Retrieval Situations and Belief Change

David E. Losada
AILab. Department of Computer Science
University of A Corunna, Spain
losada@dc.fi.udc.es

Alvaro Barreiro
AILab. Department of Computer Science
University of A Corunna, Spain
barreiro@dc.fi.udc.es

Abstract

Situational aspects are very helpful to decide relevance but they have often been left aside by Information Retrieval models. The standard logical approach to Information Retrieval, based on deciding relevance with the entailment $d \models q$, where d and q are logical representations of a document and a query respectively, does not consider retrieval situations. In this work we propose to introduce these aspects through a revision process. The relevance test is generalized to be $(S \circ d) \models q$, where $\circ$ is a Belief Revision operator and S is the logical representation of a retrieval situation. Besides, we analyze the basic properties of different belief change methods and their adequacy to the modelization of $S \circ d$. We also study the close relation between our approach and other approaches dealing with counterfactuals.

1. Introduction

It has been widely recognized that *topicality* should not be the only criterion for a relevance judgment [20, 16, 15]. An Information Retrieval (IR) system should decide relevance using more information than a simple matching of topics between document and query representations. Other factors such as semantic relations between terms, user's model and pragmatics of language should also be taken into account. For instance, the user model can contain user knowledge and his/her intentions to improve retrieval precision. With *retrieval situation* we refer to all the aspects apart from topicality that have influence on user's relevance judgment. The following example illustrates the importance of retrieval situations. Imagine two users looking for documents on "information retrieval". One of them is an IR researcher and the other is a novice student. Clearly, a document about "relevance feedback" should not be judged with the same level of relevance for both. The novice user hardly knows that relevance feedback has something to do with IR. However, systems do not have appropriate tools to deal with situational factors. As a matter of fact, introducing retrieval situations in the Boolean Model may lead to counterintuitive results. This is because retrieval situations can contradict document's content, making the document to be regarded as relevant to any query.

The necessity of a conditional logic for IR was first made explicit in [19, 20]. The idea of minimality, which is present in conditional logics, was promising for the application to IR. In [16], Nie and other researchers proposed counterfactual conditional logic to deal with retrieval situations. The capability of this kind of logics to handle contradictions was used to bring in retrieval situations. Given d and q, the logical representations of a document and an information need respectively, and S, a logical representation of a retrieval situation, relevance should be based on evaluating the conditional $d > q$ with respect to S. The technique to evaluate $d > q$ with respect to S is based on revising the situation S such that d becomes true in the revised situation $S^{'}$. Then, the query q is evaluated in the revised situation $S^{'}$. That way, the problem of contradicting information is avoided. The idea of *minimal change* is taken into account during the revision of the situation: the parts of S that are consistent with d are maintained in $S^{'}$ while the inconsistent ones are modified minimally in order to keep consistence.

F. P. Ramsey proposed a method for evaluating a conditional proposition [17]. Ramsey's method was later formulated within the Theory of Epistemic Change [9]. That formulation clearly states the close relation between the evaluation of conditionals and a belief change process. In this work we discuss the adequacy to IR of the implementation of the relevance test with Belief Revision, following directly the Ramsey test. Formally, a retrieval situation is modeled as a set of logical formulas contained in a logical theory S. A document and an information need are represented by logical formulas d and q respectively. The document is considered relevant to the information need in that situation iff $(S \circ d) \models q$ holds. However, the direct application of the Ramsey test leads to the well known *Grdenfors Triviality Result* [9]. This paper studies whether this triviality result poses any problem for IR and, in the case, which of the solutions proposed in the literature could be appropriate. Besides that, we present an analysis of the concrete change semantics that should be used for changing the retrieval situation with the document.

Finally, as logical entailment is too strict to decide relevance, we propose to use distances between logical interpretations to measure the uncertainty of $(S \circ d) \models q$, following the techniques presented in [14] which allows us to rank relevant documents.

The rest of this paper is organized as follows. Section 2 explores the connection between conditionals and change semantics, which is formalized by the Ramsey test. That section also discusses Grdenfors' Triviality result. Section 3 analyzes the kind of change semantics that should be used for IR and section 4 presents some considerations about partiality and some remarks about other IR approaches that are also free from triviality. The paper ends with some conclusions.

