Distance Semantics for Relevance-Sensitive Belief Revision

Extended Abstract[1]

Pavlos Peppas,[2] *Samir Chopra,*[3] *and Norman Foo*[4]

1 Introduction

The *AGM postulates* [4], [1] are widely regarded to have captured much of what is involved in the process of rational belief revision. Recently however, Parikh [7] observed that despite their success, the AGM postulates are rather liberal in their treatment of the notion of *relevance*. More precisely, Parikh argues that during belief revision a rational agent does not change her entire belief corpus, but only the portion of it that is relevant to the new information. This intuition of *local change*, Parikh claims, is not fully captured by the AGM postulates. To remedy this shortcoming, Parikh introduced an additional axiom, named (P), as a supplement to the AGM postulates. Loosely speaking, axiom (P) says that when new information φ is received, only part of the initial belief set K will be affected; namely the part that shares common propositional variables with the minimal language of φ. Parikh's approach is also known as the *language splitting model*.

Our main goal in this paper is to provide possible-world semantics for axiom (P). In particular, we examine new constraints on systems of spheres and, building on Grove's representation result [5], we prove that in the presence of the AGM postulates, axiom (P) is sound and complete with respect to these new semantic constraints. What is particularly pleasing about our result is that the new constraints on systems of spheres are in fact not new at all; they essentially generalize a very natural condition that predates axiom (P) and has been motivated independently by Winslett in the context of Reasoning about Action [12].

In the course of formulating semantics for axiom (P) we observed that there are in fact two possible readings of this axiom, which we call the *strong* and the *weak* versions of (P). In the full text of this article, [10], both versions are studied. Herein however we confine the discussion to the weak version of axiom (P).

---

[1] This is an extended abstract of a paper by the same title published in the proceedings of the 9^{th} *International Conference on the Principles of Knowledge Representation and Reasoning*, AAAI Press, 2004.

[2] Dept of Business Administration, University of Patras, pavlos@upatras.gr

[3] Dept of Computer and Information Science, Brooklyn College of the City University of New York, Brooklyn, NY 11210, USA, schopra@sci.brooklyn.cuny.edu

[4] School of Computer Science and Engineering, University of New South Wales, Sydney, NSW 2052, Australia, norman@cse.unsw.edu.au

Dagstuhl Seminar Proceedings 07351
Formal Models of Belief Change in Rational Agents
http://drops.dagstuhl.de/opus/volltexte/2007/1215

2 Formal Preliminaries

We assume that the reader is familiar with the AGM postulates, as well as Grove's representation result that characterizes these postulates in terms of systems of spheres (see the full text for details).

Throughout this paper we work with a finite set of propositional variables $P = \{p_1, \ldots p_m\}$. We define $\mathcal{L}$ to be the propositional language generated from P (using the standard boolean connectives $\wedge, \vee, \rightarrow, \neg$ and the special symbols $\top, \bot$) and governed by classical propositional logic $\vdash$. A sentence $\varphi \in \mathcal{L}$ is *contingent* iff $\nvdash \varphi$ and $\nvdash \neg\varphi$. For a set of sentences Γ of $\mathcal{L}$, we denote by $Cn(\Gamma)$ the set of all logical consequences of Γ, i.e., $Cn(\Gamma) = \{\varphi \in \mathcal{L}\colon \Gamma \vdash \varphi\}$. We shall often write $Cn(\varphi_1, \varphi_2, \ldots, \varphi_n)$, for sentences $\varphi_1, \varphi_2, \ldots, \varphi_n$, as an abbreviation of $Cn(\{\varphi_1, \varphi_2, \ldots, \varphi_n\})$.

