Relevance Sensitive Non-Monotonic Inference on Belief Sequences

Samir Chopra*
Knowledge Systems Group
Department of Artificial Intelligence
University of New South Wales
Sydney, NSW 2052
Australia
schopra@cse.unsw.edu.au

Konstantinos Georgatos
John Jay College of CUNY
899 Tenth Avenue
New York, NY 10019
kgeorgatos@jjay.cuny.edu

Rohit Parikh†
Brooklyn College and CUNY Graduate Center
365 Fifth Avenue
New York, NY 10016
ripbc@cunyvm.cuny.edu

ABSTRACT. We present a method for relevance sensitive non-monotonic inference from belief sequences which incorporates insights pertaining to prioritized inference and relevance sensitive, inconsistency tolerant belief revision. Our model uses a finite, logically open sequence of propositional formulas as a representation for beliefs and defines a notion of inference from maxiconsistent subsets of formulas guided by two orderings: a temporal sequencing and an ordering based on relevance relations between the putative conclusion and formulas in the sequence. The relevance relations are ternary (using context as a parameter) as opposed to standard binary axiomatizations. The inference operation thus defined easily handles iterated revision by maintaining a revision history, blocks the derivation of inconsistent answers from a possibly inconsistent sequence and maintains the distinction between explicit and implicit beliefs. In doing so, it provides a finitely presented formalism and a plausible model of reasoning for automated agents.

Introduction

Belief revision is the process of transforming a belief state K upon receipt of new information α. In dealing with this problem two issues arise. One is the

*This research was carried out while the first author was at the CUNY Graduate Center, 365 Fifth Avenue, New York, NY 10016

†The first and third authors would like to thank NSF, CUNY-FRAP and NYU for their support of our research.

non-monotonicity of the process of inference from beliefs. If a belief state K is revised by some new formula α, then a formula β which was derivable earlier may no longer be derivable. This is the logical problem to which there are two fundamental approaches. In the *logic-constrained* or *horizontal* approach [GR95], the belief representation is a closed and consistent *theory* K, and given a new proposition α, [AGM85] propose postulates for $K * \alpha$, the theory revised with α. In this approach the belief state is itself sophisticated, and constructing the revised belief state requires some ingenuity. But the usual constructions for AGM revision operators using *selection functions* [AGM85] and *epistemic entrenchments* [GM88] often fail to provide an adequate account of *iterated revision*. Also, AGM-like postulates do not specify how we *came* to believe K and after revision, it is assumed that $K * \alpha$ is a generic theory. But these postulates ignore the fact that α was our *last* information.

The *vertical* approach, in contrast, uses trivial (and repeatable) operations of revision and expansion on finite, logically open, belief representations, but utilizes a *sophisticated* notion of non-monotonic inference, see, e.g. [Doy79], [Bre91]. We suggest, in conformance with the vertical approach, that $K * \alpha$ be taken to be the *belief sequence* $K; \alpha$, i.e., a *finite*, logically open, sequence of propositions with α being the most recent. The suggestion that the most perspicuous way to represent our beliefs is a finite, logically open, set of sentences, i.e., a *belief base* has been made (amongst others) by [Han92], [Neb92]; the notion that a sequence of formulas captures the importance of temporal ordering and of maintaining a revision history is noted by [Ry91] and [Leh95]. We do not treat a belief sequence as a set but rather as a linear order much like an entrenchment except that *two* orderings, level of relevance and temporal order govern the sequence.

In this study, then, a belief base is a sequence σ of all formulas which have been received, in their usual temporal order. A typical sequence σ may contain formulas which contradict each other and hence when the question is raised whether some formula β is derivable from σ, some judgment will have to be made as to *which* formulas in σ are to be used to answer the query; clearly not all can be so used.