2. The Ramsey test and Grdenfors triviality result

Counterfactuals are conditional statements of the form "if p were true, q would hold", where p is assumed to be false in the current state of affairs. They are also called conditionals for short. Along this work we will use both names interchangeably. Ramsey anticipated a test for evaluating conditionals [17] that was summarized by Grdenfors as follows [9]: in order to find out whether a conditional proposition is acceptable in a given state of belief, one first adds the antecedent of the conditional to the given stock of beliefs. Second, if the antecedent together with the former stocks of beliefs leads to a contradiction, then one makes some adjustments, as small as possible without modifying the antecedent, such that consistency is maintained. Finally, one considers whether or not the consequent of the conditional is accepted in this adjusted state of belief.

The Ramsey test has received renewed interest since Alchourrn, Grdenfors and Makinson (AGM) developed the theory of epistemic change [1]. In [9] the Ramsey test was formulated within the AGM framework. The notion of change needed by the Ramsey test is formalized by an operation of revision:

> Evaluating a counterfactual $a > b$ in a given knowledge base T is equivalent to test whether b is a logical consequence of $T \circ a$, where $T \circ a$ represents the knowledge base T revised with the formula a by the revision operator $\circ$.

Despite the intuitive aspect of the Ramsey test, which connects the semantics of Belief Revision and the evaluation of conditional implications, its strict application leads to the *Grdenfors Triviality Result* [9]. Roughly speaking, this result claims that there are no significant logic being compatible with the Ramsey test.

The central point of Grdenfors' result stands on the inconsistency between the *Preservation Principle*, which is a fundamental criterion in all revision methods, and the *Monotonicity Principle*, which follows directly from the Ramsey test (see proof in [9]). The former principle formalizes the notion of information economy, i.e. we do not want to give up beliefs unnecessarily:

> **Preservation Principle:** If $\psi \wedge \mu$ is satisfiable then $\psi \circ \mu \equiv \psi \wedge \mu$.

Monotonicity formalizes the notion that if a knowledge base is a consequence of another knowledge base, the same holds for their respective revisions with any formula A:

> **Monotonicity Principle:** Given K and $K^{'}$ two knowledge bases such that $K \models K^{'}$, then $K \circ A \models K^{'} \circ A$

If both principles hold the associated logic is *trivial*. A logic is trivial if we cannot articulate four formulas, such that three of them are pairwise inconsistent and the fourth one is consistent with each of the former three. Formally, a logic L is said to be *trivial* if there are not four sentences ψ, χ, μ and ϕ in the language for L, such that the sentences $\psi \wedge \chi$, $\psi \wedge \mu$ and $\chi \wedge \mu$ are inconsistent in L, and the sentences $\phi \wedge \psi$, $\phi \wedge \chi$ and $\phi \wedge \mu$ are consistent in L. Otherwise the logic L is *non-trivial*

Now, we discuss whether a trivial logic is suitable for IR. A retrieval situation is represented as a set S of propositional formulas. A typical formula belonging to S could be a material implication formalizing a semantic relation between two terms. A document and an information need are represented by logical formulas d and q respectively. To decide relevance we have to check if the counterfactual $d > q$ is true in the situation S. Applying directly the Ramsey test, we have to check whether $(S \circ d) \models q$ holds. However, Grdenfors Triviality result holds, imposing unacceptable constraints. The logic would not allow us to represent three pairwise disjoint documents that are consistent with a query.

Example: Let us consider three pairwise inconsistent documents represented by the formulas d_1, d_2 and d_3 respectively and a query represented by the formula q:

$$d_1 = databases \wedge \neg programming \wedge operating\ systems$$
$$d_2 = \neg databases \wedge programming \wedge operating\ systems$$
$$d_3 = databases \wedge programming \wedge operating\ systems$$

$$q = operating\ systems$$

This is a totally feasible case in IR and, clearly, the query is consistent with each document and a trivial logic would not allow us to represent the four items. As a consequence of the triviality result, a logic with the Ramsey rule and with a change operator satisfying the preservation principle is trivial. Then, that kind of logic cannot represent the above sentences and, as a result, its use is unacceptable for IR.