A theory T of $\mathcal{L}$ is any set of sentences of $\mathcal{L}$ closed under $\vdash$, i.e., $T = Cn(T)$. In this paper we focus only on *consistent* theories. Hence from now on, whenever the term "theory" appears unqualified, it is understood that it refers to a consistent theory. We denote the set of all consistent theories of $\mathcal{L}$ by $\mathcal{K}_{\mathcal{L}}$. A theory T of $\mathcal{L}$ is *complete* iff for all sentences $\varphi \in \mathcal{L}$, $\varphi \in T$ or $\neg\varphi \in T$. We denote the set of all consistent complete theories of $\mathcal{L}$ by $\mathcal{M}_{\mathcal{L}}$. In the context of systems of spheres, consistent complete theories essentially play the role of possible worlds. Following this convention, in the rest of the article we use the terms "possible world" (or simply "world") and "consistent complete theory" interchangeably. For a set of sentences Γ of $\mathcal{L}$, $[\Gamma]$ denotes the set of all consistent complete theories of $\mathcal{L}$ that contain Γ. Often we use the notation $[\varphi]$ for a sentence $\varphi \in \mathcal{L}$, as an abbreviation of $[\{\varphi\}]$. For a theory T and a set of sentences Γ of $\mathcal{L}$, we denote by $T + \Gamma$ the closure under $\vdash$ of $T \cup \Gamma$, i.e., $T + \Gamma = Cn(T \cup \Gamma)$. For a sentence $\varphi \in \mathcal{L}$ we often write $T + \varphi$ as an abbreviation of $T + \{\varphi\}$.

In the course of this paper, we often consider *sublanguages* of $\mathcal{L}$. Let P' be a subset of the set of propositional variables P . By $\mathcal{L}^{P'}$ we denote the sublanguage of $\mathcal{L}$ defined over P'. In the limiting case where P' is empty, we take $\mathcal{L}^{P'}$ to be the language generated by $\top, \bot$ and the boolean connectives. For a sublanguage $\mathcal{L}'$ of $\mathcal{L}$ defined over a subset P' of P, by $\overline{\mathcal{L}'}$ we denote the sublanguage defined over the propositional variables in the complement of P' i.e., $\overline{\mathcal{L}'} = \mathcal{L}^{(P-P')}$. For a sentence χ of $\mathcal{L}$, by $\mathcal{L}_\chi$ we denote the *minimal* sublanguage of $\mathcal{L}$ within which χ can be expressed (i.e., $\mathcal{L}_\chi$ contains a sentence that is logically equivalent to χ, and moreover no *proper* sublanguage of $\mathcal{L}_\chi$ contains such a sentence).[5] Finally we note that in the forthcoming discussion, we often project operations defined earlier for the entire language $\mathcal{L}$, to one of its sublanguages $\mathcal{L}'$. When this happens, all notation will be subscripted by the sublanguage $\mathcal{L}'$. For example, for a set of sentences $\Gamma \subset \mathcal{L}'$, the term $Cn_{\mathcal{L}'}(\Gamma)$ denotes the logical closure of Γ *in* $\mathcal{L}'$. Similarly, $[\Gamma]_{\mathcal{L}'}$ denotes the set of all maximally consistent supersets of Γ *in* $\mathcal{L}'$. When no subscript is present, it is understood that the operation is relevant to the original language $\mathcal{L}$.

[5] It is not hard to verify that for every χ, $\mathcal{L}_\chi$ is unique – see [7] for details.

3 Relevance-Sensitive Belief Revision

When revising a theory T by a sentence φ it seems plausible to assume that only the beliefs that are *relevant* to φ should be affected, while the rest of the belief corpus is unchanged. For example, an agent that is revising her beliefs about planetary motion is unlikely to revise her beliefs about Malaysian politics. This simple intuition is not fully captured in the AGM paradigm. To see this consider the trivial revision function $*_t$ defined below:

$$T *_t \varphi = \begin{cases} T + \varphi & \text{if } \varphi \text{ is consistent with } T \\ Cn(\varphi) & \text{otherwise} \end{cases}$$

It is not hard to verify that $*_t$ satisfies all the AGM postulates, and yet whenever $\neg\varphi \in T$, it has the rather counter-intuitive effect of throwing away *all* beliefs in T that are not consequences of φ, regardless of whether these beliefs are related to φ or not.

