

Towards a Belief Revision Based Adaptive and Context Sensitive Information Retrieval System

RAYMOND Y.K. LAU

Department of Information Systems, City University of Hong Kong
Tat Chee Avenue, Kowloon, Hong Kong SAR
raylau@cityu.edu.hk

PETER D. BRUZA

School of Information Systems, Queensland University of Technology
GPO Box 2434, Brisbane, QLD 4001, Australia
p.bruza@qut.edu.au

DAWEI SONG

Knowledge Media Institute, The Open University
Walton Hall, Milton Keynes, MK7 6AA, United Kingdom
d.song@open.ac.uk

In an adaptive information retrieval (IR) setting, the information seekers' beliefs about which terms are relevant, or non-relevant, will naturally fluctuate. This article investigates how the theory of belief revision can be used to model adaptive IR. More specifically, belief revision logic provides a rich representation scheme to formalize retrieval contexts so as to disambiguate vague user queries. In addition, belief revision theory underpins the development of an effective mechanism to revise user profiles in accordance with information seekers' changing information needs. It is argued that information retrieval contexts can be extracted by means of the information flow text mining method so as to realize a highly autonomous adaptive IR system. The extra bonus of a belief-based IR model is that its retrieval behavior is more predictable and explanatory. Our initial experiments show that the belief-based adaptive IR system is as effective as a classical adaptive IR system. To our best knowledge, this is the first successful implementation and evaluation of a logic-based adaptive IR model which can efficiently process large IR collections.

Categories and Subject Descriptors:

H.3.3 [Information Storage and Retrieval]: Information Search and Retrieval --- Retrieval Models, Selection process; H.3.4 [Information Storage and Retrieval]: Systems and Software --- User profiles and alert services, performance evaluation (efficiency and effectiveness); H.3.m [Information Storage and Retrieval]: Miscellaneous --- Theoretical Study of Information Retrieval

General Terms: Theory, Algorithms, Experimentation

Additional Key Words and Phrases: Belief revision, retrieval context, information flow, text mining, adaptive information retrieval.

1. INTRODUCTION

Distributed computer-based information systems have undergone huge growth in recent years. The growing availability of large, dynamic, heterogeneous, and distributed sources of information via the World Wide Web (Web) has contributed greatly to “information overload” (Maes 1994; Levy 2005). Accordingly, there is a growing demand for the development of highly autonomous and adaptive information retrieval (IR) systems which can automatically select relevant information items on behalf of their users.

Figure 1 highlights the main functional components of an adaptive IR system. An information seeker first translates their implicit information needs into queries. Recurring queries are often stored in a user profile. On the other hand, information objects from specific information sources

such as the Web are characterized by a particular indexing scheme. These characterizations are also stored in the local cache of the adaptive IR system. The matching mechanism of the adaptive IR system tries to match the user's information needs with incoming information objects by comparing the corresponding queries and document characterizations. The information objects deemed relevant by the adaptive IR system are dispatched to the user in the form of a retrieval result set. After reviewing the information objects, the user can then provide relevance feedback to the adaptive IR system. The learning mechanism of the adaptive IR system will use this feedback information to revise and refine the initial user profile. As a user's information needs as well as the underlying retrieval context may change over time, the adaptive IR system should continuously revise its user profile based on the user's most recent relevance feedback so as to maintain the effectiveness of the information matching processes.




Figure 1 – An Overview of an Adaptive Information Retrieval System

In general, information search is divided into information retrieval (IR) and information filtering (IF) although they share some common activities (Belkin and Croft 1992). IR often refers to the situation that an information seeker takes an active role to specify their ad hoc queries, whereas IF is concerned with the removal of irrelevant information from an incoming stream of information based on the information seeker's long term and recurring retrieval goals stored in a user profile. As the proposed adaptive IR system can support both ad hoc interactive retrieval tasks and long-term recurring retrieval tasks, the more general term “IR” is applied to our prototype system.

IR involves uncertainty both in terms of query and document representation. By way of illustration, given the query “Java”, an IR system may return documents about “Computer Programming” or documents about “Merapi”, a volcano on the island of Java. If the information seeker is a computer programmer (i.e., in the context of information technology), documents about java programming are relevant. However, if the information seeker is a volcanologist (i.e., in the context of volcanology), documents about “Merapi” are more likely to be considered relevant. From the perspective of the IR system, there is uncertainty in determining which set of documents is relevant because the queries captured in the user profile are often incomplete and the implicit retrieval context may not be readily available. In fact, the issues of *partiality* and *uncertainty* are inherent in any IR processes (van Rijsbergen 1986; Lalmas and Bruza 1998), and we believe that these issues contribute significantly to information overload.

One of the justifications for the development of a belief revision based adaptive IR system is that the expressive power of logic is believed to be able to model most of the fundamental aspects in information retrieval (Chiaromella and Chevallet 1992; Lalmas and Bruza 1998; van Rijsbergen 1986; Sebastiani 1998). Previous psychological study has shown that the postulates characterizing preferential logic are compatible with the characteristics of human reasoning (Neves et. al. 2000). It has also been reported that possibilistic rather than probabilistic reasoning is closer to the kind of approximate reasoning exercised by human experts (Raufaste and Neves 1998). The AGM belief revision logic which underpins the proposed adaptive IR model has close connection with the preferential logic and possibilistic logic, and hence it shows promise in modeling the human approximate reasoning processes in IR.

The logical uncertainty principle (van Rijsbergen 1986), which is a generalization of the Ramsey test (Gärdenfors 1988), has spawned fruitful theoretical investigations into logic-based approaches for IR. The logical uncertainty principle states that:

“given any two sentences x and y , a measure of the uncertainty of $x \rightarrow y$ relative to a given data set is determined by the minimal extent to which we have to add information to the data set, to establish the truth of $x \rightarrow y$.”

With a logic-based adaptive IR model, sentence x can be taken as the representation of an information seeker's needs, and sentence y is the characterization of a document (Nie 1986). In addition, a data set can be interpreted as a retrieval context which characterizes a particular information matching situation (Nie et. al. 1995). A retrieval context may consist of an information seeker's background, their long term search goals, tasks at hand, knowledge about a retrieval domain, etc. The subject of context has received a great deal of attention in the field of information retrieval (Cool 2001; Cool and Spink 2002; Cohen and Singer 1996; Lawrence 2000; Nie et. al. 1995). Recently, there are series of ACM SIGIR workshops (IRiX) which examine the prominent features of contexts (e.g., task features, document features, environment features, etc.) and applying these features to improve IR (Ingwersen and Järvelin 2005). Belief revision can be taken as a means of directly implementing the logical uncertainty principle for adaptive IR because of its close connection with the Ramsey test (Gärdenfors 1988). More specifically, a retrieval context is formally represented by a belief set K and the changing retrieval context is modeled by the corresponding belief revision function K_α^* (Alchourrón et. al. 1985), where α is the logical representation of a user's relevance feedback. Then, a document d can be evaluated with respect to the refined retrieval context K and the query q . Such a matching process is underpinned by $q \sim d$, where \sim is the expectation inference relation (Gärdenfors and Makinson 1994), a class of nonmonotonic inference relation. As a matter of fact, the idea of applying belief revision and nonmonotonic reasoning to practical applications has also been explored in other application domains (Bessant et. al. 1998).

One contribution of this paper is the illustration of an effective way for discovering retrieval contexts based on the computation of *information flow* through semantic space models (Song and Bruza 2001; Song and Bruza 2003). As retrieval contexts may evolve over time, this paper also discusses a formal approach of revising queries and contexts based on the AGM belief revision

framework. Furthermore, despite the discussions of the benefits of logic-based IR models in IR literature for the past two decades (Chiaramella and Chevallet 1992; Hunter 1997; Lalmas and Bruza 1998; Lau et. al. 2001; Losada and Barreiro 2001; van Rijsbergen 1986; Sebastiani 1998), few successful empirical evaluations of logic-based IR models have emerged. One important contribution of this paper is to report the evaluation of our belief revision logic based adaptive IR system based on large IR benchmark collections. Last but not least, the merit of improved explanatory power of a logic-based adaptive IR model is confirmed based on a usability study.

The remainder of the paper is organized as follows. In Section 2, a review of logic-based IR models is presented, followed by a preliminary of the AGM belief revision logic highlighted in Section 3. Section 4 proposes an information theoretic method for inducing an information seeker's interests, and Section 5 describes the information flow based text mining method for retrieval context discovery. Based on the examples in context revision and information matching, Section 6 illustrates the operational details of the belief revision based adaptive IR model. Section 7 reports our empirical evaluation of the belief-based adaptive IR model. Finally, we offer concluding remarks and describe future directions of this research work.

2. LOGIC-BASED MODELS FOR IR

Investigations into logic-based IR have attempted to formalize the notion of “aboutness” by axiomatising its properties in terms of a neutral, theoretical framework (Bruza and Huibers 1994; Huibers and Wondergem 1998; Bruza et. al. 2000). The motivation for this has been to study the aboutness relation from a theoretical stance in order to better understand what properties of this relation promote effective retrieval (as well as which properties do not). The neutral, underlying framework is important as it allows aboutness to be studied independent of the idiosyncrasies of a given IR model. There is as yet no consensus regarding the property of aboutness except that it should be logic-based (Huibers and Wondergem 1998; Bruza et. al. 2000). The notion of aboutness in IR has been applied to examine the postulates characterizing the AGM belief revision logic to see if the belief revision framework is applicable in the context of adaptive IR (Lau et. al. 1999). Therefore, the adaptive IR model presented in this paper has its roots in earlier theoretical work in logic-based IR.

Huibers and van Linder (1996) attempted to formalize intelligent information retrieval agents based on modal logic. Modal operators were introduced to address the essential concepts such as aboutness, non-aboutness, and information preclusion (Bruza and Huibers 1994) in IR. For example, one kind of retrieval agent is defined based on the notion of aboutness $d \models_a q$ (i.e., a document d to be about a query q). Strictly speaking, $d \models_a q$ is established iff the agent knows that the query q is satisfied in at least one document model of d . It is believed that such a satisfiability relation should be developed based on non-classical logic (Huibers and van Linder 1996). Moreover, the retrieval agent considers a document d to be non-about q , denoted $d \not\models_a q$, iff it knows that d implies the negation of q . For the belief revision based IR approach presented in this paper, a document is about (or not about) a query to a certain degree and this gradated information matching is modeled by epistemic entrenchment (Gärdenfors and Makinson 1988). One advantage of a formal approach for the development of information agents is that the agents'

retrieval behavior is explanatory. This is a facet we intend to exploit as typical IR systems are not scrutable.

In statistical analysis, the relationships among key phrases are established by frequency ratios, whereas in semantic analysis, the relationships are established by meaning. It is believed that semantic information is critical in matching a user's needs to information objects (Hunter 1996). For automating the use of semantic information, it is necessary to specify when a particular specialization, generalization, or synonym relationship should be used. Accordingly, an expressive formal framework is required to capture and reason about the semantic information. Hunter (Hunter 1996; Hunter 1997) proposed to use nonmonotonic logics, particularly default logic, to process semantic information about terms, and hence to identify the semantic relationships between queries and documents. For example, given the default rule $(\frac{oil \wedge cooking \wedge \neg petroleum}{\neg petroleum})$ and a query $(olive \wedge oil \wedge cooking)$, the original query will be

refined to exclude any information items about petroleum. It was suggested that the default rules of term relationships could be manually elicited from domain experts by asking them to illustrate the synonym, polysemy, generalization, and specialization relationships (Hunter 1997). One weakness of default logic is that it does not augur well for large scale implementation. For instance, the computational complexity of model checking in default logic is NP-Complete (Liberatore and Schaerf 1998). Moreover, Hunter's default logic approach for IR assumes that default rules of term relationships are manually elicited from domain experts. This is the classical knowledge acquisition bottle-neck. This article extends the approach of logic-based semantic information processing by employing the information flow method to automatically discover the semantic relationships among terms.

Logical imaging has been applied to develop IR models (Crestani and van Rijsbergen 1995; Crestani 1998). The goal is to evaluate the probability of the conditional $(d \rightarrow q)$, whereas d and q stand for a document and a query respectively. Logical imaging has its root in non-classical logic but it is also based on the kinematics of probability distributions over terms. When the probability $\Pr(d \rightarrow q)$ is evaluated, the formula d will be imaged on the closest world(s) t , where t is a term representing a world in the logical imaging IR model. Then, the formula q is evaluated in these closest world(s). To capture the uncertainty in IR processes, the worlds (i.e., terms) are characterized by a probability distribution. These prior probabilities can be induced based on the Inverse Document Frequencies (IDF) of terms in a collection. The IR logical imaging paradigm consists of several methods to deal with the kinematics of probabilities associated with the worlds. For instance, imaging on the d -world(s) is taken as transferring the priori probabilities from the non d -world(s) to the closest d -world(s) according to a distance measure derived from the mutual information between pairs of terms. For standard imaging, the probability associated with a non d -world is simply transferred to the closest d -world. Then, for each term appearing in a query, the posterior probability of the term is summed to derive the Retrieval Status Value (RSV) of the document with regards to the query q . So, for standard imaging, the RSV is derived by: $\Pr(d \rightarrow q) = \sum_t \Pr(t) \times \tau(t_d, q)$, where $\tau(t_d, q) = 1$ if a query term appears in a d world (i.e., d

and q have overlapping terms); otherwise it is zero. For general imaging, standard imaging is generalized in the sense that there could be more than one closest world where d is true. For general logical imaging, the percentage of probability transferred from each non d -term to a d -

term can be defined separately via individual opinionated probability distribution. The main difference between the logical imaging IR model and the belief revision based IR model proposed here is that term weights representing a user's preferences are induced with respect to *epistemic entrenchment* which satisfies possibilistic (Dubois and Prade 1991) rather than probabilistic axioms in the belief revision based IR model. Above all, the entrenchment degrees of terms are derived according to a user's preferences over the underlying terms, and the kinematics of the entrenchment degrees are also driven by the changes of a user's information preferences.