2.1. Avoiding triviality

Since Grdenfors presented his Triviality Result, several researchers have studied different methods to avoid it. There have been two fundamental lines of work in this sense. Approaches on the first line drop the Preservation Principle. Although this principle is commonly accepted by Belief Revision operators, it is not satisfied by other change methods, such as Belief Update mechanisms [11]. However, the decision between update or revision should be made according to the correspondence between their change semantics and the expected behaviour in the domain of application. In section 3 we show that the semantics of update is not suitable for our purposes. Other researchers have followed the line of defining weaker forms of the Ramsey test. The key objective of this policy is to avoid monotonicity. The dependency between conditionals and change semantics is captured in a less strict way:

> **Ramsey test:** $A > B$ is *accepted* in K iff $K \circ A \models B$.

The proof that monotonicity is a consequence of the Ramsey test [9] stands on the fact that the conditional connective $>$ belongs to the language. In contrast, the previous relaxed form of the Ramsey test blocks that proof. A problem of this more relaxed formulation is that it cannot represent nested conditionals. This means that we cannot represent conditionals with the form "if (if A then B) then C". For IR this implies that we cannot represent documents using conditional sentences. However, we are not interested here in conditionals for that representational task. In fact, in [16] the counterfactual implication was proposed to make the relevance test for IR while other aspects such as relationships between terms were modeled with material implications. Summarizing, we will consider counterfactuals at the meta-level, to which the triviality result does not apply. As the relevance test can be done using a revision process at the meta-level, we can follow the previous formulation of the Ramsey test such that we can decide relevance using a logic without conditionals. This goes in the line of [3], where the authors evaluate a conditional sentence without explicitly defining the conditional operator.

3. Choosing a change semantics

The key choice now is the kind of semantics that should be used to change the retrieval situation S with the document d. This section studies different change semantics present in the literature.

3.1. Belief Revision rationality postulates

The properties that a *reasonable* Belief Revision operator should have were formalized as a set of rationality postulates by Grdenfors, Alchourrn and Makinson [2, 8, 1]. These postulates were formulated in a very general way and did not assume any concrete representation. In fact, they used deductively closed sets of sentences in some unspecified language. These sets are known as belief sets. The AGM postulates were later adapted to knowledge bases (finite sets of propositional sentences) by Katsuno and Mendelzon (KM) [12]. The set of KM postulates is depicted in fig.1. Belief Revision rationality postulates have been widely recognized as a paradigm and it is interesting to analyze them trying to identify the implications that they impose when applied to IR. In this sense, it is important to recall that we want to check whether $(S \circ d) \models q$ holds. Thus, we revise the retrieval situation S with the document d. The first postulate, sometimes called the *success* postulate, states that after the revision, the document should be a logical consequence of the revised situation. This is a basic postulate in Belief Revision and formalizes the notion that the knowledge base and the new information have different status because after the revision the new information has to be accepted no matter what happens with the old information. From the IR perspective, the important consideration is whether the document should prevail over the retrieval situation. Obviously, if we want to implement the counterfactual proposed in [16] we have to follow the Ramsey test whose result forces us to make the revision $S \circ d$ and not $d \circ S$. However one could think that the proposal could have been to test whether $S > q$ in d. In any case, we think it is interesting to discuss both options. If $s \wedge d$ is satisfiable both tests produce the same result. This follows directly from R2. When there are contradictions between the representation of a document and the representation of a situation, one of the representations is changed to keep consistence. Let us analyze the choice with an example. A user can think that "medicine" has nothing to do with "computer science". This can be modeled as the belief `medicine` $\rightarrow \neg$ `computer science` belonging to S. An hypothetical document dealing with both issues could be represented as `medicine` $\wedge$ `computer science` $\wedge \ldots$ and would come into contradiction with S. The revision $S \circ d$ entails both medicine and computer science (from R1), so that a query about one of them would retrieve the

$(R1) \psi \circ \mu$ implies μ

$(R2)$ If $\psi \wedge \mu$ is satisfiable, then $\psi \circ \mu \equiv \psi \wedge \mu$

$(R3)$ If μ is satisfiable, then $\psi \circ \mu$ is also satisfiable

$(R4)$ If $\psi_1 \equiv \psi_2$ and $\mu_1 \equiv \mu_2$ then $\psi_1 \circ \mu_1 \equiv \psi_2 \circ \mu_2$