In order to block revision functions like $*_t$ Parikh introduced in [7] a new axiom, named (P), as a supplement to the AGM postulates. The main intuition that axiom (P) aims to capture is that an agent's beliefs can be subdivided into disjoint compartments, referring to different subject matters, and that when revising, the agent modifies only the compartment(s) affected by the new information:

(P) If $T = Cn(\chi, \psi)$ where χ, ψ are sentences of disjoint sublanguages $\mathcal{L}_1$, $\mathcal{L}_2$ respectively, and $\varphi \in \mathcal{L}_1$, then $T * \varphi = (Cn_{\mathcal{L}_1}(\chi) \circ \varphi) + \psi$, where $\circ$ is a revision operator of the sublanguage $\mathcal{L}_1$.

It was shown in [7] that (P) is consistent with the first six AGM postulates (known as the *basic* AGM postulates). The results presented later in this paper entail that (P) is in fact consistent with all eight AGM postulates (usually numbered (K*1) - (K*8)).

4 Two Readings of Axiom (P)

Before proceeding with the formulation of semantics for axiom (P), it is worth taking a closer look at it.

Consider two sentences $\chi, \psi \in \mathcal{L}$, such that $\mathcal{L}_\chi \cap \mathcal{L}_\psi = \emptyset$, and let T be the theory $T = Cn(\{\chi, \psi\})$. Moreover, let φ be any sentence in $\mathcal{L}_\chi$. According to axiom (P), anything outside $\mathcal{L}_\chi$ will not be affected by the revision of T by φ. This however is only one side of axiom (P). The other side concerns the part of the theory T that *is* related to φ, which according to axiom (P) will change to $Cn_{\mathcal{L}_\chi}(\chi) \circ \varphi$, where $\circ$ is a revision function defined over the sublanguage $\mathcal{L}_\chi$. It is this second side of axiom (P) that needs closer examination.

Axiom (P) is open to two different interpretations. According to the first reading, which we call the *weak* version of axiom (P), the revision function $\circ$ that modifies the relevant part of T – call it the *local* revision function – may *vary* from theory to theory, even when the relevant part $Cn(\chi)$ stays the same. To give a concrete example, let a, b, c be propositional variables, let T be the theory $T = Cn(a \wedge b, c)$, and let T' be the theory $T' = Cn(a \wedge b, \neg c)$. Denote by $\mathcal{L}_1$ the sublanguage defined over $\{a, b\}$ and by $\mathcal{L}_2$ the sublanguage defined over $\{c\}$. Moreover, let φ be the sentence $\varphi = \neg a \vee \neg b$. The part of T and T' that is relevant to φ (in the sense of the language-splitting model) is the same for both theories, namely $Cn(a \wedge b)$. Nevertheless, according to the *weak* version of axiom (P), the local revision operators $\circ$ and $\circ'$ that modify the two *identical* relevant parts of T and T' respectively, may very well *differ*. For example, it could be the case that $Cn_{\mathcal{L}_1}(a \wedge b) \circ (\neg a \vee \neg b) = Cn_{\mathcal{L}_1}(\neg a \wedge b)$, and $Cn_{\mathcal{L}_1}(a \wedge b) \circ' (\neg a \vee \neg b)$ $= Cn_{\mathcal{L}_1}(a \wedge \neg b)$, from which it follows that $T * \varphi = Cn(\neg a, b, c)$, and $T' * \varphi =$ $Cn(a, \neg b, \neg c)$. In other words, the weak version of axiom (P) allows the local revision function to be *context-sensitive*. In the scenario described above, the presence of c in T leads to a local revision function $\circ$ for $Cn_{\mathcal{L}_1}(a \wedge b)$ that produces $Cn_{\mathcal{L}_1}(\neg a \wedge b)$ as the result of revising by $\neg a \vee \neg b$; on the other hand, the presence of $\neg c$ in T', induces a local revision function $\circ'$ for $Cn_{\mathcal{L}_1}(a \wedge b)$ that produces $Cn_{\mathcal{L}_1}(a \wedge \neg b)$ for the same input. Therefore, while c (or $\neg c$) remains unaffected during the (global) revision by $\neg a \vee \neg b$ (since it is not relevant to the new information), its presence influences the way that the relevant part of the theory is modified.