The notions of beliefs, desires, and intentions have been applied to characterize an information seeker's (e.g., a librarian) high level IR goals and belief revision has been exploited to simulate the changes of mental states of an information seeker (Logan et. al. 1994). Natural Language Processing (NLP) technique was used to induce the system's beliefs about an information seeker's information needs based on the continuous interactions between the information seeker and the IR system. As IR can be taken as comprising many low level sub-tasks, the corresponding IR system is designed as a multi-agent system with each autonomous agent performing a particular IR sub-task. As the authors have already indicated, computational complexity is the main obstacle against applying such a belief-based multi-agent system to support real-world IR activities (Logan et. al. 1994). Our belief revision mechanism is implemented based on a computationally efficient transmutation method (Williams 1997).

Dalal's belief revision operator (Dalal 1988) was applied to document ranking in IR (Losada and Barreiro 1999; Losada and Barreiro 2001; Losada and Barreiro 2003a). Dalal's revision makes use of the cardinality of the *symmetric difference* between two interpretations I and J as a measure of the distance between these interpretations¹ (i.e., $\text{dist}(I, J)$). For example, the semantic distance between the set of models of ψ (i.e., $M(\psi)$) and I is defined by $\text{dist}(M(\psi), I) = \min_{J \in M(\psi)} \text{dist}(J, I)$. Thereby, a faithful assignment of a total pre-order \leq_ψ is defined as $I \leq_\psi J$ iff $\text{dist}(M(\psi), I) \leq \text{dist}(M(\psi), J)$. In IR, if a user's information needs N and a document Doc are represented by formulae q and d respectively, the semantic similarity between N and Doc can be approximated by the symmetric distance of the corresponding models. For example, for each $m \in M(d)$, $\text{dist}(M(q), m) = \min_{J \in M(q)} \text{dist}(J, m)$ is computed. An average measure can then be applied to compute the symmetric distance between $M(q)$ and $M(d)$ by

$$\text{sim}(D, N) = \frac{\sum_{m \in M(d)} \text{dist}(M(q), m)}{|M(d)|}. \text{ However, it is extremely costly to compute the symmetric}$$

difference between sets of models even with a moderate number of atoms (Losada and Barreiro 2001). We employ a formula-based representation for our belief revision model so that the belief-based adaptive IR system is computationally tractable. Recently, Losada and Barreiro have also adopted a formula-based approach to implement their belief revision based IR matching function and demonstrated some successes in the TREC-3 routing task (Losada and Barreiro 2003b). One main difference between our belief revision based adaptive IR model and Losada and Barreiro's IR model is that we apply the AGM belief revision framework to model an IR system's changing beliefs about retrieval situations, whereas Losada and Barreiro's work

¹ An interpretation is a mapping function from the propositional symbols into the set {true, false}. A model of a logical expression is an interpretation that maps the logical expression into true.

(Losada and Barreiro 1999; Losada and Barreiro 2001; Losada and Barreiro 2003a) focuses on a belief-based ranking function.

3. THE AGM BELIEF REVISION LOGIC

The AGM belief revision framework, one of the most influential works in belief revision theory, is coined after its founders Alchourrón, Gärdenfors and Makinson (Alchourrón et. al. 1985). In this framework, belief revision processes are taken as the transitions among belief states. A belief state (set) K is represented by a theory of a classical language L . A belief is represented by a sentence of L supplemented with an entrenchment degree indicating the degree of firmness of such a belief. Three principle types of belief state transitions are identified and modeled by the corresponding belief functions:

- Expansion K_{α}^{+} is the process of accepting a new belief α that does not contradict existing beliefs in a belief set K ;
- Contraction K_{α}^{-} is the removal of a belief α and all other beliefs that logically imply α from a belief set K ;
- Revision K_{α}^{*} is the incorporation of a belief α that may contradict existing beliefs in a belief set K ; this operation corresponds to adding a new query α into an existing user profile K where not α was specified.

Belief expansion K_{α}^{+} can be interpreted as adding the relevance feedback information α into a user profile which stores the representation of an existing retrieval situation K . Since the relevance feedback information α is consistent with the current retrieval situation K (e.g., the existing queries), the new retrieval requirements (i.e., the relevance feedback) can simply be added to the user profile without invoking complicated updating operation. Similarly, belief contraction K_{α}^{-} can be taken as removing outdated queries from the user profile based on the relevance feedback α . Belief revision K_{α}^{*} is the most common and also the most complicated belief change operation. Since the new retrieval requirements α (e.g, requiring documents about Java) is contradictory to the existing retrieval situation K (e.g., not requiring documents about Java), some existing information stored in the user profile must be given up. This belief revision operation should be executed according to sound principles such that useful information stored in the user profile can still be maintained.

The AGM framework comprises sets of postulates to characterize the belief functions for *consistent* and *minimal* belief revision. In the context of IR, the belief revision function K_{α}^{*} can be applied to regulate (according to the principles of minimal and rational revision) the changes of a retrieval situation K given an information seeker's relevance feedback α . For example, given an initial retrieval situation $K = \{\text{sculpture} \rightarrow \text{art}, \neg\text{sculpture}\}$, it is clear that the information seeker does not want items about "sculpture". The sentence ($\text{sculpture} \rightarrow \text{art}$) means

requiring items about “sculpture” implies (i.e., logically entails) requiring items about “art” for this particular information seeker. If the information seeker later informs the IR system that their preference has changed to “sculpture”, the revised retrieval situation becomes $K' = \{\text{sculpture} \rightarrow \text{art}, \text{sculpture}, \text{art}\}$ via executing the belief revision function $K_{\text{sculpture}}^*$. Based on this example, it can be observed that one advantage of the belief revision based IR model is that an information seeker’s shifting interest (e.g., documents about “art”) can be automatically deduced by the IR system. Moreover, this change can be explained based on the logical axiom of *modus ponens* (i.e., $\text{sculpture} \wedge (\text{sculpture} \rightarrow \text{art}) \vdash \text{art}$). The symbol \wedge represents the logical AND, and the symbol \vdash denotes the logical derivability relation. Above all, the change applied to a user profile is carried out according to sound AGM principles such as minimal belief change. For instance, there is no reason of removing the belief ($\text{sculpture} \rightarrow \text{art}$) even though the belief ($\neg \text{sculpture}$) does not hold any more.

The AGM framework also specifies the constructions of the belief functions based on various mechanisms. One of them is *epistemic entrenchment* \leq (Gärdenfors and Makinson 1988). It captures the notions of significance, firmness, or defeasibility of beliefs. Formally, an epistemic entrenchment ordering is a total pre-order of the sentence (e.g., α, β, γ) in L , and is characterized by the following postulates (Gärdenfors and Makinson 1988):

- EE1: $\forall \alpha, \beta, \gamma \in K : \alpha \leq \beta \leq \gamma$ implies $\alpha \leq \gamma$ (Transitivity)
- EE2: $\forall \alpha, \beta \in K : \alpha \vdash \beta$ implies $\alpha \leq \beta$ (Dominance)
- EE3: $\forall \alpha, \beta \in K : \alpha \leq \alpha \wedge \beta$ or $\beta \leq \alpha \wedge \beta$ (Conjunctiveness)
- EE4: If $K \neq K_{\perp}$, $\alpha \notin K$ iff $\forall \beta \in K : \alpha \leq \beta$ (Minimality)
- EE5: $\forall \beta \in K : \beta \leq \alpha$ implies $\vdash \alpha$ (Maximality)

whereas $\alpha \leq \beta$ means that β is at least as entrenched as α . The notation K_{\perp} in EE4 means a belief set with inconsistent beliefs, which is not a desirable state for a rational agent. Intuitively, epistemic entrenchment relations induce preference orderings of beliefs according to the importance of these beliefs in the face of change. If inconsistency arises during a belief revision operation, the least significant beliefs (i.e., beliefs with the lowest entrenchment degree) are given up in order to restore consistency. The postulates of epistemic entrenchment are valid in the context of IR (Lau et. al. 1999). For a computer-based implementation of epistemic entrenchment and hence the AGM belief functions, Williams (1995) developed finite partial entrenchment rankings to represent epistemic entrenchment orderings.

Definition 1. A finite partial entrenchment ranking is a function \mathbf{B} that maps a finite subset of sentences in L into the unit interval $[0, 1]$ such that the following conditions are satisfied for all $\alpha \in \text{dom}(\mathbf{B})$:

- (PER1) $\{\beta \in \text{dom}(\mathbf{B}) : \mathbf{B}(\alpha) < \mathbf{B}(\beta)\} \vdash \alpha$;

(PER2) If $\vdash \neg \alpha$ then $\mathbf{B}(\alpha) = 0$;

(PER3) $\mathbf{B}(\alpha) = 1$ if and only if $\vdash \alpha$.

Essentially, PER1 states that the set of sentences ranked strictly higher than a sentence α cannot entail α . This property corresponds to the Dominance property of epistemic entrenchment (Gärdenfors and Makinson 1988). $\mathbf{B}(\alpha)$ is referred to as the degree of entrenchment of an explicit belief α . The set of explicit beliefs of \mathbf{B} is $\{\alpha \in \text{dom}(\mathbf{B}): \mathbf{B}(\alpha) > 0\}$, and is denoted $\text{exp}(\mathbf{B})$. The set of implicit beliefs $K = \text{Cn}(\text{exp}(\mathbf{B}))$ is denoted $\text{content}(\mathbf{B})$, where Cn is the classical consequence operator. For example, $\mathbf{B}(\text{java}) = 0.6$ represents the entrenchment degree (i.e., firmness) of the belief about information item “java”. The finite partial entrenchment ranking $\mathbf{B} = \{(\text{computer} \wedge \text{java}, 0.6), (\text{computer} \wedge \text{java} \rightarrow \text{programming}, 0.5), (\text{programming}, 0.5)\}$ satisfies the property PER1 to PER3 because $\{(\text{computer} \wedge \text{java}), (\text{computer} \wedge \text{java} \rightarrow \text{programming})\} \vdash \text{programming}$, and $\mathbf{B}(\text{programming}) = 0.5$ which has the same entrenchment degree as the set of sentences entailing it. However, the ranking $\mathbf{B} = \{(\text{computer} \wedge \text{java}, 0.6), (\text{computer} \wedge \text{java} \rightarrow \text{programming}, 0.5), (\text{programming}, 0.4)\}$ does not satisfy the properties of finite partial entrenchment ranking since beliefs with higher entrenchment degree in the ranking entail (\vdash) a belief with lower entrenchment degree (i.e., violating PER1). In particular, the belief “programming” does not have the same firmness as the set of beliefs which logically entails it. In order to describe the epistemic entrenchment ordering $\leq_{\mathbf{B}}$ generated from a finite partial entrenchment ranking \mathbf{B} , it is necessary to rank implicit beliefs.

Definition 2. Let $\alpha \in L$ be a contingent sentence. Let \mathbf{B} be a finite partial entrenchment ranking and $\beta \in \text{exp}(\mathbf{B})$. The degree of entrenchment of an implicit belief α is defined by:

$$\text{degree}(\mathbf{B}, \alpha) = \begin{cases} \sup(\{B(\beta) \in \text{range}(\mathbf{B}) : \text{cut}_{\leq}(\beta) \vdash \alpha\}) & \text{if } \alpha \in \text{content}(\mathbf{B}) \\ 0 & \text{otherwise} \end{cases}$$

where the notation \sup refers to supremum (i.e., the least upper bound) of a possibility distribution (Dubois and Prade 1995). The $\text{cut}_{\leq}(\beta)$ operation extracts a set of explicit beliefs which is at least as entrenched as β according to a finite partial entrenchment ranking \mathbf{B} . \vdash is the classical inference relation. More precisely, a cut operation is defined by: $\text{cut}_{\leq}(\beta) = \{\gamma \in \text{dom}(\mathbf{B}): \mathbf{B}(\beta) \leq \mathbf{B}(\gamma)\}$. For example, given the belief set $\mathbf{B} = \{(\text{google}, 0.8), (\text{google} \rightarrow \text{search-engine}, 0.5)\}$, the operation $\text{cut}_{\leq}(\text{google})$ will return a single belief “google”. Moreover, $\text{degree}(\mathbf{B}, \text{search-engine}) = 0.5$ is derived according to Definition 2 because the minimal entrenchment degree in the strongest cut of \mathbf{B} that logically entails “search-engine” is 0.5.