Here the problem of *computational tractability* arises. Typically, σ may be quite large and the logical problem of determining which part of σ is to be used and the problem of determining whether β or its negation follows from that part is likely to be **NP**-complete at the very least. Thus some strategy for reducing this cost is needed. Human agents do not normally report computational difficulties in answering queries; if they are unable to answer a query they typically attribute their difficulty to ignorance and not to computational limitations. Hence they must be using some such cost reducing strategy and a realistic attempt to model human reasoning would require that we explore such strategies.

Since revision becomes simple under this approach, the notion of inference must be correspondingly more sophisticated. We describe a method for *non-monotonic inference* from belief sequences which does this, but departs significantly from previous approaches in one respect. It makes heavy use of *relevance relations* amongst formulas in a belief base. These relevance relations serve as a particular instance of a more general model of context-sensitive inference.

The particular method presented here incorporates the insights in earlier proposals made by [Geo96], [Par96] and [CP99]. [Geo96] uses the linear order of a belief sequence as a prioritization to generate a variety of inference relations, shows that these schemes are non-monotonic and, therefore, induce a method for belief revision. [Par96] shows that if we have a theory referring to two or more disjoint subjects, then our language can be partitioned into corresponding sub-languages, and it is suggested that new information about *one* of them should not affect any other. This ensures a relevance or context sensitive, *localized* notion of belief revision and serves as one way of capturing a more general notion of relatedness amongst propositions in a belief base (as studied by [Was99]). [CP99] consider sets of theories called B-structures, which are individually consistent, but can be *jointly* inconsistent, to capture the intuition that real agents often reason with an inconsistent, yet usable, set of beliefs which is divided into individually consistent compartments.

Our (current) method of inference blocks the derivation of explicitly inconsistent beliefs from a possibly inconsistent belief sequence by using a notion of inference from *maxiconsistent subsets of relevant formulas.* Choosing maxiconsistent subsequences in order to avoid inconsistency was used in [Geo96], while relevance is determined, as in [Par96] and [CP99] by a specialized notion of language overlap or by other context determined features. The formula whose inference from the sequence is to be determined imposes a prioritization on the formulas present in the sequence by virtue of its relevance relations with them (thus reorganizing the temporal ordering present in the sequence). The resulting procedure for inference serves as a generalization of the methods presented in [Geo96] and [CP99]. In this way, we hope to present a model for belief revision that is a plausible representation of real agents' reasoning.

In the first section of the paper, we present preliminary definitions and establish the notion of relevance that we will work with. In the second section we define our notion of inference and examine its properties.

Notation: In the following, $\mathcal{L}$ is a finite propositional language with the usual logical connectives $(\neg, \vee, \wedge, \rightarrow, \leftrightarrow)$. The constants *true, false* are in $\mathcal{L}$. Greek letters α, β, γ denote arbitrary formulas while Roman lower case letters p, q, r denote propositional atoms. $\alpha \Leftrightarrow \beta$ means that $\alpha \leftrightarrow \beta$ is a tautology. Cn will denote the usual classical consequence relation. We reserve the letter σ for belief sequences.

The General Method: An Introduction

In this paper we implement the notion of relevance via a set Δ of non-logical symbols *relevant to a query* β, and a *triggering function* τ which determines *how* this set Δ is to be used. We illustrate this by an example. Suppose we receive some query about the affair of the Cuban boy Elian who was taken by his mother from Cuba for an abortive (for her) escape attempt and who was subsequently kept by some relatives in Miami despite appeals from US Attorney General Janet Reno to turn the boy over to his father.

Wanting to answer the query from our existing information σ, we might then take Δ to consist of the symbols {*Elian, Reno, Miami*} and only use formulas from our database σ which contain *any one* of these symbols. We would then be using the triggering function which says: Use all elements of σ which overlap Δ. However, a more sophisticated strategy might use the *same* set {*Elian, Reno, Miami*} but ask for formulas from σ which contain the symbol *Elian* and at least one of the other two symbols. We would then exclude information about other problems which Janet Reno has dealt with, and also exclude information about Elian's childhood or about travel agents booking vacations to Miami.