$(R5)(\psi \circ \mu) \wedge \psi$ implies $\psi \circ (\mu \wedge \psi)$

$(R6)$ If $(\psi \circ \mu) \wedge \psi$ is satisfiable, then $\psi \circ (\mu \wedge \psi)$ implies $(\psi \circ \mu) \wedge \psi$

Figure 1. Belief Revision rationality postulates for knowledge bases

document. This is what intuitively would be expected. On the other hand, the revision $d \circ S$ entails `medicine` $\rightarrow \neg$ `computer science` (from R1), but cannot entail both `medicine` and `computer science` (from R3). This would prevent a query containing the *missing* term from retrieving the document, even though the document contains that term. Therefore, an important effect of chosing $S \circ d$ is that the user can learn relationships that he/she did not know. In fact, the knowledge represented in S can be erroneous and documents can help to rectify it.

Postulate R2 says that if there is no contradiction between a document and a retrieval situation, the revision is the result of their conjunction. As mentioned above, R2 is known as the Preservation Principle and the idea of information economy captured by R2 is suitable for IR.

Postulate R3 says that although the retrieval situation were unsatisfiable, the result of the revision is satisfiable if the document is also satisfiable. An unsatisfiable retrieval situation could come from a user whose knowledge is contradictory. Even in this case, a satisfiable document would keep the revision satisfiable, so that it would not entail any query. On the other hand, an unsatisfiable document would be considered relevant to any query. Basically, this latter case does not affect in an unreasonable way because an unsatisfiable document is a kind of somewhat *pathological* and any *normal* document will be represented by a formula that is neither valid nor unsatisfiable. An interesting discussion about this point can be found in [18].

Postulate R4 states the Principle of the Irrelevance of Syntax and the last two postulates represent the condition that revision is accomplished with minimal change [11]. To illustrate the idea behind these postulates, let us suppose that there is some metric for measuring the distance between the models of ψ, $Mod(\psi)$, and any interpretation I. Postulate (R5) says that if an interpretation I is minimal with respect to a set, $Mod(\psi)$, and I also belongs to a smaller set, $Mod(\mu \wedge \phi)$, then I must also be minimal within the smaller set $Mod(\mu \wedge \phi)$. A violation of the postulate (R6) would imply that an interpretation I may be closer to the KB than J within a certain set, while J is closer than I within some other set. The application of these postulates to the IR domain ensures that a situation is changed minimally in order to accept a document.

3.2. Belief Update

This section briefly analyzes why Belief Update operators are not suitable for our application. The Preservation Principle is a basic notion for Belief Revision methods but it is rejected by Belief Update mechanisms. This means that even though S is consistent with d, the update of S with d is not guaranteed to be equivalent to $S \wedge d$. This contrasts with the expected behaviour for IR, where the final representation should reflect both sources of knowledge without dropping any one. We should only remove parts of the knowledge represented in the retrieval situation S in case of contradiction with the document representation d. To sum up, belief update operators are not a suitable tool for modeling the change that a document makes to a retrieval situation because we want that the preservation principle be fulfilled within that change operation.

Other consideration against the use of update operators is that they do not ensure that consistence is kept. The update of S with d can be inconsistent and, in that case, the document would be considered relevant with respect to any query. Let us recall that we use the result of $S \circ d$ to decide relevance via the entailment $(S \circ d) \models q$ and, if $S \circ d$ is inconsistent, the entailment holds whatever q is. Clearly, this is an unacceptable case for IR and it is not a pathological one. As a matter of fact, retrieval situations coming into contradiction with document's content do really exist.

3.3. Which Belief Revision operator for IR?

In section 3.1 we showed how rationality postulates (R1)-(R6) conform a suitable framework for modeling the change that a document makes to a retrieval situation. Since the Theory of Change was formulated, several Belief Revision operators have been proposed. In this direction, a very interesting study was made by Katsuno and Mendelzon [12]. That work, besides proving the equivalence between the postulates (R1)-(R6) and the original AGM ones for belief sets, presents a review of different Belief Revision operators. It contains an analysis of the way that several operators behave with respect to the postulates. The main conclusion is that only Dalal's revision operator [4], $\circ_D$, satisfies (R1)-(R6) in a nontrivial way. Other well known operators fail to fulfil the postulates. Since postulates capture well the semantics of the change that a document makes to a retrieval situation, we think that Dalal's method is the best way to carry out the revision.