In a belief revision based adaptive IR system, queries and query contexts are represented by a set of beliefs. When an information seeker's needs and the underlying retrieval context change, the entrenchment degrees of the corresponding beliefs are raised or lowered in the adaptive IR system's knowledge base. Raising or lowering the entrenchment degree of a belief is conducted via a belief revision operation $\mathbf{B}^*(\alpha, i)$ where α is a sentence and i is the new entrenchment

degree. For example, if an information seeker is interested in documents about “volcano” now, the belief revision operation $\mathbf{B}^*(\text{volcano}, 0.8)$ can be invoked to revise such a belief into the adaptive IR system’s knowledge base to represent their current interest. The entrenchment degree of 0.8 can be computed based on our preference induction mechanism described in Section 4. The Rapid Anytime Maxi-adjustment (RAM) method (Lau 2003) developed based on the Maxi-adjustment method (Williams 1997) is proposed to implement the belief revision operation $\mathbf{B}^*(\alpha, i)$:

Definition 3. Let α be a contingent sentence, $j = \text{degree}(\mathbf{B}, \alpha)$ and $0 \leq i < 1$. The (α, i) Rapid Maxi-adjustment of \mathbf{B} is $\mathbf{B}^*(\alpha, i)$ defined by:

$$\mathbf{B}^*(\alpha, i) = \begin{cases} B^-(\alpha, i) & \text{if } i < j \\ (B^-(\neg\alpha, 0))^+(\alpha, i) & \text{if } i > j \\ B^+(\alpha, i) & \text{if } i = j \text{ and } j > 0 \text{ and } \alpha \notin \text{exp}(B) \\ B & \text{otherwise} \end{cases}$$

For all $\beta \in \text{dom}(\mathbf{B})$, $\mathbf{B}^-(\alpha, i)$ is defined as follows:

1. For β with $B(\beta) > j$, $B^-(\alpha, i)(\beta) = B(\beta)$.

2. For β with $i < B(\beta) \leq j$,

$$B^-(\alpha, i)(\beta) = \begin{cases} i & \text{if } \{\gamma : B^-(\alpha, i)(\gamma) > B(\beta)\} \cup \\ & \{\delta : B^-(\alpha, i)(\delta) = B(\beta) \wedge \text{Seq}(\delta) \leq \text{Seq}(\beta)\} \mid \neg\alpha \\ B(\beta) & \text{otherwise} \end{cases}$$

3. For β with $B(\beta) \leq i$, $B^-(\alpha, i)(\beta) = B(\beta)$.

For all $\beta \in \text{dom}(\mathbf{B}) \cup \{\alpha\}$, $\mathbf{B}^+(\alpha, i)$ is defined as follows:

1. For β with $B(\beta) \geq i$, $B^+(\alpha, i)(\beta) = B(\beta)$.

2. For β with $j \leq B(\beta) < i$,


$$B^+(\alpha, i)(\beta) = \begin{cases} i & \text{if } i < \text{degree}(B, \alpha \rightarrow \beta) \\ \text{degree}(B, \alpha \rightarrow \beta) & \text{otherwise} \end{cases}$$

3. For β with $B(\beta) < j$, $B^+(\alpha, i)(\beta) = B(\beta)$.

The intuition of the RAM method is that if the new entrenchment degree i of a sentence α is less than its existing degree j , the belief revision operation $\mathbf{B}^*(\alpha, i)$ invokes a contraction process, that is $\mathbf{B}^-(\alpha, i)$. In other words, the entrenchment degree of α will be lowered. If the new degree i of α is higher than its existing degree j , an expansion operation $\mathbf{B}^+(\alpha, i)$ should be initiated. However, to ensure that the set of beliefs remains consistent in an agent’s knowledge base, $\neg\alpha$

must first be assigned the lowest entrenchment degree (i.e., contracting it from the theory). Therefore, the contraction operation $\mathbf{B}^-(\neg\alpha, 0)$ must be executed before raising the degree of α to i (i.e., adding the belief α to the knowledge base). If the new degree i of α is equal to its existing degree j and α is not explicitly stored in the knowledge base (i.e., $\alpha \notin \text{exp}(\mathbf{B})$), it will be added to the knowledge base by executing the belief expansion operation $\mathbf{B}^+(\alpha, i)$. Since α has already been an implicit belief in the knowledge base, inconsistency will not occur because of the belief expansion operation.

During raising or lowering the entrenchment degree of α , the degrees of other sentences (e.g., β) are adjusted in a minimal way such that the properties PER1, PER2, and PER3 are maintained. With reference to Definition 3, both the belief contraction operation $\mathbf{B}(\alpha, i)$ and the belief expansion operation $\mathbf{B}^+(\alpha, i)$ are further defined by three sub-operations. For belief extraction, the notation $\mathbf{B}(\alpha, i)(\beta)$ refers to the new entrenchment degree of β after contracting α from the finite partial entrenchment ranking \mathbf{B} . The sub-operations 1 and 3 state that existing beliefs β with entrenchment degree strictly higher than the existing degree of α , or β with entrenchment degree less than or equal to the new entrenchment degree of α , will not be affected by the belief contraction operation. For other beliefs β , their entrenchment degree may either be lowered to i (if β together with other more entrenched beliefs in \mathbf{B} entail α) or remain unchanged (if β together with other more entrenched beliefs in \mathbf{B} do not entail α).



Belief Contraction: $\mathbf{B}(\alpha, i)$

Figure 2 – The Belief Contraction Operation

The sub-operation 2 of belief contraction can be depicted by Figure 2, whereas the entrenchment degree of α is lowered to i . As can be seen, from Figure 2, after such a belief change, the belief β together with other beliefs more entrenched than β will logically entail α , which leads to the violation of property PER1 (i.e., the dominance property) for finite partial entrenchment ranking. Therefore, the entrenchment degree of β should also be lowered to i (the minimal change) in order to restore \mathbf{B} to the normal state. With reference to Figure 2, an example in the context of IR is followed. Assuming that the formulas α , β , χ stand for “Web”, “html”, and “agent” respectively, the implication $\beta \rightarrow \alpha$ in Figure 2 indicates that if the information seeker is interested in documents about “html”, they will be interested in documents about “Web” as well. In this example, the belief (html \rightarrow web) is the most entrenched belief although this may not necessarily be true in every world-world IR scenario. If the information seeker explicitly states

that they are no longer interested in the “Web” (i.e., lowering the entrenchment degree of α from j to i in Figure 2), their degree of interest about “html” as modeled in an IR system’s knowledge base should be lowered as well such that the property of finite partial entrenchment ranking (e.g., PER1) is maintained. The intuition is that if $(\text{html} \rightarrow \text{web})$ is a very certain belief for this particular information seeker, and they are not interested in the Web any more, it makes sense to drop the belief of “html”; otherwise the belief “Web” will still hold because the beliefs $\{(\text{html} \rightarrow \text{web}), \text{html}\}$ always implies the belief “Web”. On the other hand, the belief “agent” is not affected by this belief contraction operation (i.e., its degree of significance should be the same as before).

The *Seq* function defined in the sub-operation 2 assigns a unique sequence number to a sentence β if there is more than one sentence with the same entrenchment degree (i.e., at the same entrenchment rank). Under such a circumstance, it does not matter which sentence is contracted first (if they violate PER1) because these sentences are equally preferred or not preferred from an information seeker’s point of view. This is the main difference between the Maxi-adjustment method (Williams 1997) and the RAM method (Lau 2003) for implementing the AGM belief revision. On the other hand, the belief χ will not be affected by the belief contraction operation for α because its entrenchment degree is higher than the original degree of α , and this rule is defined in the sub-operation 1 of $\mathbf{B}^-(\alpha, i)$ in Definition 3.

Similarly, for belief expansion $\mathbf{B}^+(\alpha, i)$, the sub-operations 1 and 3 state that existing beliefs β with entrenchment degree higher than or equal to the new entrenchment degree of α , or β with entrenchment degree strictly less than the existing degree of α , will not be affected by the belief expansion operation. For other beliefs β , their entrenchment degree may either be raised to i or $\text{degree}(\mathbf{B}, \alpha \rightarrow \beta)$ dependent on which one is lower (sub-operation 2).

If a finite partial entrenchment ranking \mathbf{B} has x natural partitions, it only requires $\log_2 x$ classical satisfiability checks (Lang 1997). Therefore, given the propositional Horn logic as the representation language, the transmutation based RAM method for AGM belief revision only involves polynomial time complexity in the worst case. By means of the anytime approximation of an AGM belief revision operation and other optimization techniques, the belief revision based adaptive IR system can be scaled up to support large scale IR. For example, since a belief revision operation is computationally expensive, the revision of a belief will be deferred until its cumulative change is larger than a pre-defined system threshold. We use SICStus Prolog, a commercially available Prolog system, to carry out theorem proving (\vdash). Our Java-based belief revision engine utilizes the Jasper Java interface of SICStus Prolog to communicate with its inference engine. The detailed examples of how our belief revision works within the IR context will be given later in Section 6.

4. INDUCING USER PREFERENCES AND ENTRENCHMENT ORDERINGS

Conceptually, given a retrieval context, an IR system needs to establish a focus (i.e., a user's specific interest) over such a context. With reference to human information processing theory, this means passing a stimuli to the long term memory to trigger the *spreading activation process* (Card et. al. 1983). Since a user may have difficulty in specifying their specific interest, a highly autonomous IR system should be able to automatically induce a user's interest based on the user's interactions with the system. For example, if the user has requested some documents recently, the retrieved documents will form the basis of supervised learning for the user's preferences. In terms of document representation, a document is pre-processed according to traditional IR techniques to extract a set of tokens (e.g., stems, n-grams, or phrases) as its characterization (Salton and McGill 1983). At the symbolic level, each token t is mapped to a positive literal of the classical propositional language L .

Conversely, a user's information need is induced based on a set of relevant documents D^+ and a set of non-relevant documents D^- directly or indirectly judged by the information seeker (e.g., based on archived documents or viewing time). Essentially, three types of tokens can be extracted; positive tokens represent what items the information seeker would like to retrieve; negative tokens indicate what the information seeker does not want; neutral tokens are not good indicators of their information needs. The following preference induction method is used to extract various types of tokens and induce the corresponding preference values. It is developed based on the information theoretic *Keyword Classifier* which was successfully applied to adaptive information filtering (Kindo et. al. 1997):

$$pre(t) = \varepsilon \times \tanh\left(\frac{df(t)}{pos} \times \Pr(rel | t) \times \log_2 \frac{\Pr(rel | t)}{\Pr(rel)} - \frac{df(t)}{neg} \times \Pr(nrel | t) \times \log_2 \frac{\Pr(nrel | t)}{\Pr(nrel)}\right)$$

where $-1 < pre(t) < 1$ is the preference value for a term t . The proposed preference induction mechanism is similar to a linear classifier. If a term appears in many positive documents (i.e., $df(t) > pos$) and the presence of the term is likely a contributing factor (i.e., $\Pr(rel|t) > \Pr(rel)$) towards the relevance of the document, a high positive preference value will be derived. Similarly, the classifier can also take into account the presence of negative feedback. The relative weight of the positive evidence and the negative evidence is controlled by adjusting the parameters of pos and neg respectively.

It should be noted that expected cross entropy EH (Koller and Sahami 1997) bears much similarity with the above preference induction formula. Expected cross entropy is defined by $EH(t, C) = \Pr(t) \sum_{c \in C} \Pr(c|t) \times \log_2 \frac{\Pr(c | t)}{\Pr(c)}$, where the set C includes {relevant, non-relevant}.

The terms pos and neg are the learning thresholds for positive terms and negative terms respectively. The function \tanh is the hyperbolic tangent function. The adjustment factor ε ensures that the induced entrenchment degree is less than the maximal entrenchment degree because all the induced beliefs should be contractable (i.e., they are not tautologies).

$\Pr(rel | t) = \frac{df(t_{rel})}{df(t)}$ is the estimated conditional probability (based on the training data) that a document is relevant given that it contains a term t . It is expressed as the fraction of the number of relevant documents which contain the term t (i.e., $df(t_{rel})$) over the total number of documents which contain term t (i.e., $df(t)$). Similarly, $\Pr(nrel | t) = \frac{df(t_{nrel})}{df(t)}$ is the estimated conditional probability that a document is non-relevant if it contains term t . The term $df(t_{nrel})$ represents the number of non-relevant documents which contain term t . In addition, $\Pr(rel) = \frac{|D^+|}{|D^+| + |D^-|}$

and $\Pr(nrel) = \frac{|D^-|}{|D^+| + |D^-|}$ are the priori probabilities that a document is relevant and non-

relevant respectively. A positive $pre(t)$ indicates that the underlying term t is a positive token, whereas a negative preference value implies that t is a negative token. If the preference value of a token is below a threshold λ , the token is considered neutral. A positive token is mapped to a positive literal such as l , whereas a negative token is mapped to a negated literal such as $\neg l$. The entrenchment degree $\mathbf{B}(\alpha_t)$ of an explicit belief α_t is computed according to:

$$\mathbf{B}(\alpha_t) = \begin{cases} \frac{|pre(t)| - \lambda}{1 - \lambda} & \text{if } |pre(t)| > \lambda \\ 0 & \text{otherwise} \end{cases}$$

Document	Document Content	User Judgment
Doc1	java, technology, program	Relevant
Doc 2	java, technology, program	Relevant
Doc 3	java, volcano, program, technology	Relevant
Doc 4	java, volcano, program, technology	Relevant
Doc 5	java, volcano, program, technology	Relevant
Doc 6	program	Non-Relevant
Doc 7	computer, program, technology	Non-Relevant
Doc 8	computer, program, technology	Non-Relevant
Doc 9	computer, internet, program, technology	Non-Relevant
Doc 10	computer, internet, program, technology	Non-Relevant

Table 1 - An example of training documents

Table 1 depicts an example of ten training documents and their associated relevance judgments from an assumed user. Table 2 shows the result of applying the aforementioned preference induction method to the sample of training documents depicted in Table 1. The last column in Table 2 shows the derived entrenchment degrees associated with the beliefs representing some term preferences. In this example, the parameters: $|D^+| = |D^-| = 5$, $\varepsilon = 0.95$, $\lambda = 0.3$, $pos = 5$, and $neg = 5$ are assumed in the preference induction process. $|D^+|$ and $|D^-|$ represent the sizes of the sets of known relevant documents and non-relevant documents respectively. The parameters ε ,

pos, and *neg* are applied to the preference induction formula to control the maximal preference value, the weight of positive evidence, and the weight of negative evidence respectively. The parameter λ is used to filter out insignificant beliefs of an information seeker's preference. The parameters ε , λ , *pos*, and *neg* are estimated based on empirical evaluation. If a parameter value leads to satisfactory retrieval performance in the pilot runs, it will be adopted for the experiments related to that particular collection. As the contents of D^+ and D^- evolve according to an information seeker's changing preferences, the entrenchment degrees of the corresponding beliefs are raised or lowered in the user profile of the adaptive IR system. The changes applied to the epistemic entrenchment ordering of beliefs will then generate different nonmonotonic consequence relations which underpin the adaptive IR system's decisions about document relevance at various points of time.