To prevent such an influence we need to resort to the *strong* version of axiom (P) which makes the local revision function $\circ$ *context-independent.* According to the strong interpretation of (P), for any two theories $T = Cn(\chi, \psi)$ and $T' =$ $Cn(\chi, \psi')$, such that $\mathcal{L}_\chi \cap \mathcal{L}_\psi = \mathcal{L}_\chi \cap \mathcal{L}_{\psi'} = \emptyset$, there exists a *single* local revision function $\circ$ such that $T * \varphi = (Cn_{\mathcal{L}_\chi}(\chi) \circ \varphi) + \psi$ and $T' * \varphi = (Cn_{\mathcal{L}_\chi}(\chi) \circ \varphi) + \psi'$, for any $\varphi \in \mathcal{L}_\chi$.

To make explicit the two possible reading of axiom (P) we rephrase it in terms of conditions (R1) and (R2) below:

(R1) If $T = Cn(\chi, \psi)$, $\mathcal{L}_\chi \cap \mathcal{L}_\psi = \emptyset$, and $\varphi \in \mathcal{L}_\chi$, then $(T * \varphi) \cap \overline{\mathcal{L}_\chi} = T \cap \overline{\mathcal{L}_\chi}$.

(R2) If $T = Cn(\chi, \psi)$, $\mathcal{L}_\chi \cap \mathcal{L}_\psi = \emptyset$, and $\varphi \in \mathcal{L}_\chi$, then $(T * \varphi) \cap \mathcal{L}_\chi =$ $(Cn(\chi) * \varphi) \cap \mathcal{L}_\chi$.

Condition (R1) corresponds to the weak version of axiom (P). It simply states that the part of the theory T that is *not related* to the new information φ is not affected by the revision. Adding (R2) to condition (R1) gives us the strong version of axiom (P). To see this, consider a revision function $*$ (which defines a revision policy for *all* the theories of $\mathcal{L}$), and let $T = Cn(\chi, \psi)$ and T' $= Cn(\chi, \psi')$ be two theories such that $\mathcal{L}_\chi \cap \mathcal{L}_\psi = \mathcal{L}_\chi \cap \mathcal{L}_{\psi'} = \emptyset$. Consider now any sentence $\varphi \in \mathcal{L}_\chi$. The relevant part to φ of T and T' is in both cases the same. Then, according to (R2), the way that this relevant part is modified in both T and T' is also the same; namely, as dictated by the revision function $*$ itself when applied to $Cn(\chi)$ (once again, notice that $*$ is defined for all theories, including T, T', *and* $Cn(\chi)$).

Notice that the strong version of axiom (P) brings about a new feature in the picture of classical AGM revision: it makes associations between the revision policies of *different* theories; none of the AGM postulates have this property (they all refer to a *single* theory T).

As already mentioned in the introduction, herein we shall discuss only the weak version of axiom (P), that is condition (R1). The full text however also covers condition (R2) .

5 The Special Case of Complete Theories

In formulating semantics for condition (R1) we start with the special case of consistent complete theories as belief sets. This will help us build up our intuition and better digest the general case that follows.

Let T be a consistent complete theory, and let S_T be a system of spheres centered on $[T]$. The intended reading of S_T is that it represents *comparative similarity* between possible worlds i.e., the further away a world is from the center of S_T, the less similar it is to $[T]$.[6] Yet the definition of a system of spheres gives no indication about *how* similarity between worlds should be measured. In [9] a specific criterion of similarity is considered, originally introduced in the context of *Reasoning about Action* with Winslett's *Possible Models Approach* (PMA) [12]. This criterion, called *PMA's criterion of similarity*, measures "distance" between worlds based on propositional variables. In particular, let r, r' be any two possible worlds of $\mathcal{L}$. By $\mathit{Diff}(r, r')$ we denote the set of propositional variables that have different truth values in the two worlds i.e., $\mathit{Diff}(r, r') = \{p_i \in P : p_i \in r$ and $p_i \notin r'\} \cup \{p_j \in P : p_j \notin r$ and $p_j \in r'\}$. A system of spheres S_T is a *PMA system of spheres* iff it satisfies the following condition [9] (throughout this paper, the symbols r and r' always represent consistent complete theories):

(PS) If $\mathit{Diff}(T, r) \subset \mathit{Diff}(T, r')$ then there is a sphere $V \in S_T$ that contains r but not r'.