Be that as it may, once we have decided on Δ and τ we will get some subsequence σ' of σ consisting of all formulas α of σ which have been triggered by τ using Δ, i.e. those for which $\tau(\Delta, \alpha) = 1$. σ' is still in temporal order so that the most recent formulas are last[1]. We now construct a maxi-consistent set Γ by starting from the most recent formulas in σ' and work backwards, adding all formulas which are consistent with formulas already picked. We now say that $\sigma \models_{\Delta,\tau} \beta$ iff Γ classically implies β.

Since the set Δ of relevant symbols and the sequence σ' are easily computed from the query β, and the sequence σ' is much smaller than σ, the construction of Γ and the logical problem of checking the derivability of β from Γ become much more tractable.

We now get several nice properties including a *weak monotony.* Of course we do not expect *in general* that if $\sigma \models_{\Delta,\tau} \beta$ then also $\sigma * \alpha \models_{\Delta,\tau} \beta$. However if α is *not triggered* then it will make no difference to the construction of σ', and the implication of $\sigma * \alpha \models_{\Delta,\tau} \beta$ from $\sigma \models_{\Delta,\tau} \beta$ will go through. Thus, for instance, if our database is extended by information about New York weather, queries about Elian will have the same answers that they did before. Full *adjunction* will also hold for fixed Δ, τ. Thus if $\sigma \models_{\Delta,\tau} \beta$ and $\sigma \models_{\Delta,\tau} \gamma$, then $\sigma \models_{\Delta,\tau} (\beta \wedge \gamma)$ will also hold. Similarly, if $\sigma \models_{\Delta,\tau} \beta$, and $\sigma \models_{\Delta,\tau} (\beta \rightarrow \gamma)$, then $\sigma \models_{\Delta,\tau} \gamma$ will hold. Theorems 1-5 below describe some of these properties.

In practice the set Δ will naturally depend on the formula β about which we are querying. Thus the same σ may use a set Δ_1 to answer queries about β, and another set Δ_2 to answer queries about γ. In that case it may well happen that the shifts in Δ may prevent rules like adjunction from holding in practice and we might be able to derive each of β, γ without also deriving their conjunction.[2]

[1] In a later application we also allow relevance to move some formulas ahead of others, but the temporal order remains relevant.

[2] A conversation initiated by Herman Kahn illustrates this principle. Kahn asked a class what the US President would do if the Soviets bombed some large US city. The class responded that the President would launch a massive attack on the Soviet Union. On the further question what the Soviet Union would do in response and what the consequences

Belief Sequences

We begin with the definition of a belief sequence:

Definition 1 *A belief sequence is a sequence of formulas under a temporal ordering, i.e., a sequence of formulas, $\sigma = \beta_1 \ldots \beta_n$ where for any pair of beliefs β_i, β_j if $i < j$, β_j is more recent than β_i. Given two sequences, σ_1, σ_2 we say that $\sigma_1 \sqsubseteq \sigma_2$ if σ_2 is obtained from σ_1 by the concatenation of zero or more formulas; σ_1 will then be an initial segment of σ_2*

Under a temporal ordering the most recent formulas occur at the tail of the sequence. We assume that each formula α in the sequence is expressed in its *smallest language* L_α as defined below. Note that the linear temporal order can be replaced with some other linear order expressing prioritization. For example, one could order the propositions on the basis of the trustworthiness of their source.

Definition 2 *Given a sequence σ, set Δ of relevant propositional symbols and triggering function τ let $\tau(\Delta, \sigma)$ be the sequence (in the same order as in σ) of formulas of σ triggered by τ using Δ.*

Definition 3 *Given a sequence σ of formulas let $maxi(\sigma)$ be the set of formulas of σ selected as follows. We work backwards through σ, adding formulas which are consistent with formulas chosen so far, and discarding those which are not. Specifically, if $\sigma = \beta_1, ..., \beta_n$, then we define sets Γ_i by $\Gamma_0 = \emptyset$ and for each $j \leq n$, letting $i = n - j$,*
if β_i is consistent with Γ_j then $\Gamma_{j+1} = \Gamma_j \cup \{\beta_i\}$.
If β_i is not consistent with Γ_j then $\Gamma_{j+1} = \Gamma_j$.
Then $maxi(\sigma)$ is simply Γ_n.