Term	df(t_{rel})	df(t_{nrel})	pre(t)	α_t	B(α_t)
java	5	0	0.724	java	0.605
computer	0	4	-0.631	\neg computer	0.473
volcano	3	0	0.510	volcano	0.300
internet	0	2	-0.361	\neg internet	0.087
technology	5	4	0.266	-	-
program	5	5	0	-	-

Table 2 - An Example of Induced Epistemic Entrenchment

5. MINING CONTEXTUAL KNOWLEDGE

Contextual knowledge refers to the semantic relationships among concepts. This section illustrates the approaches we adopt to automatically derive two types of semantic relationships: information flow and information preclusion.

5.1 Information Flow Based Text Mining

Information flow computations through a high dimensional semantic space have been proposed as a means for computing term associations, both explicit and implicit. (Song and Bruza 2001; Song and Bruza 2003). Information flow is motivated from the conceptual level of cognition (Gärdenfors 2000). Encouraging results have been obtained with information flow based query expansion using a semantic space created by Hyperspace Analogue to Language (Bruza and Song 2002). In the belief revision based adaptive IR system, the information flow method is applied to discover initial contextual information.

Hyperspace Analogue to Language (HAL) is an exemplar of models emerging from cognitive science generally referred to as "semantic space" (Lowe 2001). These models run over a corpus of text and build representations of words in a (reduced) high dimensional space. The appeal of these models to IR related applications is that they have an encouraging track record of replicating human information processing, for example, semantic word association norms. IR has a long history of exploring term associations. Typically the underlying basis is probabilistic, or more specifically information theoretic, for example EMIM. Granted there has been notable pragmatic success computing term associations in this way, however the fundamental question

remains whether such associations actually accord with those the user would make. For this reason, we prefer to use semantic space models due to their cognitive compatibility, especially in relation to human word association norms. Two prominent semantic space models are Hyperspace Analogue to Language (HAL) (Burgess et. al. 1998) and Latent Semantic Analysis (LSA) (Landauer et. al. 1998). We prefer to use HAL as it does not involve the additional computational cost of dimensionally reducing the semantic space. Moreover, in the research cited above which employed HAL to underpin query expansion, the performance differential over two prominent probabilistic approaches was marked. Although the experiments reported do not settle the question whether the computation of term associations should be cognitively, or probabilistically motivated, there was more than enough encouragement to pursue further investigations with HAL in an IR setting.

HAL produces representations of words in a high dimensional vector space; these vector representations seem to correlate with the equivalent human representations. For example, word associations computed on the basis of HAL vectors seem to mimic human word association judgments (Burgess et al. 1998). HAL is “*a model that acquires representations of meaning by capitalizing on large-scale co-occurrence information inherent in the input stream of language*”. The space comprises high dimensional vector representations for each term in the vocabulary. Briefly, given an n -word vocabulary, the HAL space is a $n \times n$ matrix constructed by moving a window of length l over a corpus by one word increments ignoring punctuation, sentence and paragraph boundaries. All words within the window are considered as co-occurring with each other with strengths inversely proportional to the distance between them. After traversing the corpus, an accumulated co-occurrence matrix for all the words in a target vocabulary is produced. The co-occurrence matrix is added to its transpose to result in a symmetric matrix the rows of which are termed HAL vectors.

As an example of a HAL vector derived from a large corpus, consider part of the normalized HAL vector for “*superconductors*” computed from a corpus of Associated Press news:

```
superconductors = < U.S.:0.11 american:0.07 basic:0.11 bulk:0.13 called:0.15 capacity:0.08 carry:0.15  
ceramic:0.11 commercial:0.15 consortium:0.18 cooled:0.06 current:0.10 develop:0.12 dover:0.06  
electricity:0.18 energy:0.07 field:0.06 goal:0.06 high:0.34 higher:0.06 improved:0.06 japan:0.14 loss:0.13  
low:0.06 make:0.07 materials:0.25 new:0.24 require:0.09 research:0.12 researching:0.13 resistance:0.13  
retain:0.06 scientists:0.11 semiconductors:0.10 states:0.11 switzerland:0.06 technology:0.06  
temperature:0.48 theory:0.06 united:0.10 university:0.06>
```

This example demonstrates how a word is represented as a weighted vector whose dimensions comprise other words. The weights represent the strengths of association between “superconductors” and other words seen in the context of the sliding window: the higher the weight of a word, the more it has lexically co-occurred with “superconductors” in the same context(s). The quality of HAL vectors is influenced by the window size; the longer the window, the higher the chance of representing spurious associations between terms. Burgess *et al.* (1998) use a window size of eight or ten in their studies.

More formally, a concept² c_i is a vector representation: $c_i = \langle w_{c_i p_1}, w_{c_i p_2}, \dots, w_{c_i p_n} \rangle$ where p_1, p_2, \dots, p_n correspond to words in the vocabulary and are called *dimensions*, n is the dimensionality of the HAL space, and $w_{c_i p_i}$ denotes the weight of p_i in vector of c_i . A dimension is termed a *property* of concept c_i if and only if its weight is greater than zero. A property p_i of a concept c_i is termed quality property iff $w_{c_i p_i} > \hat{\partial}$, where $\hat{\partial}$ is a non-zero threshold value. Let $QP_{\hat{\partial}}(c)$ denote the set of quality properties of concept c . $QP_{\mu}(c)$ denotes the set of quality properties of concept c with above mean positive weight. (Mean positive weight is calculated as the mean of all dimensions greater than zero). For notational convenience, $QP(c)$ will be used to denote $QP_0(c)$. The latter notation simply denotes the set of dimensions of concept c with positive weight.

Concept combination is an important issue as combinations of words may represent a single underlying concept, for example, *space program*. An important intuition in concept combination is that one concept can dominate the other. For example, the term “space” can be considered to dominate the term “program” because it carries more of the information in the phrase. Given two concepts $c_1 = \langle w_{c_1 p_1}, w_{c_1 p_2}, \dots, w_{c_1 p_n} \rangle$ and $c_2 = \langle w_{c_2 p_1}, w_{c_2 p_2}, \dots, w_{c_2 p_n} \rangle$, the vector representation of the combined concept is denoted $c_1 \oplus c_2$. Dominance is assumed proportional to the inverse document frequency (IDF) of the concept in question. For example, “space” is deemed to dominate “program” as its IDF is higher than that of “program”. The combination of concepts represented by HAL vectors can be computed by a heuristic form of vector addition (Song and Bruza 2003). For the purposes of this article it is sufficient to bear in mind that the result of concept combination is a vector representation.

A HAL vector can be considered to represent the information “state” of a particular concept (or combination of concepts) with respect to a given corpus of text. The degree of information flow between “space program” and “satellites”, say, is directly related to the degree of inclusion between the respective information states represented by HAL vectors. Total inclusion leads to maximum information flow. Assuming a vector space, information flow is computed by a function; the domain of which is vector pairs (u, v) and with range $[0, 1]$. The vector u is denoted the *source* of the information flow and vector v the target. Intuitively the function tries to express the amount of how much of a source concept implies a target. For example, one would expect significant information flow between “space program” and “satellites”, denoted $space \oplus program \mid\text{-} satellites$. When information flow is above a certain threshold ω then an information flow is established between the source and target concepts:

$$t_1, \dots, t_m \mid\text{-} t_j \text{ iff } \text{flow}(\bigoplus_{i=1}^m c_i, c_j) > \omega$$

² The word “concept” is used somewhat loosely; it can be envisaged as “term” in the traditional IR sense. The word “concept” employed for the semantic space computed by HAL can be viewed as a computational approximation, albeit rather primitive, of human conceptual space. See (Gärdenfors 2000) for more details.

where c_i denotes the HAL vector of the concept t_i . (For ease of exposition, $\bigoplus_{i=1}^m c_i$ will be simplified to c_i because combinations of concepts are also concepts). The threshold ω is set empirically. The degree of flow is computed in terms of the ratio of intersecting quality properties of c_i and c_j to the number of quality properties in the source c_i

$$\text{flow}(c_i, c_j) = \frac{\sum_{p_l \in (QP_{\mu}(c_i) \wedge QP(c_j))} w_{c_i p_l}}{\sum_{p_k \in QP_{\mu}(c_i)} w_{c_i p_k}}$$

An information flow $t_1, \dots, t_m | - t_j$ is converted to the corresponding belief $(t_1 \wedge t_2 \wedge, \dots, \wedge t_m \rightarrow t_j, i)$. The entrenchment degree i is computed by multiplying $\text{flow}(\bigoplus c_i, c_j)$ with an adjustment factor. By way of illustration, after applying the information flow based text mining to the Reuters-21578 collection, it is found that the concept ‘‘NEC’’ exhibits strong information flow to other concepts such as (computer, 0.9415), (electronics, 0.8355), (Japan, 0.7623), etc. In this way, contextual rules such as (NEC \rightarrow Japan, 0.7623) is established.

5.2 Inducing Information Preclusion Relations

The other important semantic relationships *information preclusion* (Bruza and Huibers 1994) can be acquired from a corpus through supervised learning. An information preclusion relation such as $t_i \perp t_j$ indicates that a token t_i precludes another token t_j driven by an information seeker’s information needs. For example, $car \perp boat$ may hold if a user interested in documents about ‘‘car’’ is not interested in documents about ‘‘boat’’ with respect to a particular retrieval task. An information preclusion relation is represented by a rule $t_i \rightarrow \neg t_j$. For a term t from the set of positive tokens (e.g., by looking up the system table as depicted in Table 2), if $df(t_{rel}) > \gamma$ and $df(t_{nrel}) = 0$, the term t is added to the antecedent set L. Similarly, for a term t satisfying $df(t_{nrel}) > \gamma$ and $df(t_{rel}) = 0$, it is added to the consequent set R. Then, for each term $t_i \in L$, generate a rule $t_i \rightarrow \neg t_j$ for each $t_j \in R$. The entrenchment degree of such a rule is derived by: $\Pr(t_i) \times \Pr(t_j) \times \delta$. The adjustment factor δ is estimated based on empirical tests during the pilot runs of a document collection. In practice, the computations of both information flow and information preclusion are conducted offline in order to maintain good online information retrieval performance. As information flow is an unsupervised learning method (i.e., the class label for a document is not required), it can be done before interactive information retrieval takes place. However, the mining of information preclusion associations can only be conducted when certain number of training documents is available.

6. PROFILE REVISION AND SEMANTIC-BASED DOCUMENT MATCHING

Based on an information seeker’s explicit and implicit feedback, the belief revision based adaptive IR system can induce a set of beliefs representing the information seeker’s current information needs as described in Section 4. In addition, contextual knowledge is discovered via information flow based text mining. The second step of learning a possibly changing retrieval

situation is to revise the corresponding beliefs into the system's user profile (i.e., the knowledge base) via the belief revision processes. In general, the belief-based adaptive IR system will maintain a separate user profile for each individual information seeker who is interested in a specific topic. Furthermore, these profiles can be aggregated to form a generic profile for a group of information seekers who share similar interests (e.g., a default profile for new system users). The following examples demonstrate how user profile revision (i.e., learning) and information matching are conducted in the belief-based adaptive IR system.

6.1 Learning in Retrieval Situation RS_1

Table 3 illustrates the changes applied to the adaptive IR system's user profile (i.e., the knowledge base K) before and after a query is received. This example describes a retrieval situation that an information seeker with the general interest in science (e.g., perhaps a student with major in science subjects) issues the query “java”. Initial contextual knowledge is acquired before the query is received. The initial belief about the information seeker’s background e.g., (science, 0.9) is obtained via the dialog between the user and the system, and a default entrenchment degree is assigned to this relatively certain information. In addition, inference rule such as (java \wedge computer \rightarrow programming) is obtained by executing the information flow based text mining method. Information preclusion relationship such as (volcanology \rightarrow \neg computer) is induced according to the method described in Section 5.2. In this example, it is assumed that the information seeker has reviewed some documents before, and therefore a training set is available. As illustrated in this example, a retrieval context is not simply a set of terms as viewed by some researchers (Lawrence 2000; Cohen and Singer 1996), it has a richer representation which includes the relationships among terms triggered by particular retrieval tasks (Hirst and St-Onge 1998). One distinct advantage of the belief-based adaptive IR system is that the expressive power of the belief revision logic allows the semantics of a retrieval context to be represented, and thereby opening the door to more effective information matching.

Belief (α)	Before Revision	After Revision
science	0.900	0.900
volcanology \rightarrow \neg computer	0.830	0.830
java \wedge computer \rightarrow programming	0.713	0.713
java \wedge volcano \rightarrow merapi	0.713	0.713
science \wedge volcanology \rightarrow volcano	0.695	0.695
science \rightarrow computer	0.427	0.427
java	0.000	0.900
computer	0.427	0.427
\neg volcanology	0.427	0.427
programming	0.000	0.427
\neg computer	0.000	0.000
merapi	0.000	0.000
volcano	0.000	0.000

Table 3 – The First Retrieval Situation RS_1

To incorporate the user's current interest (i.e., a query) into the adaptive IR system's knowledge base, the Rapid Anytime Maxi-adjustment (RAM) operation $\mathbf{B}^*(\text{java}, 0.9)$ is invoked. After such a learning process, the system can automatically deduce that the user may be interested in documents about “computer programming”. The revised user profile will be used to match against incoming documents. In the next subsection, we propose an entrenchment-based similarity measure as the matching function. By viewing a recommended document, the user may eventually find that she is interested or not interested in the document. The adaptive IR system can then further revise its beliefs about the user's preferences, which will be illustrated by the Table 4 later, based on the user’s relevance feedback and the preference induction mechanism described in Section 4. The upper sections (above the double lines) of Table 3 and Table 4 list all the explicit beliefs, and the lower sections show all the implicit beliefs deduced based on the explicit beliefs stored in the system’s knowledge base. The second column in Table 3 and 4 indicates the entrenchment degrees of the corresponding beliefs before belief revision takes place, and the third column shows the entrenchment degrees of those beliefs after belief revision pertaining to a retrieval situation is conducted. The entrenchment degrees of the implicit beliefs are computed according to Definition 2 of Section 3. Only the implicit beliefs relevant for our discussion are shown in the tables. A belief with zero entrenchment degree is contracted from the knowledge base K . The beliefs with changing entrenchment degrees are highlighted in these tables.