According to condition (PS), the less a world r differs from the initial belief set T in propositional variables, the closer it is to the center of S_T.[7] It turns out that, in the special case of consistent complete belief sets, condition (PS) is precisely the constraint that is needed to bring about condition (R1):

Theorem 1. *Let $*$ be a revision function satisfying the AGM postulates, T a consistent complete theory of $\mathcal{L}$, and S_T the system of spheres centered on $[T]$, corresponding to $*$ by means of (S*). Then $*$ satisfies (R1) at T iff S_T satisfies (PS).*

As already mentioned in the introduction, what is quite appealing about Theorem 1 is that it characterizes (R1), not in terms of some technical non-intuitive condition, but rather by a natural constraint on similarity between

[6] Perhaps "comparative plausibility" would have been a better term in the present context. However we shall tolerate this slight abuse of terminology mainly to comply with [9].

[7] Notice that condition (PS) places no constraints on the relative order of worlds that are *Diff*-incomparable.

possible worlds, that in fact predates (R1) and was motivated independently in a different context [12]. Moreover, as we will show in the next section, the essence of this characterization of (R1) in terms of constraints on similarity, carries over into the general case of incomplete belief sets (albeit with some modifications).

6 The General Case

To elevate Theorem 1 to the general case, we first need to extend the definition of *Diff* to cover comparisons between a world r and an arbitrary, possibly *incomplete*, theory T. The generalization of *Diff* that we shall use herein takes into account the notion of a T*-splitting* introduced by Parikh in his language-splitting model [7].

Let T be a theory of $\mathcal{L}$ and $P_1, P_2, \ldots, P_n$ a partition of the set P of all propositional variables in $\mathcal{L}$. We say that $\{P_1, P_2, \ldots, P_n\}$ is a T*-splitting* iff there exist sentences $\varphi_1 \in \mathcal{L}^{P_1}$, $\varphi_2 \in \mathcal{L}^{P_2}$, … , $\varphi_n \in \mathcal{L}^{P_n}$, such that $T = Cn(\varphi_1, \varphi_2, \ldots, \varphi_n)$. Parikh has shown in [7] that for every theory T there is a unique *finest* T-splitting, i.e. one which refines[8] every other T-splitting. We shall denote the finest T-splitting of T by $\mathcal{F}(T)$. Using the notion of a finest T-splitting, we define the difference between a (possibly incomplete) theory T of $\mathcal{L}$ and a world r as follows: $Diff(T, r) = \bigcup\{P' \in \mathcal{F}(T) : \text{for some } \varphi \in \mathcal{L}^{P'}, T \vdash \varphi \text{ and } r \vdash \neg\varphi\}$. It is not hard to verify that in the special case of a consistent complete theory T, the above definition of *Diff* collapses to the one given in the previous section.

Having generalized *Diff*, condition (PS) can be applied verbatim to *any* system of spheres S_T, including the ones related to incomplete belief sets T. It turns out however that in the general case (PS) is too strong (i.e. there are revision functions satisfying (R1) whose corresponding system of spheres fails to satisfy (PS) – see the full text, [10], for details).

Working towards an appropriate weakening of (PS) we shall need some additional notation and definitions.