If τ, Δ are clear from the context we write $ch(\sigma)$ for the set $maxi(\tau(\Delta, \sigma))$. This is the maxi-consistent subset of $\tau(\Delta, \sigma)$ chosen as in definition 3 from the sequence defined in definition 2. Obviously ch depends on both Δ and τ and so we will be careful to indicate this fact when Δ or τ may vary. Use of the notation ch will therefore indicate that we are keeping Δ and τ fixed. Note that for each $\beta \in \sigma$, either $\beta \in maxi(\sigma)$ or else β is inconsistent with $maxi(\sigma)$.

Definition 4 *$\sigma \models_{\Delta,\tau} \beta$ iff $ch(\sigma) \models \beta$.*

Theorem 1 *(i) Let α be consistent with $ch(\sigma)$ and triggered by Δ, τ, then $ch(\sigma; \alpha) = ch(\sigma) \cup \{\alpha\}$ and conversely. (ii) If α is not triggered by Δ, τ, then $ch(\sigma; \alpha) = ch(\sigma)$.*

would be, the class realized that both sides would be wiped out. Kahn then asked about the initial stage, "So why would the President launch an attack on the Soviet Union?" This dialogue illustrates that local thinking, which we are formalizing in this paper has the obvious pitfall that second order or third order effects and consequences may be overlooked in a particular chain of reasoning. They may then be brought out by a more detailed query.

Proof. (ii) is obvious as then $\tau(\Delta, \sigma; \alpha) = \tau(\Delta, \sigma)$. Also the reverse direction of (i) is obvious as $ch(\sigma; \alpha)$ is required to be consistent.

This leaves the forward direction of (i) as the only nontrivial case. So suppose that α is consistent with $ch(\sigma)$ and α is triggered by Δ, τ. Then $\tau(\Delta, \sigma; \alpha) = \tau(\Delta, \sigma); \alpha$. Clearly when we are creating $maxi(\tau(\Delta, \sigma; \alpha))$ we will start with α (which is assumed consistent) and working backward through $\tau(\Delta, \sigma)$, we will select formulas which are consistent with formulas already chosen. We now claim that even though we have α in hand, we will do just the same as when we were choosing the formulas of $maxi(\tau(\Delta, \sigma)$. For if a formula was rejected before, it was inconsistent with formulas already chosen and so it will be inconsistent with these plus α. If a formula *was* chosen before, then it ended up in $maxi(\tau(\Delta, \sigma)$ and is therefore consistent with already chosen formulas plus α by assumption. Hence it will be chosen this second time as well. ∎

We have given this one proof in detail. The proofs of the theorems below are analogous and given in a more sketchy way. In all four cases, if μ is not triggered, then the conclusion is obvious. So we only consider the case where it is triggered.

Theorem 2 *Assume that α is triggered and $\alpha \models \mu$. Then $\sigma; \mu; \alpha \models_{\Delta,\tau} \beta$ iff $\sigma; \alpha \models_{\Delta,\tau} \beta$.*

Proof. It is easily seen that $ch(\sigma; \mu; \alpha) = ch(\sigma; \alpha) \cup \{\mu\}$ and both sides contain α. Hence, μ being weaker than α provides no additional power. ∎

Theorem 3 *Assume that α is triggered and $\alpha \models \neg\mu$. Then $\sigma; \mu; \alpha \models_{\Delta,\tau} \beta$ iff $\sigma; \alpha \models_{\Delta,\tau} \beta$.*