6.2 Semantic Similarity Matching

Similarity measures are often used to induce document rankings for IR (Losada and Barreiro 1999; Salton and McGill 1983). An entrenchment-based similarity measure $Sim(RS, doc)$ is developed to approximate the *semantic correspondence* between a retrieval situation RS and an information object doc :

$$Sim(RS, doc) \approx Sim(B, d) \\ = \frac{\sum_{l \in d} [\text{degree}(B, l) - \text{degree}(B, \neg l)]}{|S|}$$

A retrieval situation refers to a user query and the associated retrieval context. It should be noted that the proposed similarity measure $Sim(RS, doc)$ not only considers the syntactic aspects in information matching but also takes into account the semantics among information items by means of nonmonotonic inference conducted according to the degree function $degree(\mathbf{B}, \alpha)$ defined in Definition 2 of Section 3. The above similarity measure combines the advantages of quantitative ranking and symbolic reasoning in a single formulation. It is not a simple overlapping model since the function $degree(\mathbf{B}, l)$ invokes nonmonotonic inference about the relevance of a document characterization d with respect to the knowledge base $content(\mathbf{B})$ which

represents a retrieval situation RS . The basic idea is that a document doc is characterized by a set of positive literals $d = \{l_1, l_2, \dots, l_n\}$. If the system's knowledge base $K = content(\mathbf{B})$ logically entails an atom l_i , a positive contribution is made to the overall similarity score because of the partial semantic correspondence between RS and doc . This kind of logical entailment is non-classical and is underpinned by the expectation inference relation $|\sim$ (Gärdenfors and Makinson 1994). On the other hand, if K implies the negation of a literal $l_i \in d$, it indicates certain semantic distance between RS and doc . Therefore, the similarity value is reduced by a certain degree. The set S of active information carriers partially characterizing the document doc is defined by $S = \{l_i \in d: degree(\mathbf{B}, l_i) > 0 \vee degree(\mathbf{B}, \neg l_i) > 0\}$. At the implementation level, both queries and the associated retrieval context are represented by the beliefs stored in the adaptive IR system's knowledge base K . With reference to the first retrieval situation RS_1 , if the following three documents are evaluated by the adaptive IR system, the result of document matching will be:

$$d_1 = \{\text{computer, programming}\}$$

$$d_2 = \{\text{volcanology, computer, programming}\}$$

$$d_3 = \{\text{merapi, volcano}\}$$

$$\text{Sim}(\mathbf{B}, d_1) = 0.427$$

$$\text{Sim}(\mathbf{B}, d_2) = 0.142$$

$$\text{Sim}(\mathbf{B}, d_3) = 0$$

$$\therefore doc_3 \leq doc_2 \leq doc_1$$

where $doc_i \leq doc_j$ means doc_j is at least as preferable as doc_i with respect to the retrieval situation. Such a ranking corresponds to our intuition about document preference with respect to the retrieval situation RS_1 .

6.3 Learning in Retrieval Situation RS_2

If the retrieval context is changed because the information seeker is actually a science student specializing in “volcanology”, the system's knowledge base K before and after incorporating such a contextual change is depicted in Table 4. The new information about the user's background is revised into K via the belief revision operation $\mathbf{B}^*(\text{volcanology}, 0.9)$. In this case, the entrenchment degree is a default value applied to the background information entered by the user.

Belief (α)	Before Revision	After Revision
science	0.900	0.900
java	0.900	0.900
volcanology \rightarrow \neg computer	0.830	0.830
java \wedge computer \rightarrow programming	0.713	0.713
java \wedge volcano \rightarrow merapi	0.713	0.713
science \wedge volcanology \rightarrow volcano	0.695	0.695
science \rightarrow computer	0.427	0.000
volcanology	0.000	0.900
\negcomputer	0.000	0.830
merapi	0.000	0.695
volcano	0.000	0.695
\negvolcanology	0.427	0.000
programming	0.427	0.000
computer	0.427	0.000

Table 4 – The Second Retrieval Situation RS_2

According to the RAM method, $(\mathbf{B}(\neg\text{volcanology}, 0))^+(\text{volcanology}, 0.9)$ is executed. As the belief $(\neg\text{volcanology}, 0.427)$ is implied by the knowledge base K , the belief revision operation must first lower the entrenchment degree of $(\neg\text{volcanology})$ to zero before adding the explicit belief $(\text{volcanology}, 0.9)$ into K such that the representation of the retrieval situation remains consistent and coherent. The belief $(\text{volcanology}, 0.9)$ then invokes a number of expansion operations leading to the insertions of implicit beliefs $(\text{Merapi}, 0.695)$, $(\neg\text{computer}, 0.830)$ and $(\text{volcano}, 0.695)$. Before $(\neg\text{computer}, 0.830)$ is inserted, $\mathbf{B}(\text{computer}, 0)$ has to be first executed. In doing so, the explicit belief $(\text{science} \rightarrow \text{computer}, 0.427)$ (i.e., the least entrenched belief) is contracted from the theory base $\text{exp}(\mathbf{B})$. If a dispatch threshold $disp$ is used, the system can make binary decision of document selection. A document d will be selected by the system if $\text{Sim}(\mathbf{B}, d) > disp$ is established. For instance, if $disp = 0.1$ is chosen, the system will select doc_1 and doc_2 for the user in the first retrieval situation, but select doc_3 in the second situation. The document ranking in retrieval situation RS_2 is as follows:

$$\text{Sim}(\mathbf{B}, d_1) = -0.830$$

$$\text{Sim}(\mathbf{B}, d_2) = 0.035$$

$$\text{Sim}(\mathbf{B}, d_3) = 0.695$$

$$\therefore \text{doc}_1 \leq \text{doc}_2 \leq \text{doc}_3$$

6.4 Explaining the System's IR Decision

Figure 3 is a screen shot of the belief-based adaptive IR system. It demonstrates how the belief-based adaptive IR system facilitates the generation of human comprehensible explanations to justify the system's document selection decisions. In this example, the content of the system's knowledge base (theory base) is shown at the lower right corner. This represents the system's beliefs about the current retrieval situation (i.e., query and the associated retrieval context). In particular, the hypothetical user's interest is about "Internet", and the contextual information includes $\{(internet \rightarrow softbot, 0.85), (softbot \rightarrow spider, 0.85), (spider \rightarrow crawler, 0.85), (crawler \rightarrow \neg music, 0.023)\}$. On the other hand, a document characterized by $d = \{internet, spider, music, mp3\}$, and the system's justifications of selecting that document are depicted at the upper left window. As illustrated by the explanation window, the IR system selects this document because the token "internet" is contained in the document and this is explicitly requested by the user. Moreover, the token "spider" appears in the document, and such a token is associated with the token "internet"; this relationship is dynamically discovered by applying the information flow mining to a set of documents recently browsed by the user (or against a document collection as conducted in our experiments). Accordingly, the token "spider" contributes a positive value to the overall document similarity score $Sim(RS, doc)$.

The screenshot shows the AIFS interface with an 'Explanation for document 1' window open. The window contains a table of evidence and a 'Belief' table. The 'Explanation for document 1' window has a title bar with a close button (X) and a scroll bar on the right. Below the title bar is a table with columns: Document, Certainty Degree, and Support or Counter Evidence. The table lists: internet (0.3, Item "internet" is requested), spider (0.3, Item "internet" is requested, and item "spider" is associated with it), music (0.023, Item "internet" is requested, which precludes item "music"), and mp3 (0.0, Neither support nor rejection). Below the table are statistics: No. of Supporting Evidence: 2, No. of Counter Evidence: 1, Degree of Semantic Correspondence: 0.192, and Document Threshold: 0.0. A conclusion states: 'The degree of semantic correspondence between the document and the retrieval context is high.' and a decision: 'Select document'. At the bottom of the window are buttons for 'Relevant', 'Not relevant', and 'Close'. Below the explanation window is a 'Belief' table with columns: Degree and Belief. The table lists: 0.85 (internet->softbot), 0.85 (softbot->spider), 0.85 (spider->crawler), 0.3 (internet), and 0.023 (crawler->-music). At the bottom right of the main window is a 'Save theory base' button.

Document	Certainty Degree	Support or Counter Evidence
internet	0.3	Item "internet" is requested
spider	0.3	Item "internet" is requested, and item "spider" is associated with it
music	0.023	Item "internet" is requested, which precludes item "music"
mp3	0.0	Neither support nor rejection

Degree	Belief
0.85	internet->softbot
0.85	softbot->spider
0.85	spider->crawler
0.3	internet
0.023	crawler->-music

Figure 3 – Explaining the System's Document Selection Decision

On the other hand, the token "music" from the document contributes a negative value (-0.023) to the overall document score because there is a small semantic gap between the user's information

need and part of the document as evidenced by the contextual information {(internet→softbot, 0.85), (softbot→spider, 0.85), (spider→crawler, 0.85), (crawler→¬music, 0.023)}. After reviewing the system’s judgment, the user can provide relevance feedback by clicking the “Relevant” or the “Not Relevant” button at the bottom of the document selection/explanation window. This action will trigger some belief revision operations to raise or lower the entrenchment degrees of the corresponding beliefs in the system’s knowledge base K (i.e., the user profile). Periodically, the training documents (documents judged by the user) will also be used to discover information preclusion relationships.

7. SYSTEM EVALUATION

7.1 General Experimental Procedures

The evaluation procedure of our belief revision based adaptive IR system (BR) is based on the adaptive information filtering benchmark task used in the seventh Text REtrieval Conference (TREC-7) (Hull 1998). A set of TREC topics was used to represent the diverse initial information needs of a hypothetical information seeker, and a relevance feedback file is employed to simulate human relevance feedback for documents. The TREC relevance feedback file consists of three fields, the first field is a TREC topic number, the second field contains the TREC-AP document ID, and the last field indicating whether a particular document is relevant (constant "1") or not (constant "0") with respect to a particular TREC topic. A stream of incoming documents (e.g., newswires) are presented to a filtering system; the system needs to make a binary decision of document disposal (e.g., accept or reject) immediately whenever a document arrives. In this sense, the adaptive filtering task is quite different from the traditional routing task where a batch mode of ranking operation is conducted (Robertson 2002). If the filtering system decides to retrieve a document, the relevance judgment information associated with the document is available to the system to revise a user profile. Our experimental procedure differs from the TREC-7 adaptive filtering procedure in that the relevance judgment information of each document is available to our adaptive IR system to revise a user profile after the document dissemination phase, whereas only the relevance judgment information of a retrieved document is available to a filtering system to revise a user profile in TREC-7. In addition, our belief-based adaptive IR system can make use of information flows to develop the appropriate retrieval context before adaptive filtering begins. In each experimental run, documents are filtered with respect to the specific interest (a single topic) of a hypothetical information seeker. We used the same performance measures (F1 and F3 utilities) as adopted in the TREC-7 adaptive filtering task to evaluate the system:

$$F1 = 3 \times |\text{Ret_Rel}| - 2 \times |\text{Ret_Nrel}|$$

$$F3 = 4 \times |\text{Ret_Rel}| - |\text{Ret_Nrel}|$$

where $|\text{Ret_Rel}|$ and $|\text{Ret_Nrel}|$ refer to the number of relevant and non-relevant documents retrieved by an IR system respectively. The larger a F1 or F3 score is, the better an IR system performs. In terms of the document collections, we have tested our system based on both the TREC-AP collection (Hull 1998) and the Reuters-21578 collection.

A baseline adaptive IR system (VS) was also developed based on the vector space model (Salton and McGill 1983) and the Rocchio learning method (Rocchio 1971). In the VS system, the following term weighting formula was used to compute the weight w_t of a term t (Salton 1991):

$$w_t = \frac{\left(0.5 + 0.5 \frac{tf_t}{\max tf}\right) \cdot \log_2 \frac{N}{N_t}}{\sqrt{\left(0.5 + 0.5 \frac{tf_t}{\max tf}\right)^2 \cdot \left(\log_2 \frac{N}{N_t}\right)^2}}$$

where tf_t is the occurrence frequency of term t in a document, and N_t and N represent the number of documents containing term t and the total number of documents in a collection respectively. The Rocchio learning method (Rocchio 1971) was used in the VS system:

$$\bar{Q}_{i+1} = \alpha \times \bar{Q}_i + \frac{\beta}{|rel|} \times \sum_{d \in rel} \frac{\vec{d}}{\|\vec{d}\|} - \frac{\gamma}{|nrel|} \times \sum_{d \in nrel} \frac{\vec{d}}{\|\vec{d}\|}$$

where \bar{Q}_{i+1} is the new prototype vector generated at time point $i + 1$, and \bar{Q}_i represents the original vector before learning takes place. The set of relevant (*rel*) and non-relevant (*nrel*) documents are those documents filtered by the VS system during a training cycle. The simulated human relevance judgment for each document (AP document or Reuters-21578 document) is obtained from the relevance feedback file. The notation \vec{d} represents a document vector of TFIDF weights and $\|\vec{d}\|$ is Euclidian vector length. The parameters $\alpha = 1$, $\beta = 0.75$, and $\gamma = 0.25$ were applied in this experiment.