Consider a theory T and let r be a world not compatible with T i.e., $r \notin [T]$. Clearly $Diff(T, r) \neq \emptyset$. Is there another world r' that differs from T on exactly the same propositional variables, i.e., $Diff(T, r) = Diff(T, r')$? If T is complete, the answer is obviously "no": for any set of propositional variables P', there can only be *one* world r such that $Diff(T, r) = P'$. If however T is *incomplete* (i.e., $[T]$ contains more than one world), this is no longer the case. For example, suppose that $T = Cn(a \leftrightarrow b, c \leftrightarrow d)$ – where a, b, c, d, are the propositional variables of the language – and let r, r' be the possible worlds $r = Cn(\{\neg a, b, c, d\})$, and $r' = Cn(\{a, \neg b, c, d\})$. It is not hard to see that, although r and r' are different, $Diff(T, r) = Diff(T, r') = \{a, b\}$. The two worlds r and r', have also another thing in common: they agree on the propositional variables *outside* $Diff(T, r)$. We call such worlds *external* T*-duals*:

[8] A partition Z refines another partition Z', if for every element of Z there is a superset of it in Z'.

Definition 1. *Let r, r' be possible worlds, and let T be a theory of $\mathcal{L}$. The worlds r and r' are* external T-duals *iff $\mathit{Diff}(T,r) = \mathit{Diff}(T,r')$ and $r \cap (P - \mathit{Diff}(T,r)) = r' \cap (P - \mathit{Diff}(T,r'))$.*

Multiple T-duals (external and *internal* ones as we will see later) add more structure to a system of spheres, and render condition (PS) too strong for the general case. The possibility of placing external T-duals in *different* spheres, opens up new ways of ordering worlds that still induce relevance-sensitive revision functions without however submitting entirely to the demands of (PS).

Let us elaborate on this point. Consider a system of spheres S_T centered on the theory T, and let r, r' be any two worlds such that $\mathit{Diff}(T,r) \subset \mathit{Diff}(T,r')$. Theorem 1 tells us that in the special case of complete theories, to ensure local change (alias, condition (R1)) the world r should be placed (strictly) closer to the center $[T]$ of S_T than r'. In the general case however, and with the aid of external T-duals, one can perhaps afford to be a bit more liberal about the location of r; perhaps all that is needed is that at least one external T-dual r'' of r (and not necessarily r itself) be closer to $[T]$ than r'. It turns out that, in fact, this is pretty much the case, expect that the world r'' "covering" for r (in relation to r') is not just any external T-dual of r but a very specific one: it is the external T-dual of r that agrees with r' on all literals in $\mathit{Diff}(T,r)$. We shall call this external T-dual of r, the *r'-cover for r at T*, and we shall denote it by $\vartheta_T(r,r')$:

Definition 2. *Let T be a theory of $\mathcal{L}$, let r, r' be two possible worlds such that $\mathit{Diff}(T,r) \subset \mathit{Diff}(T,r')$, and let r'' be an external T-dual of r. The world r'' is the r'-cover for r at T iff $r'' \cap \mathit{Diff}(T,r) = r' \cap \mathit{Diff}(T,r)$. We shall denote the r'-cover for r at T by $\vartheta_T(r,r')$.*

Equipped with the notion of "covering" we propose the following weaker version of (PS):

(Q1) If $\mathit{Diff}(T,r) \subset \mathit{Diff}(T,r')$ then there is a sphere $V \in S_T$ that contains $\vartheta_T(r,r')$ but not r'.

It is not hard to verify that (PS) entails (Q1), and that (Q1) collapses to (PS) when the initial belief set T is complete. Moreover, (Q1) is *strictly* weaker than (PS) (see the full text for details). In fact it is too weak for our purpose. Condition (Q1) needs to be coupled with another condition, named (Q2), to produce the intended correspondence with (R1) in the general case. This second condition uses the notion of an *internal T-dual* defined below:

Definition 3. *Let r, r' be possible worlds, and let T be a theory of $\mathcal{L}$. The worlds r and r' are* internal T-duals *iff $\mathit{Diff}(T,r) = \mathit{Diff}(T,r')$, and $r \cap \mathit{Diff}(T,r) = r' \cap \mathit{Diff}(T,r')$.*

Condition (Q2) below essentially places any two internal T-duals at the same distance from the center $[T]$ of a system of spheres S_T:

(Q2) If r and r' are internal T-duals, then they belong to the same spheres in S_T; i.e., for any sphere $V \in S_T$, $r \in V$ iff $r' \in V$.