Dalal's revision operator is model-based, i.e. it builds an order among logical interpretations. This order is used to select the models for the revised theory. Roughly speaking, it is a cardinality-based operator that counts the differing terms between interpretations. IR classical approaches usually stand on a indexing vocabulary where each index term or keyword is supposed to be a good representative of a concept. On the logical side, this is represented by a propositional alphabet where each propositional letter represents an index term. Therefore, when changing a retrieval situation with a document by Dalal's operator, the models of the changed situation will be those models of the document having less keywords in *disagreement* with respect to the models of the situation. Clearly, this behaviour is related to IR, where there are several similarity measures counting term matches between representations. In this line, a recent work [14] presents the use of Dalal's operator to get a measure of the uncertainty of $d \models q$. Our use of Dalal's operator is different here because we are interested in the final result of the revision. If the revised situation entails the query, the document is considered relevant. In contrast, the key point in [14] was the ordering among interpretations induced by Dalal's operator but not the final result of the revision process. The ordering among interpretations was used to build a ranking of documents given a query.

A very interesting study about the complexity of several methods for updating and revising knowledge bases was made by Eiter and Gottlob [7]. Specifically, their work focus on the problem of given a knowledge base T, an update formula p and a formula q, decide if q is derivable from $T \circ p$. This is precisely our relevance test. An important point is that Dalal's approach gives better complexity results than the other operators. In the general case complexity is in the class $P^{NP[O(log n)]}$. However, if we assume that p, q and all the formulas belonging to T are Horn formulas and that the size of the update formula is bounded by a constant, the complexity is polynomial, $\mathcal{O}(||T|| \cdot ||q||)$, where $||T||$ and $||q||$ represent the size of the theory and the size of the query respectively.

Some results reached by del Val in [5, 6] are especially interesting for the implementation of the revision $S \circ_D d$. Del Val's work deeply analyzes the sources of complexity of several revision and update methods and, as a result, he proposed a syntactic characterization for the logical formulas that take part in the change operation. If the formulas representing the knowledge base and the new information have specific forms the new knowledge base can be efficiently computed. Specifically, the complexity results for Dalal's Belief Revision operator are very encouraging. The work [6] depicts an algorithm that, given a knowledge base and a new information both stored in Conjunctive Normal Form (CNF), computes the result of Dalal's revision as a CNF formula. This algorithm is executed in polynomial time. The same result of complexity can be obtained for Disjunctive Normal Form (DNF) representations. A CNF formula has the form $d_1 \wedge d_2 \wedge \ldots$ where each d_j is a disjunction $l_1 \vee l_2 \vee \ldots$, where each l_j is a literal, i.e. a propositional letter or its negation. In an analogous way, a DNF formula has the form $c_1 \vee c_2 \vee \ldots$ where each c_j is a conjunction $l_1 \wedge l_2 \wedge \ldots$, where each l_j is a literal, i.e. a propositional letter or its negation. Following the mentioned results, if we have the representation of the retrieval situation and the representation of the document both stored in CNF or both in DNF, we can obtain the representation of the revised situation in polynomial time. Retrieval situations will typically be composed of a set of material implications formalizing relationships between terms, e.g. `linux` $\rightarrow$ `operating systems`. This kind of formulas can be represented by an equivalent disjunctive formula, i.e. $\neg$ `linux` $\vee$ `operating systems`. Then, S can be represented by the conjunction of a set of disjunctive formulas, e.g ($\neg$ `linux` $\vee$ `operating systems`) $\wedge$ ($\neg$ `pascal` $\vee$ `programming`) $\wedge$ `...`. In short, retrieval situations can be easily represented in CNF. Regarding the document representation d, classical vectors with binary weights can be represented as a conjunction of literals such that a positive literal stands for a weight 1 of the involved term in the vector and a negative one stands for a weight 0, e.g. $d = pascal \wedge \neg linux \wedge software \wedge \ldots$. The resulting representation is in CNF and we can apply Del Val's algorithm for computing $S \circ_D d$ in an efficient way. We think that the previous contraints in the form of the representations are not too limiting for our present application and the efficiency of the algorithms is promising for dealing with large amounts of data.