Information flow mining was applied to the subset of a document collection (e.g., AP90 of the entire TREC-AP collection) to develop the initial retrieval context. For each topic, only 20 significant information flow relations were loaded into the BR system's knowledge base before an adaptive filtering task began. No information preclusion mining was applied to the experiments reported in this paper. To develop a comparable setting for the VS system, WordNet (Miller et. al. 1990) was used to expand each initial query (e.g., terms extracted from the title field of a TREC topic or from the topic description file of the Reuters-21578 collection). In particular, a maximum of 20 terms were added to the initial query based on the synonym sets (synsets) found in WordNet. When WordNet-based semantic query expansion was performed, the most common sense (as defined by WordNet) was consulted first, followed by the less common senses. For instance, the title “acquisitions” of TREC topic 002 was expanded with terms such as “contracting”, “getting”, “transferred possession”, “learning”, “accomplishment”,

“attainment”, etc. An expanded query was then loaded to the VS system before the adaptive filtering task was performed. All the experimental runs were based on the configuration of a single Pentium III 800MHz CPU with 256MB main memory. Both the BR and the VS system were developed using Java SDK 1.4. For the experiments reported in this paper, the parameters: $\varepsilon = 0.95$, $\lambda = 0.3$, $pos = 5$, and $neg = 1000$ were applied to the BR system. These parameters were estimated based on empirical evaluation. For example, several topics from each collection were randomly selected, and then various parameter values were tried to obtain the reasonable filtering performance. The best combination of parameter values in these pilot runs were then applied to the BR system in subsequent runs.

Last but not least, document dissemination thresholding is also a challenging task in adaptive information filtering. Nevertheless, an optimal thresholding strategy often depends on the particular performance measure (Arampatzis and van Hameren 2001; Zhang and Callan 2001). In order to retain the computational efficiency of our logic-based adaptive IR system, we adopted the simple thresholding strategy by tracking the average document scores between the set of relevant documents and the set of non-relevant documents, and then set the threshold somewhere between these two points at each learning cycle (Callan 1998). More specifically, simple heuristic rules were used to calibrate our dissemination threshold. For the typical situation where the average document score of the relevant documents (Avg^+) is higher than that of the non-relevant documents (Avg^-), the dissemination threshold is estimated according to: $Avg^+ + [(Avg^+ - Avg^-) \times \theta]$, where θ is an adjustment factor and is set to -0.25 for our experiments. On the other hand, if the average document score of the non-relevant documents is greater than or equal to the average document score of the relevant documents at a particular learning cycle, our system will take a conservative decision by not altering the threshold at all. According to our observation, this situation could happen during the early stages of filtering. The initial dissemination threshold for a topic was established based on empirical testing. The same thresholding strategy was applied to both the BR and the VS systems respectively. We believe that employing more sophisticated dynamic thresholding method may improve the performance of our belief-based adaptive IR system. However, we will leave this task as part of our future work.

7.2 Evaluation Based on the TREC-AP Collection

The TREC-AP comprises the Associated Press (AP) newswires covering the period from 1988 to 1990 with a total number of 242,918 documents (728 MB). As with the TREC-7 adaptive filtering task, a set of 50 topics (from topic 001 to 050) were used to represent the diverse initial information needs of a hypothetical user. An example of the TREC topic 008 is shown below:

<top>

<head> Tipster Topic Description

<num> Number: 008

<dom> Domain: International Economics

<title> Topic: Economic Projections

<desc> Description:

Document will contain quantitative projections of the future value of

some economic indicator for countries other than the U.S.</top>

A TREC topic contains several fields with each field marked up by the corresponding tag. The TREC topic number field is marked up by the <num> tag. Our experiments only used the title field marked up by the <title> tag to represent the initial query of a particular filtering run. Exactly the same relevance feedback file of TREC-7 was used in our experiment. Table 5a and Table 5b show the results of our runs for the TREC topics 001 – 050. The first column in these tables identifies the TREC topic number and the second column shows the number of true relevant documents judged by the TREC forum; the remaining three columns show the BR system's performance. The last three columns depict the performance difference between the BR system and the VS system. A positive figure in columns 6 or 7 indicates that the BR system outperforms the VS system. The last column shows how many seconds more (a positive number) are consumed by the BR system. As either information flow mining or WordNet based query expansion is conducted before the adaptive filtering processes take place, their computational time is not included in Table 5a and Table 5b.

BR's Performance					BR vs. VS		
Topic	Rel.Doc.	F1	F3	Time	Δ F1	Δ F3	Δ Time
1	344	-4	18	50452	-66	-153	32899
2	459	-227	14	45361	-10	-130	24216
3	220	-94	-2	50083	-60	-35	35252
4	94	-39	-7	35012	74	-98	22390
5	68	7	21	24879	-39	-62	12249
6	270	-96	-13	38635	-335	-530	23927
7	311	-5	0	19572	192	-54	2490
8	66	-32	-16	28064	176	66	15324
9	107	-48	-19	31424	-22	-91	18444
10	265	370	660	71335	586	403	55939
11	427	-23	151	46962	858	-161	30299
12	624	-6	22	27370	-103	-149	2097
13	86	51	88	24760	-96	-138	12208
14	116	-1	12	23274	22	-39	9742
15	107	-87	-41	35985	146	-52	22615
16	142	-69	23	36756	-112	-46	23285
17	224	-73	151	67759	-248	-89	52111
18	130	-41	-13	26162	1581	1550	13222
19	269	-3	6	25886	25	-50	9829
20	158	-53	6	37530	-119	-192	23851
21	29	-109	-27	34365	-49	-79	22519
22	1238	-38	496	75793	-16	132	42735
23	237	250	445	37725	-72	9	22628
24	300	-111	92	48147	-107	-11	30979
25	65	-43	-9	32608	60	-65	20133

Table 5a – BR vs. VS based on TREC-AP collection (Topics 1 -25)

BR's Performance					BR vs. VS		
Topic	Rel.Doc.	F1	F3	Time	Δ F1	Δ F3	Δ Time
26	61	-25	0	25315	49	-13	13205
27	20	-5	5	40551	-8	1	28876
28	66	-22	4	37944	-8	1	25685
29	7	-4	-2	23646	12	6	12248
30	1	0	0	18845	2	1	7576
31	1	-22	-11	23449	8	4	12166
32	6	-2	-1	22858	-2	-1	11496
33	13	0	0	22917	-1	-3	11328
34	2	-8	-4	24293	-5	-6	12964
35	1	-2	-1	20848	5	0	9578
36	10	-7	4	27193	-7	-6	15757
37	7	11	18	19665	36	23	8282
38	276	-12	4	25000	-13	-4	7703
39	4	-4	-2	20103	19	1	8808
40	118	-29	-12	26093	309	-243	11092
41	30	-3	1	25459	0	0	13015
42	151	-33	-14	38163	62	11	22969
43	102	-133	-64	29452	-146	-83	16153
44	152	-4	-2	21682	7	-59	8422
45	52	-135	-50	46152	21	-7	33236
46	74	-7	4	25081	-110	-175	12715
47	89	0	0	23135	-10	-25	10058
48	30	4	7	21602	49	27	9686
49	55	-8	-4	19266	12	-59	6931
50	6	-11	-3	30170	-11	-3	18635

Table 5b – BR vs. VS based on TREC-AP collection (Topics 26 -50)

The mean and standard deviation of the F1 scores achieved by the BR system over the 50 TREC topics are -19.7 and 84.2 respectively. The mean and standard deviation of the F3 scores achieved by the BR system are 38.7 and 133.9 respectively. On the other hand, the mean and standard deviation of the F1 scores achieved by the VS system over the 50 TREC topics are -113.5 and 544.2 respectively. The mean and standard deviation of the F3 scores achieved by the VS system are 49.9 and 270.8 respectively. As a matter of fact, each TREC topic exhibits unique semantic features such as *clarity* (Cronen-Townsend et. al. 2002), and therefore retrieval effectiveness may vary across topics. The clarity score measures the degree of ambiguity of a query with respect to a collection of documents, and it is computed based on the relative entropy between the query and the collection language models (Cronen-Townsend et. al. 2002). It seems that a topic by topic comparison between the two systems will lead to a more meaningful analysis. There are 24 out of the 50 TREC topics in which the BR system performs better than or equal to the VS system in terms of the F1 score. However, in terms of the F3 score, there are only 16 topics in which the BR system performs better than or equal to the VS system. By

testing the hypotheses $H_{\text{null}}: \mu_{\text{BR}} - \mu_{\text{VS}} = 0$ and $H_{\text{alternative}}: \mu_{\text{BR}} - \mu_{\text{VS}} > 0$ with paired one tail t -test on the F1 scores, the null hypothesis is rejected at 11% level of significance ($t(49) = 1.218$, $p = .11$). However, in terms of the F3 scores, a significant statistical difference between the BR system and the VS system cannot be established. Therefore, a conclusive statement cannot be made if the filtering performance of the BR system is better than that of the VS system based on the TREC-AP collection. However, their filtering performance is comparable.

The reason why the BR system outperforms the VS system for some TREC topics is that inference about a retrieval context can be conducted by the BR system even though only a short query is available to the system initially (e.g., only a few terms extracted from the TREC title field is made available). For example, the following information flows for TREC topic 008 (Economic Projections) are obtained via text mining. These information flows are revised into the BR system's knowledge base before adaptive filtering takes place:

economic \oplus projections |– (administration:0.96 growth:0.90 percent:0.87 budget:0.80
forecast:0.78 deficit:0.76 inflation:0.73 national:0.70 analysts:0.68 spending:0.67 drought:0.65
sprinkel:0.65 prices:0.63 outlook:0.63 rates:0.61 interest:0.61 optimistic:0.60 federal:0.59
estimates:0.59 exports:0.56)

The notation |– indicates the flow of information from the left to the right. As explained in Section 5.1, the above information flows are translated to explicit beliefs such as ((economic \wedge projections) \rightarrow growth, 0.9), ((economic \wedge projections) \rightarrow forecast, 0.78), ((economic \wedge projections) \rightarrow inflation, 0.73), etc. In fact, all the terms are stemmed by our systems. Nevertheless, for readability reason, the non-stemmed terms are shown as examples in this paper. As shown in this example, only twenty significant information flows of a topic are applied to the BR system.

A concrete example of how belief-based informational inference taking place in the BR system follows: the TREC-AP document (AP880319-0174 – “New Financial Help Given Argentina”), which is judged relevant for the TREC topic 008 by the TREC experts, is retrieved by the BR system even though the terms “economic projections” do not appear in the document. The reason is that the positive beliefs about the information seeker's needs such as (budget, 0.8), (forecast, 0.78), (deficit, 0.76), (inflation, 0.73), (interest, 0.61), etc. are inferred by the BR system based on the given information flows, and these beliefs just happen to be the characterization of the particular document (i.e., the corresponding terms are found in the document). On the other hand, the same document cannot be retrieved by the VS system because there is no overlapping between the query terms (“economic projections”) and the terms characterizing the document.

It should be noted that the VS system can also learn the relevant terms such as “budget”, “forecast”, etc. and revising these terms into its user profile after a period of learning. However, the VS system cannot work as effective as the BR system does during the initial period of

filtering because the VS system cannot conduct context-sensitive query expansion. One may wonder why query expansion based on WordNet (Miller et. al. 1990) does not help in this case. For instance, the term “projection” has nine senses in WordNet (version 1.6). The first synonym set (the most common sense) is “prediction, anticipation, foresight”, the second synonym set is “visual communication”, and the third synonym set is “plan, program, programme”, etc. Nevertheless, none of these senses leads to an expanded query containing terms such as “forecast”, “budget”, “inflation” which are relevant with respect to the financial domain. In general, WordNet is a generic lexicon and it is very difficult to perform *context-sensitive* query expansion based on it. The proposed belief-based adaptive IR system (BR) is able to conduct context-sensitive informational inference to enhance retrieval effectiveness. However, the current information flow mechanism is not perfect. As can be seen, the information flow ((economic \wedge projections) \rightarrow percent) is a very general implication, and it may lead to some non-relevant documents to be retrieved. Such an information flow may be regarded as the noise generated from the text mining process. This weakness partly explains why there is still a large variation in terms of the performance of the BR system across the TREC topics.

Another interesting observation is about TREC topic 018 – “Japanese Stock Market Trends”. Although there is a significant boost on retrieval effectiveness, the negative F1 and F3 scores as achieved by the BR system are not really outstanding. After careful examination of the retrieval results of the VS system, we find that quite a number of non-relevant documents which are about stock markets in countries other than Japan or stock information rather than market trend are retrieved by the VS system. As a result, the utility scores of the VS system is much lower than that of the BR system. Although a very general term such as “market” will probably contribute zero to the overall document score in the VS system, other query terms (terms contained in the TREC title field) such as “stock” contributes a positive document score. Since a relatively low startup dissemination threshold and a primitive threshold updating strategy are used in our systems, a document just about “U.S. stock market” will be assigned a document score slightly higher than the dissemination threshold by the VS system. On the other hand, a noisy term such as “market” will make a negative contribution to the overall document score in the BR system according to our semantic information matching function defined in Section 6.2. Therefore, a document will not be selected by the BR system unless it contains quite a number of positive terms such as “Nikkei”, “Japan”, etc. As a result, fewer mistakes are made by the BR system for this relatively noisy topic.