Notice that in the special case that T is complete, no world r has internal or external T-duals (other than itself). Consequently, in that case, (Q1) reduces to (PS), while (Q2) degenerates to a vacuous condition.

The promised correspondence between (R1) and the two conditions (Q1) - (Q2) is given by the theorem below:

Theorem 2. *Let $*$ be a revision function satisfying the AGM postulates, T a consistent theory of $\mathcal{L}$, and S_T a system of spheres centered on $[T]$, that corresponds to $*$ by means of (S*). Then $*$ satisfies (R1) at T iff S_T satisfies (Q1) - (Q2).*

7 Conclusion

The main contribution of this paper is a system-of-spheres characterization of Parikh's axiom (P). What is quite appealing about this result is that the semantic conditions (Q1) - (Q2) that characterize (the weak version of) axiom (P) are quite natural constraints on similarity between possible worlds. In fact, conditions (Q1) - (Q2) essentially generalize a measure of similarity that predates axiom (P), and was motivated independently in the context of Reasoning about Action by Winslett. This intuitive nature of the semantics is more evident in the special case of consistent complete belief sets. An interesting by-product of our study is the identification of the two possible readings of axiom (P), both of which are plausible depending on the context.

It should be noted that apart from Winslett, other authors have also made specific proposals for measuring distance between possible worlds (see for example, [2], [3], and [11]).[9] It would be a worthwhile exercise to investigate whether any of these measures of distance also yield some kind of "local change effect" for their associated revision functions.

References

1. Alchourron, C., Gardernfors, P., and Makinson, D. "On the logic of theory change: Partial meet functions for contraction and revision", *Journal of Symbolic Logic* 50:510–530, 1985.
2. Borgida, A. "Language features and flexible handling of exceptions in information systems", *ACM Transactions on Database Systems* 10:563–603, 1985.
3. Dalal, M. "Investigations into a theory of knowledge base revisions: Preliminary report", In *Proceedings of the 7th AAAI Conference*, 475–479, 1988.
4. P. Gardenfors. "Epistemic Importance and Minimal Changes of Belief", *Australasian Journal of Philosophy*, vol 62, pp 136-157, 1984.
5. Grove, A. "Two modellings for theory change", *Journal of Philosophical Logic* 17:157–170, 1988.

[9] For a more general study on the notion of distance in belief revision, refer to [6].

6. Lehmann, D., Magidor, M., and Schlechta, K. "Distance semantics for belief revision", *Journal of Symbolic Logic* 66(1):295–317, 2001.
7. Parikh, R. "Beliefs, belief revision, and splitting languages", In Lawrence Moss, J. G., and de Rijke, M., eds., *Logic, Language, and Computation Volume 2 - CSLI Lecture Notes No. 96*. CSLI Publications, 266–268, 1999.
8. Peppas, P., and Wobcke, W. 1992. "On the use of epistemic entrenchment in reasoning about action", In *Proceedings of the 10th European Conference on Artificial Intelligence*, 403–407. John Wiles and Sons.
9. P. Peppas, N. Foo, and A. Nayak "Measuring similarity in belief revision", *Journal of Logic and Computation* 10(4), 2000.
10. P. Peppas, S. Chopra, and N. Foo, "Distance Semantics for Relevance-Sensitive Belief Revision", *Proceedings of the 9th International Conference on the Principles of Knowledge Representation and Reasoning*, Whistler, Canada, June 2004.
11. Satoh, K. "Nonmonotonic reasoning by minimal belief revision", In *Proceedings of the International Conference on Fifth Generation Computer Systems*, 455–462, 1988.
12. Winslett, M. Reasoning about action using a possible models approach. In *Proceedings of the 7th AAAI Conference*, 89–93, 1988.