3.4. Partial relevance

As shown in previous sections our proposal stands on deciding relevance using the entailment $(S \circ_D d) \models q$. However this criterion is still too strict because it cannot represent partial relevance. In this respect, we propose to use the results of [14] to get a measure of the uncertainty of the entailment. In that work, the measure of distance between interpretations formalized by Dalal's Belief Revision operator [4], $\circ_D$, was used to define a similarity measure between documents and queries. Formally, the Belief Revision process $q \circ_D d$ builds an order induced by the query between the document's models. This order was extended to define an order between documents given a query. Therefore, the uncertainty of the entailment $d \models q$ is captured. Besides, the ranking subsumes the one imposed by the inner product query-document similarity measure. For the present work we generalize the model to consider retrieval situations: once the retrieval situation has been revised, i.e. $S \circ_D d$, a measure of the uncertainty of the entailment $(S \circ_D d) \models q$ is obtained within the Belief Revision process $q \circ_D (S \circ_D d)$. Note that the proposal developed in [14] is a particular case of the previous formulation where the retrieval situation is an empty theory. That is, when $S = \emptyset$, $(S \circ_D d) \equiv d$.

4. Discussion

Some IR models have tried to estimate the uncertainty of the implication by the conditional probability $P(q/d)$. However the limitation of these approaches was shown by Lewis' triviality results [13]. Basically, only four values of probability are obtained. An important point is that Lewis' triviality is a particular case of Grdenfors triviality [9]. Crestani and Van Rijsbergen developed a model [3] which evaluates the uncertainty of the conditional implication based on an *imaging* process. The model exploits term-term relationships and it is free from triviality. However, it does not consider retrieval situations. Nie and other researchers [16] took into account retrieval situations but their proposal is a general framework within which IR may be considered. In some sense we describe a way of implementing their proposal and thus, the complexity of the problems involved can be analyzed. In fact, Lewis' system of spheres for conditional logic and the orderings used within the Theory of Change are mutually definable [10].

A fundamental aspect captured by logic is partiality. In fact, the previous model is useless with a total information assumption. Basically, retrieval situations are ignored when documents are complete theories. A document is a complete theory if for each propositional formula p either $d \models p$ or $d \models \neg p$ holds. In that case, the document has only one model and, as a result of R1, the revision $S \circ d$ has to be equivalent to d. This means that the information modelized in S is always ignored. On the other hand, when a document is a partial theory, the representation of the retrieval situation can help to complete the document's representation. Van Rijsbergen already pointed out that the most natural assumption is to consider documents as partial descriptions [20]. In fact, we can think about a realistic IR system using this policy: the document analysis phase produces the set of keywords of each document and instead of negating the rest of the system's index terms, the system maintains partial descriptions. In logic this corresponds with the fact that the system does not make a closed world assumption (CWA) but allows documents as partial theories. Given a user's information need expressed as a query and a retrieval situation that contains user's knowledge, the retrieval situation helps to complete the document representation before analyzing the query. This makes a user-oriented completion of the document. For instance, a common user may not know that the language ml has something to do with computer science (cs) but an experienced user would probably now that they are related concepts. As a consequence, the first user's retrieval situation will not contain anything about that relation and the retrieval situation of the experienced user will contain $ml \rightarrow cs$. Then, given a document dealing with ml, if both users articulate a query asking about cs documents, the first user would not access the document and the second one would. This goes in the line that unexperienced users receive generic documents while experienced users receive specific ones. In fact, it does not have much sense to present a ml document to a user that does not know what it is. He/she is probably looking for more general documents about computer science. Then, the precision of the set of retrieved documents is improved, saving the user from inspecting documents that, almost certainly, will not interest him.

5. Conclusion

High expressive IR systems need a modelization of documents and information needs closer to their actual semantics. In order to achieve that, conventional IR models have to go deeper into the modelization of situational factors. The introduction of retrieval situations in IR models can be accomplished using counterfactual conditional logic. A counterfactual $d > q$ evaluated in a retrieval situation S was proposed in the literature to decide the relevance of the document d with respect to the query q in the given retrieval situation. In this work we have studied the implementation of the counterfactual via a change method using the results of the Ramsey Test. We can conclude that the use of a logic with a conditional connective seems to be an overshoot because the test can be done using the classic entailment and a revision operator. Furthermore, this election blocks Grdenfors Triviality Result and, therefore, the resulting logic is non-trivial.