As a whole, the performance of the BR system is encouraging although a conclusive statement about the retrieval effectiveness of the BR system cannot be made due to the high variances of the differences between the BR and the VS systems across the TREC topics. In general, the BR system is likely to perform well in a TREC topic if the information flows can accurately characterize the particular retrieval context of that topic. However, more research work is required to identify the exact conditions under which the BR system will outperform the VS system. The BR system spent 32,695.6 seconds to process a TREC topic and 0.135 second for a document on average. The VS system spent 14,096.3 seconds to process a TREC topic and 0.058 second for a document on average. Although the VS system is faster than the BR system, the belief revision based IR system is still remarkably efficient even for processing such a large document collection. According to the system log, the most computationally expensive

operations of the BR system are the belief revision operations. Basically, each belief revision operation invokes the underlying theorem prover (SICStus Prolog) to check if the dominance property (EE2) is maintained in the BR system's knowledge base. Several optimization techniques were applied to the BR system to make it scalable. For instance, only the relatively entrenched beliefs are selected for the belief revision processes since each belief revision operation is computationally expensive. Moreover, a belief revision operation will not be invoked unless the change of entrenchment degree of a belief is greater than a pre-defined system threshold. The frequency of invoking a learning cycle where term preference value computation and belief revision take place is also under tight control. Finally, by using the "anytime" algorithm for the implementation of belief revision, the elapsed time of each learning cycle will not exceed a pre-defined time limit.

7.3 Evaluation Based on the Reuters-21578 Collection

The BR system was also evaluated based on the Reuters-21578 collection with the Lewis-Split subset which contains 19,813 documents (13 MB). An example of a Reuters-21578 document is shown below:

```
<REUTERS TOPICS="YES" LEWISSPLIT="TRAIN" CGISPLIT="TRAINING-SET"
  OLDID="5544" NEWID="1">
<DATE>26-FEB-1987 15:01:01.79</DATE>
<TOPICS><D>cocoa</D></TOPICS>
<TEXT>
<TITLE>BAHIA COCOA REVIEW</TITLE>
<BODY>Showers continued throughout the week in
the Bahia cocoa zone, alleviating the drought since early
.....</BODY></TEXT>
</REUTERS>
```

A Reuters-21578 document contains many fields with each field marked up by a pair of tags. For instance, the main body of the news is delimited by the <BODY> and </BODY> tags, and the Topic field is delimited by the <TOPICS> and </TOPICS> tags. To apply the TREC adaptive filtering task to the Reuters-21578 collection, a collection pre-processing procedure is required. For instance, we need to construct a relevance feedback file similar to the one used in TREC-7. The topic field of each newswire document is parsed to extract the topic codes representing particular Reuters-21578 topics. For example, if the commodity code "cocoa" is found in the <Topics> field, a relevance feedback record will be created. The relevance feedback record consists of the Reuters-21578 record ID, the system generated topic ID corresponding to the topic code, and a constant "1" indicating that the document is relevant for the particular topic. Such a file format is exactly the same as the relevance feedback file used in TREC-7. Since the <Topics> field may contain more than one topic code, multiple relevance feedback records could be generated based on one Reuters-21578 document. If a combination of record ID and topic ID

cannot be found from the relevance feedback file, the corresponding Reuters document is considered non-relevant for the particular topic.

The Reuters-21578 topic file (part of the Reuters-21578 collection) which consists of 135 pre-defined topics was used to construct the initial queries. Each topic code retrieved from the topic file is used to extract the corresponding topic description from the Reuters-21578 category description file. For instance, the corresponding topic description for the corporate code “acq” is “acquisitions”. This topic description is used to develop the initial query (user profile) for that particular Reuters-21578 topic. Usually, a topic description consists of one or two terms. Only the first 20 topics (e.g., from “acq” to “corn”) were used in our experiment. Similar to the TREC-AP based experiments, an initial query in the BR system was contextualized by the information flows which were extracted before the adaptive filtering task took place. As for the VS system, each initial query was expanded with terms found from the synsets of WordNet. After collection pre-processing, the topic field of each Reuters-21578 document is removed, and so the topic codes are not available to our adaptive IR systems during the adaptive filtering processes. The F1 and F3 scores achieved by the BR system and the comparison between the BR and the VS systems are depicted in Table 6. There are 14 out of the 20 Reuters topics (i.e., more than a half) in which the BR system performs better than or equal to the VS system in terms of the F1 or the F3 scores.

BR's Performance					BR vs. VS		
Topic	Rel.Doc.	F1	F3	Time	Δ F1	Δ F3	Δ Time
acq	2366	2096	3673	21142	848	1489	20109
alum	57	24	52	734	-24	-22	405
austdlr	4	0	0	408	0	0	142
austral	0	-6	-3	397	2	1	147
barley	51	120	190	411	43	44	100
bfr	0	0	0	362	0	0	112
bop	105	0	16	612	-50	-134	244
can	3	-29	-12	385	-17	-6	119
carcass	68	26	58	918	17	41	546
castor-meal	0	0	0	386	0	0	134
castor-oil	2	0	0	399	0	0	143
castorseed	1	3	4	2251	0	0	1987
citruspulp	1	3	4	2369	2	1	2110
cocoa	73	174	262	726	7	31	378
coconut	6	-22	4	433	-34	-17	172
coconut-oil	7	0	0	447	11	-2	178
coffee	139	267	476	3853	-40	60	3424
copper	65	111	218	734	23	84	387
copra-cake	3	-4	-2	439	-4	-2	175
corn	237	472	801	1351	141	333	903

Table 6 – BR vs. VS based on Reuters-21578 collection

The means and standard deviations of the F1 and the F3 scores achieved by the BR system over the 20 Reuters topics are ($\mu=161.8$, $\sigma=471.5$) and ($\mu=287.1$, $\sigma=823.0$) respectively. On the other hand, the means and standard deviations of the F1 and the F3 scores achieved by the VS system are ($\mu=115.5$, $\sigma=284.8$) and ($\mu=192.0$, $\sigma=489.1$) respectively. By testing the hypotheses $H_{\text{null}}: \mu_{\text{BR}} - \mu_{\text{VS}} = 0$ and $H_{\text{alternative}}: \mu_{\text{BR}} - \mu_{\text{VS}} > 0$ with paired one tail t -test on the F3 scores, the null hypothesis is rejected at 11% level of significance ($t(19) = 1.26$, $p = .11$). However, the statistical difference between the BR system and the VS system in terms of the F1 scores can only be established at 15% level of significance ($t(19) = 1.07$, $p = .15$). Although, a conclusive statement cannot be made if the filtering performance of the BR system is better than that of the VS system, their filtering performance is once again comparable. The BR system spent 1,937.9 seconds to process a Reuters-21578 topic and 0.098 second for a document on average. The VS system spent 342.1 seconds to process a topic and 0.017 second for a document on average.

Similar to the situation when adaptive filtering was applied to the TREC-AP collection, context-sensitive information inference is the main reason why the BR system outperforms the VS system for some Reuters-21578 topics. For example, the following information flows for the topic “acq” (acquisitions) were extracted from the training set of the Reuters-21578 collection:

acquisitions |– (company:0.871 corp:0.845 dlrs:0.842 offer:0.714 shares:0.709 stock:0.684 merger:0.625 bank:0.614 shareholders:0.585 agreement:0.583 buy:0.577 tender:0.577 board:0.571 purchase:0.558 international:0.552 proposed:0.549 cash:0.541 completed:0.540 debt:0.538 bid:0.523)

The BR system can make use of the beliefs such as (acquisitions \rightarrow company, 0.871), (acquisitions \rightarrow offer, 0.714), (acquisitions \rightarrow shares, 0.709), (acquisitions \rightarrow stock, 0.684), (acquisitions \rightarrow merger, 0.625), etc. to contextualize the initial query “acquisitions”. For the Reuters newswire document (96 – “Investment Firms Cut Cyclops Stake”), neither the term “acquisitions” nor the synonyms extracted from WordNet appears in the document. However, the document is characterized by terms such as “shares” and “stock”. Therefore, the BR system was able to retrieve the document. On the other hand, the VS system was not able to retrieve the same document because there was no overlapping between its query vector and the document vector during the initial period of adaptive filtering. Similarly, document (128 – “Liebert Corp Approves Merger”) is characterized by terms such as “merger”, “shareholders”, “stock”, “shares”, etc. rather than the term “acquisitions”. The BR system was able to retrieve this document because the positive beliefs (shares, 0.709), (merger, 0.625), (shareholders, 0.585), etc. were inferred by the system to represent the current retrieval situation. On the other hand, the same document was rejected by the VS system since there was no overlapping between its query vector and the document vector.

7.4 Usability Study for the Belief-based Adaptive IR System

In order to evaluate the explanatory power offered by the belief-based adaptive IR system, an experiment was conducted to compare the explanation capability of the BR system and that of the VS system respectively. The explanatory power of an IR system refers to the ability of the

system to provide some cues about its IR decisions. Traditional quantitative IR systems only provide document scores (e.g., retrieval status values RSV) as a means of justifying their IR decisions. However, this black-box IR approach may not be sufficient to support mission critical IR activities (e.g., retrieving information about possible terrorist attacks) because the information seekers would like to see more justifications about the IR results to avoid paying the high cost of a false alarm.

A group of twenty college students taking the elective course of Web Intelligence participated in this experiment which is part of the usability study for our belief-based adaptive IR system. These students attended six weeks' lectures in IR theories and techniques before participating in the usability experiment. One group of students used the BR system (i.e., the experimental group) to retrieve documents from a small collection, and the another group (i.e., the control group) used the VS system to conduct the same task. At the beginning of the experiment, ten subjects were randomly chosen and assigned to the experimental group, and the remaining subjects were assigned to the control group. The test collection comprised fifty ACM technology news articles (TechNews) covering various topics in IT; standard document pre-processing procedures such as stop word removal and stemming were applied. The ACM TechNews are delivered to its subscribers three times a week. An example of TechNews is shown below:

"Softbots Stride Forward - The rapid progress of intelligent agent, or "softbot," technology has scientists convinced that people who do not have time to participate in conferences or perform business activities will be able to employ computer-generated avatars that act on their behalf within a decade. ..."

Both the user profiles of the BR system and the VS system were initialized with the topic "intelligent agents". A training set of TechNews, which was different from the test set, was used to develop the contextual knowledge about the query. When the adaptive filtering task began, each subject used their assigned IR system to retrieve documents and reviewed the system's justifications of document selection. For the BR system, the user interface consists of an explanation window as depicted in Figure 3. The user interface of the VS system is similar to that of the BR system except that its explanation window only displays a single document selection score. After reviewing all the fifty documents, a questionnaire was distributed to each participant to collect information about their perceived usability of the IR systems. From among the ten response items appearing on the questionnaire, three items are related to the dependent variable of "explanatory power" of an IR system, and the variable is measured using a 5-point semantic differential scale from strongly agree (5), agree (4), neutral (3), disagree (2), to strongly disagree (1). The three question items are: (a) "The IR system provides the cue why a document is retrieved"; (b) "The IR system provides an explanation of document selection"; (c) "The IR system can justify why a document is selected". The response rate of the study was 100%. The overall mean scores and standard deviations of the three question items are ($\mu_{BR}=4.03$, $\sigma_{BR}=0.21$) and ($\mu_{VS}=2.17$, $\sigma_{VS}=0.35$) for the BR system and the VS system respectively. By testing the hypotheses $H_{null}: \mu_{BR} - \mu_{VS} = 0$ and $H_{alternative}: \mu_{BR} - \mu_{VS} > 0$ with paired one tail t -test, the null hypothesis was rejected ($t(2) = 8.54$, $p < .01$). Therefore, we conclude that the explanatory power of the BR system is higher than that of the VS system according to this usability study.

8. CONCLUSIONS AND FUTURE WORK

The AGM belief revision logic provides a sound and rigorous framework to develop adaptive IR systems. In particular, the logic offers sufficient expressive power to represent IR contexts, and it also provides a sound inference mechanism to model the nonmonotonicity of information matching arising in changing retrieval contexts. On the other hand, information flow based text mining allows IR systems to discover contextual IR knowledge autonomously. The induction power brought by information flow based text mining is complementary to the nonmonotonic reasoning capability offered by the belief revision system. One distinct advantage of the belief-based adaptive IR system is that its learning and matching behavior can be understood based on the axioms characterizing the AGM logic. Our experiments show that the belief-based adaptive IR system achieves comparable performance as a classical adaptive IR system for the TREC-AP and the Reusters-21578 collections respectively. The belief-based adaptive IR system is efficient enough to deal with large and complex IR tasks, and its explanatory power is confirmed according to our usability study. To our best knowledge, this is the first comprehensive empirical evaluation of a logic-based adaptive IR model based on large IR benchmark collections and yet encouraging performance is observed. Future work includes selectively (e.g., based on the clarity score) applying text mining to enrich a retrieval context so that overall retrieval effectiveness can be improved. Moreover, more sophisticated thresholding techniques will be exploited to bootstrap the performance of the belief-based adaptive IR system.

ACKNOWLEDGMENTS

The work reported in this paper has been funded in part by the Co-operative Research Centre for Enterprise Distributed Systems Technology (DSTC) through the Australian Federal Government's CRC Programme (Department of Education, Science and Training). This work is also funded in part by the UK's Engineering and Physical Sciences Research Council (EPSRC), grant number EP/E002145/1 and Overseas Travel Grant number EP/E037402/1. The authors would also like to thank the three anonymous reviewers for valuable comments.

