions

We have presented a complete analysis of the relations between belief revision operators on one hand and circumscription and its cardinality-based variant on the other hand. Furthermore, we have pointed out the many benefits that the established correlations can deliver to the analysis of both fields.

Our results greatly extends Winslett's results on transforming her revision operator into circumscription presented in [Winslett, 1989]. Even though Winslett's analysis could be further extended to deal with other operators, our results provide us with more direct and simple translations.

Acknowledgments

We want to thank Marco Cadoli and Francesco M. Donini for helpful discussions on the content of this paper.

References

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

[Borgida, 1985] A. Borgida. Language features for flexible handling of exceptions in information systems. *ACM Transactions on Database Systems*, 10:563–603, 1985.

[Cadoli and Lenzerini, 1994] M. Cadoli and M. Lenzerini. The complexity of propositional closed world reasoning and circumscription. *Journal of Computer and System Sciences*, 48:255–310, 1994.

[Cadoli *et al.*, 1995a] M. Cadoli, F. M. Donini, P. Liberatore, and M. Schaerf. The size of a revised knowledge base. In *Proceedings of the Fourteenth ACM SIGACT SIGMOD SIGART Symposium on Principles of Database Systems (PODS'95)*, pages 151–162, 1995. Extended version available as Technical Report DIS 34-96, Dipartimento di Informatica e Sistemistica, Università di Roma "La Sapienza", November 1996.

[Cadoli *et al.*, 1995b] M. Cadoli, F. M. Donini, and M. Schaerf. On compact representations of propositional circumscription. In *Proceedings of the Twelfth Symposium on Theoretical Aspects of Computer Science (STACS'95)*, pages 205–216, 1995. Extended version as RAP.14.95 DIS, Univ. of Roma "La Sapienza", July 1995.

[Dalal, 1988] M. Dalal. Investigations into a theory of knowledge base revision: Preliminary report. In *Proceedings of the Seventh National Conference on Artificial Intelligence (AAAI'88)*, pages 475–479, 1988.

[de Kleer and Konolige, 1989] J. de Kleer and K. Konolige. Eliminating the fixed predicates from a circumscription. *Artificial Intelligence*, 39:391–398, 1989.

[Dowling and Gallier, 1984] W. P. Dowling and J. H. Gallier. Linear-time algorithms for testing the satisfiability of propositional Horn formulae. *Journal of Logic Programming*, 1(3):267–284, 1984.

[Eiter and Gottlob, 1992] T. Eiter and G. Gottlob. On the complexity of propositional knowledge base revision, updates and counterfactuals. *Artificial Intelligence*, 57:227–270, 1992.

[Eiter and Gottlob, 1993] T. Eiter and G. Gottlob. Propositional circumscription and extended closed world reasoning are Π_2^p-complete. *Theoretical Computer Science*, 114:231–245, 1993.

[Forbus, 1989] K. D. Forbus. Introducing actions into qualitative simulation. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence (IJCAI'89)*, pages 1273–1278, 1989.

[Gärdenfors, 1988] P. Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. Bradford Books, MIT Press, Cambridge, MA, 1988.

[Ginsberg, 1986] M. L. Ginsberg. Conterfactuals. *Artificial Intelligence*, 30:35–79, 1986.

[Katsuno and Mendelzon, 1989] H. Katsuno and A. O. Mendelzon. A unified view of propositional knowledge base updates. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence (IJCAI'89)*, pages 1413–1419, 1989.

[Katsuno and Mendelzon, 1991] H. Katsuno and A. O. Mendelzon. On the difference between updating a knowledge base and revising it. In *Proceedings of the Second International Conference on the Principles of Knowledge Representation and Reasoning (KR'91)*, pages 387–394, 1991.

[Kharta and Lifschitz, 1994] G. N. Kharta and V. Lifschitz. Actions with indirect effects. In *Proceedings of the Fourth International Conference on the Principles of Knowledge Representation and Reasoning (KR'94)*, pages 341–350, 1994.

[Lifschitz, 1985] V. Lifschitz. Computing circumscription. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence (IJCAI'85)*, pages 121–127, 1985.

[McCarthy, 1980] J. McCarthy. Circumscription - A form of non-monotonic reasoning. *Artificial Intelligence*, 13:27–39, 1980.

[Reiter, 1980] R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 13:81–132, 1980.

[Satoh, 1988] K. Satoh. Nonmonotonic reasoning by minimal belief revision. In *Proceedings of the International Conference on Fifth Generation Computer Systems (FGCS'88)*, pages 455–462, 1988.

[Winslett, 1989] M. Winslett. Sometimes updates are circumscription. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence (IJCAI'89)*, pages 859–863, 1989.

[Winslett, 1990] M. Winslett. *Updating Logical Databases.* Cambridge University Press, 1990.