We have analyzed different change semantics in the light of our use for IR. Both semantic and complexity considerations have led to propose Dalal's revision operator to model the change that a document representation makes to a retrieval situation. Previous results about measures of the uncertainty of the entailment have been generalized to deal with retrieval situations. As a result, a measure of partial relevance considering the retrieval situation can be obtained.

Acknowledgements: This work was supported in part by project PGIDT99XI10201B from the government of Galicia, *Xunta de Galicia*, (Spain) and by project PB97-0228 from the government of Spain.

References

[1] C. 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.

[2] C. Alchourrón and D. Makinson. On the logic of theory change: contraction functions and their associated revision functions. *Theoria*, 48:14–37, 1982.

[3] F. Crestani and C. J. van Rijsbergen. Information retrieval by logical imaging. *Journal of Documentation*, 51:3–17, 1995.

[4] M. Dalal. Investigations into a theory of knowledge base revision: Preliminary report. In *Proc. of AAAI-88, the 7th National Conference on Artificial Intelligence*, pages 475–479, 1988.

[5] A. del Val. *Belief Revision and Update*. PhD thesis, Stanford University, 1993.

[6] A. del Val. Syntactic characterizations of belief change operators. In *Proc. of IJCAI'93, the Thirteenth International Joint Conference on Artificial Intelligence*, pages 540–545, Chambery, France, 1993.

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

[8] P. Gärdenfors. Rules for rational changes of belief. In T.Pauli, editor, *Philosophical Essays Dedicated to Lennart Aqvist on His Fiftieth Birthday*, pages 88–101, 1982.

[9] P. Gärdenfors. Belief revisions and the ramsey test for conditionals. *The Philosophical Review*, XCV(1), January 1986.

[10] P. Gärdenfors and H. Rott. Belief revision. In D. Gabbay, C. Hogger, and J. Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming*, volume 4, Epistemic and Temporal Reasoning, pages 35–175. Clarendon Press, Oxford, 1995.

[11] H. Katsuno and A. O. Mendelzon. On the difference between updating a knowledge base and revising it. In J. Allen, R. Fikes, and E. Sandewall, editors, *Proc. of KR'91, the Second International Conference on Principles on Knowledge Representation and Reasoning*, pages 387–393, 1991.

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

[13] D. Lewis. Probability of conditionals and conditionals probabilities. In W. Harper, R. Stalnaker, and G. Pearce, editors, *The University of Western Ontario Series in Philosophy of Science*, pages 129–147, Dordrecht, Holland, 1981. D.Reidel Publishing Company.

[14] D. E. Losada and A. Barreiro. Using a belief revision operator for document ranking in extended boolean models. In *Proc. of SIGIR-99, the 22th ACM Conference on Research and Development in Information Retrieval*, pages 66–73, Berkeley, California, August 1999.

[15] J.-Y. Nie. Towards a probabilistic modal logic for semantic-based information retrieval. In *Proc. of SIGIR-92, the 14th ACM Conference on Research and Development in Information Retrieval*, pages 140–151, Copenhague, June 1992.

[16] J.-Y. Nie, M. Brisebois, and F. Lepage. Information retrieval as counterfactual. *The Computer Journal*, 38(8):643–657, 1995.

[17] F. P. Ramsey. General propositions and causality. In R. B. Braithwaite, editor, *Foundations of Mathematics and Other Logical Essays*, pages 237–257, Routledge and Kegan Paul, New York, 1950.

[18] F. Sebastiani. On the role of logic in information retrieval. *Information Processing and Management*, 34(1):1–18, 1998.

[19] C. J. van Rijsbergen. A non-classical logic for information retrieval. *The Computer Journal*, 29:481–485, 1986.

[20] C. J. van Rijsbergen. Towards an information logic. In *Proc. of SIGIR-89, the 12th ACM Conference on Research and Development in Information Retrieval*, pages 77–86, Cambridge, MA, June 1989.