REFERENCES

- Alchourrón, C.E., Gärdenfors, P. and Makinson, D. (1985) On the logic of theory change: partial meet contraction and revision functions. *Journal of Symbolic logic*, 50:510-530.
- Amati, G. and Georgatos, K. (1996). Relevance as deduction: a logical view of information retrieval. *Proceedings of the 2nd Workshop on Information Retrieval, Uncertainty and Logic*, pp. 21-26.
- Arampatzis, A. and van Hameren, A. (2001) The Score-Distributional Threshold Optimization for Adaptive Binary Classification Tasks. *Proceedings of the 24th annual international ACM SIGIR conference on Research and development in information retrieval*, pp.285-293.
- Belkin, N. and Croft, W. (1992) Information filtering and information retrieval: two sides of the same coin? *Communications of the ACM*, 35(12):29-38.
- Bessant, B., Gregoire, E., Marquis, P., and Saddotis, L. (1998). Combining nonmonotonic reasoning and belief revision: A practical approach. *Proceedings of the 8th International Conference on Artificial*

- Intelligence: Methodology, Systems, and Applications (AIMSA-98)*, volume 1480 of *LNAI*, pp. 115 – 128.
- Bruza, P.D. and Song, D. (2002) Inferring Query Models by Information Flow Analysis. *Proceedings of the 11th International ACM Conference on Information and Knowledge Management (CIKM 2002)*, pp.260-269.
- Bruza, P., Song, D., and Wong, K. (2000) Aboutness from a commonsense perspective. *Journal of the American Society for Information Science*, 51(12):1090-1105.
- Bruza, P. and van Linder, B. (1998) Preferential Models of Query by Navigation. In Lalmas M., Crestani F., and van Rijsbergen C., editors, *Information Retrieval: uncertainty and logics*, pp. 73-96. Kluwer Academic Publishers, Norwell, Massachusetts.
- Bruza, P. and Huibers, T. (1994) Investigating Aboutness Axioms using Information Fields. *Proceedings of the 17th Annual International Conference on Research and Development in Information Retrieval*, pp. 112-121.
- Burgess, C., Livesay, K. and Lund, K. (1998): Explorations in context space: words, sentences, discourse. *Discourse Processes*, 25:211-257.
- Callan, J. (1998) Learning while filtering documents. *Proceedings of the 21st annual international ACM SIGIR conference on Research and development in information retrieval*, pp.224-231.
- Card, S., Moran, T., and Newell, A. (1983) *The Psychology of Human Computer Interaction*. Lawrence Erlbaum, Hillsdale, N.J.
- Chiararella, Y. and Chevallet, J. (1992) About retrieval models and logic. *The Computer Journal*, 35(3):233-242.
- Cohen, W. and Singer, Y. (1996). Context-sensitive learning methods for text categorization. *Proceedings of the 19th ACM SIGIR International Conference on Research and Development in Information Retrieval*, pp. 307 – 315.
- Cool, C. (2001) The Concept of Situation in Information Science. *Annual Review of Information Science and Technology (ARIST)*, 35:5-42.
- Cool, C. and Spink, A. (2002) Issues of context in information retrieval (IR): an introduction to the special issue. *Information Processing and Management*, 38(5):605-611.
- Crestani, F. and van Rijsbergen, C. (1995) Information retrieval by logical imaging. *Journal of Documentation*, 51(1):3-17.
- Crestani, F. (1998) Logical imaging and probabilistic information retrieval. In Lalmas M., Crestani F., and van Rijsbergen C., editors, *Information Retrieval: uncertainty and logics*, pp. 247-279. Kluwer Academic Publishers, Norwell, Massachusetts.
- Cronen-Townsend, S., Zhou, Y., and Croft, W. (2002) Predicting Query Performance. *Proceedings of the 25th Annual International Conference on Research and Development in Information Retrieval*, pp. 299-306.
- Dalal, M. (1988) Investigations into a theory of knowledge base revision: preliminary report. *Proceedings of the 7th National Conference on Artificial Intelligence*, pp. 475-479.
- Dretske, F. (1981) *Knowledge and the Flow of Information*. The MIT Press, Cambridge, Massachusetts.
- Dubois, D. and Prade, H. (1991) Epistemic entrenchment and possibilistic logic. *Artificial Intelligence*, 50(2):223-239.
- Dubois, D. and Prade, H. (1995) Possibility Theory as a Basis for Qualitative Decision Theory. *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, pp. 1924-1932.

- Gärdenfors, P. (1988) *Knowledge in flux: modeling the dynamics of epistemic states*. The MIT Press, Cambridge, Massachusetts, 1988.
- Gärdenfors, P. and Makinson, D. (1988) Revisions of knowledge systems using epistemic entrenchment. *Proceedings of the Second Conference on Theoretical Aspects of Reasoning About Knowledge*, pp. 83-95.
- Gärdenfors, P. and Makinson, D. (1994) Nonmonotonic inference based on expectations. *Artificial Intelligence*, 65(2):197-245.
- Gärdenfors, P. (2000) *Conceptual Spaces: The Geometry of Thought*. The MIT Press, Cambridge, Massachusetts.
- Hirst, G. and St-Onge, D. (1998) Lexical Chains as Representations of Context for the Detection and Correction of Malapropisms. In Fellbaum, C., editor, *WordNet: an electronic lexical database*, pp.305-332. MIT Press, Cambridge, Massachusetts.
- Huibers, T. and van Linder, B. (1996) Formalising Intelligent Information Retrieval Agents. *Proceedings of the 18th British Computer Society Annual Information Retrieval Colloquium*, pp. 125-143.
- Huibers, T. and Wondergem, B. (1998) Towards an axiomatic aboutness theory for information retrieval. In Lalmas M., Crestani F., and van Rijsbergen C., editors, *Information Retrieval: uncertainty and logics*, pp. 297-318. Kluwer Academic Publishers, Norwell, Massachusetts.
- Hull, D. (1998) The TREC-7 Filtering Track: Description and Analysis. *Proceedings of the 7th Text Retrieval Conference*, pp. 33-56.
- Hunter, A. (1996) Intelligent text handling using default logic. *Proceedings of the IEEE Conference on Tools with Artificial Intelligence*, pp.34-40.
- Hunter, A. (1997). Using default logic for lexical knowledge. In Dov Gabbay, editor, *Proceedings of the first International Joint Conference on qualitative and quantitative practical reasoning, volume 1244 of Lecture Notes in Computer Science*, pp. 322-335. Springer-Verlag.
- Ingwersen, P. and Järvelin, K. (2005) SIGIR workshop reports: Information retrieval in context: IRiX. *ACM SIGIR Forum*, 39(2):31-39.
- Kindo, T., Yoshida, H., Morimoto, T. and Watanabe, T. (1997) Adaptive personal information filtering system that organises personal profiles automatically. *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence*, pp. 716-721.
- Koller, D. and Sahami, M. (1997) Hierarchically classifying documents using very few words. *Proceedings of the 14th International Conference on Machine Learning*, pp. 170-178.
- Lalmas, M. and Bruza, P. (1998) The use of Logic in Information Retrieval Modeling. *Knowledge Engineering Review*, 13(3):263-295.
- Lalmas, M. and van Rijsbergen, C. (1992). A logical model of information retrieval based on situation theory. *Proceedings of the 14th British Computer Society Annual Information Retrieval Colloquium*, pp. 1 – 13.
- Landauer, T., Foltz, P. and Laham, D. (1998): An introduction to latent semantic analysis. *Discourse Processes*, 25:259-284.
- Lang, J. (1997) Possibilistic Logic: Algorithms and Complexity, in: *Handbook of Algorithms for Uncertainty and Defeasible Reasoning*, Kohlas J. and Moral S. eds, Kluwer Academic Publishers, Norwell, Massachusetts.
- Lau, R., ter Hofstede, A., and Bruza, P. (1999). A Study of Belief Revision in the context of Adaptive Information Filtering. In Lucas Chi Kwong Hui and Dik Lun Lee, editors, *Proceedings of the 5th International Computer Science Conference on Internet Applications, volume 1749 of Lecture Notes in Computer Science*, pp. 1-10, Springer-Verlag.

- Lau, R., ter Hofstede, A., and Bruza, P. (2001). Maxi-Adjustment and Possibilistic Deduction for Adaptive Information Agents. *Journal of Applied Non-classical Logics*, 11(1-2): 169-201.
- Lau, R. (2002) The State of the Art in Adaptive Information Agents. *International Journal on Artificial Intelligence Tools*, 11(1):19-61.
- Lau, R. (2003) Context Sensitive Text Mining and Belief Revision for Adaptive Information Retrieval. *Proceedings of the 2nd IEEE/WIC International Conference on Web Intelligence (WI'03)*, October 13 - 17, 2003, Halifax, Canada, pp. 256 - 262, IEEE Press.
- Lau, R., Bruza, P., and Song, D. (2004). Belief Revision for Adaptive Information Retrieval. *Proceedings of the 27th ACM SIGIR International Conference on Research and Development in Information Retrieval*, July 25 - 29, 2004, Sheffield, U.K., pp. 130 – 137.
- Lawrence, S. (2000) Context in Web Search. *IEEE Data Engineering Bulletin*, 23(3):25-32.
- Levy, D. (2005) Users and interaction track: memex and hypertext: To grow in wisdom: vannevar bush, information overload, and the life of leisure. *Proceedings of the 5th ACM/IEEE-CS joint conference on Digital libraries*, pp.281-286.
- Liberatore, P. and Schaerf, M. (1998). The complexity of model checking for propositional default logics. *Proceedings of the Thirteenth European Conference on Artificial Intelligence (ECAI'98)*, pages 18-22.
- Logan, B., Reece, S., and Jones, K. (1994) Modelling information retrieval agents with belief revision. *Proceedings of the 17th Annual International Conference on Research and Development in Information Retrieval*, pp. 91-100.
- Losada, D. and Barreiro, A. (1999) Using a belief revision operator for document ranking in extended boolean models. *Proceedings of the 22nd Annual International Conference on Research and Development in Information Retrieval*, pp. 66-72.
- Losada, D. and Barreiro, A. (2001) A Logical Model for Information Retrieval based on Propositional Logic and Belief Revision. *The Computer Journal*, 44(5):410-424.
- Losada, D. and Barreiro, A. (2003a) Embedding Term Similarity and Inverse Document Frequency into a Logical Model of Information Retrieval. *Journal of American Society for Information Science and Technology*, 54(4):285-301.
- Losada, D. and Barreiro, A. (2003b) Propositional Logic Representations for Documents and Queries: A Large-Scale Evaluation. *Proceedings of the 25th European Conference on IR Research (ECIR)*, volume 2633 of LNCS, pp. 219-234.
- Lowe, W. (2001) Towards a theory of semantic space. *Proceedings of the 23rd Annual Conference of the Cognitive Society*, pp. 576-581.
- Maes, P. (1994) Agents that reduce work and information overload. *Communications of the ACM*, 37(7):31-40.
- Meghini, M., Sebastiani, F., Straccia, U., and Thanos, C. (1993). A Model of Information Retrieval based on Terminological Logic. *Proceedings of the 16th ACM SIGIR International Conference on Research and Development in Information Retrieval*, pp. 298 – 307.
- Miller, G. A., Beckwith, R., Fellbaum, C., Gross, D., and Miller, K. (1990). Introduction to WordNet: An online lexical database. *Journal of Lexicography*, 3(4):234–244.
- Nie, J. (1986). An outline of a general model for information retrieval systems. *Proceedings of the 9th ACM SIGIR International Conference on Research and Development in Information Retrieval*, pp. 495 – 506.
- Nie, J. (1989) An Information Retrieval Model based on Modal Logic. *Information Processing and Management*, 25(5):477-491.

- Nie, J., Brisebois, M., and Lepage, F. (1995) Information retrieval as counterfactual. *The Computer Journal*, 38(8):643-657.
- Neves, R., Bonnefon, J., and Raufaste, E. (2000) Rationality in human nonmonotonic inference. *Proceedings of the 8th International Workshop on Nonmonotonic Reasoning NMR'2000*, April 2000. Available from <http://xxx.lanl.gov/html/cs.AI/0003073>.
- Raufaste, R. and Neves, R. (1998). Empirical evaluation of possibility theory in human radiological diagnosis. *Proceedings of the 13th European Conference on Artificial Intelligence ECAI'98*, pp. 124 – 128.
- Robertson, S. (2002) Comparing the Performance of Adaptive Filtering and Ranked Output Systems. *Information Retrieval*, 5:257-268.
- Rocchio, J. (1971) Relevance Feedback in Information Retrieval. In G. Salton, editor, *The SMART retrieval system: experiments in automatic document processing*, pp.313-323.
- Salton, G. and McGill, M. (1983) *Introduction to Modern Information Retrieval*. McGraw-Hill, New York.
- Salton, G. (1991) Developments in automatic text retrieval. *Science*, 253:974-980.
- Sebastiani, F. (1998) On the role of logic in information retrieval. *Information Processing and Management*, 34(1):1-18.
- Song, D. and Bruza, P. (2001) Discovering Information Flow Using a High Dimensional Conceptual Space. *Proceedings of the 24th Annual International Conference on Research and Development in Information Retrieval (SIGIR'01)*, pp. 327-333.
- Song, D. and Bruza, P. (2003) Towards A Theory of Context Sensitive Informational Inference. *Journal of American Society for Information Science and Technology*, 54(4):326-339.
- van Rijsbergen, C., (1986) A non-classical logic for information retrieval. *The Computer Journal*, 29(6):481-485.
- Williams, M. (1995). Iterated theory base change: A computational model. *Proceedings of the Fourteenth International Joint Conference on artificial Intelligence (IJCAI'1995)*, pp. 1541-1547.
- Williams, M. (1997). Anytime belief revision. *Proceedings of the Fifteenth International Joint Conference on artificial Intelligence (IJCAI'1997)*, pp. 74-79.
- Wong, K.F., Song, D., Bruza, P.D., Cheng, C.H. (2001) Application of Aboutness to Functional Benchmarking in information retrieval. *ACM Transactions on Information Systems (TOIS)*, 19(4):337-370.
- Zhang, Y. and Callan, J. (2001) Maximum Likelihood Estimation for Filtering Thresholds. *Proceedings of the 24th annual international ACM SIGIR conference on Research and development in information retrieval*, pp.294-302.