Proof. It is easily seen that $ch(\sigma; \mu; \alpha) = ch(\sigma; \alpha)$, as α, after being selected, blocks μ. ∎

Theorem 4 *If $\sigma; \alpha \not\models_{\Delta,\tau} \neg\mu$ then $\sigma; \mu; \alpha \models_{\Delta,\tau} \mu$.*

Proof. μ is consistent with $ch(\sigma; \alpha)$. Hence $ch(\sigma; \mu; \alpha) = ch(\sigma; \alpha) \cup \{\mu\}$ and the conclusion follows. ∎

Theorem 5 *If $\sigma; \alpha \models_{\Delta,\tau} \mu$ then $\sigma; \mu; \alpha \models_{\Delta,\tau} \mu$.*

Proof. This is an immediate corollary of the last theorem. ∎

The rest of the paper will consider ways in which we construct Δ, τ in *specific cases* and the logical and other properties which obtain. It will become apparent that some applications may even use several Δ's for a query; a Δ_1 of more relevant symbols, a Δ_2 of less relevant symbols and so on. Our guiding principle will be: *inference from belief states is classical inference from maxi-consistent subsets of the belief state generated in accordance with the heuristic associated with a triggering function and a context set.* In the rest of the paper, we will pick a particular context set and a particular triggering function and then investigate their properties.

Relevance Relations

We now present a relation of *relatedness* amongst formulas in a sequence (originally proposed and used in [Par96], [CP99] for localized belief revision) as a preliminary to its modification for use in this study. *This relatedness relation will help us pick a particular context set and triggering function.* First, a distinction between different languages that a formula can be expressed in:

Definition 5 *The language $L(\alpha)$ is the set of propositional variables which actually occur in a formula α; the language L_α of α is the* smallest *set of propositional variables which can be used to express β, a formula logically equivalent to α.*

So, if $\alpha = p_1 \wedge (p_2 \vee \neg p_2)$ then β is p_1 and $L_\alpha = \{p_1\}$ while $L(\alpha) = \{p_1, p_2\}$. (L_α is unique, cf. Lemma LS1 in [Par96]). L_α has logically attractive properties, e.g. if $\alpha \Leftrightarrow \beta$, then $L_\alpha = L_\beta$. Hence we shall work exclusively with this notion.

Definition 6 *α, β are related by* syntactic *language overlap ($\mathcal{R}_s(\alpha, \beta)$) if $L(\alpha) \cap L(\beta) \neq \emptyset$. α, β are related by* logical *language overlap ($\mathcal{R}_l(\alpha, \beta)$)if $L_\alpha \cap L_\beta \neq \emptyset$.*

It is easily seen that $\mathcal{R}_l(\alpha, \beta))$ implies $\mathcal{R}_s(\alpha, \beta))$.

In earlier work on relatedness, Epstein [Eps95] imposes the following 'plausibility' requirements on any relatedness relation $\mathcal{R}(\alpha, \beta)$ (α is related to β):

R1 $\mathcal{R}(\alpha, \beta)$ iff $\mathcal{R}\ (\neg\alpha, \beta)$.
R2 $\mathcal{R}(\alpha, \beta \wedge \gamma)$ iff $\mathcal{R}\ (\alpha, \beta \rightarrow \gamma)$.
R3 $\mathcal{R}(\alpha, \beta)$ iff $\mathcal{R}(\beta, \alpha)$.
R4 $\mathcal{R}(\alpha, \alpha)$.
R5 $\mathcal{R}\ (\alpha, \beta \rightarrow \gamma)$ iff $\mathcal{R}\ (\alpha, \beta)$ or $\mathcal{R}(\alpha, \gamma)$.

Rodrigues [Rod97] has shown that the relation ($\mathcal{R}_s(\alpha, \beta)$), i.e., syntactic language overlap, is the smallest relation satisfying Epstein's conditions. However, we might also consider a condition not considered by Epstein:

R6 if $\mathcal{R}(\alpha, \beta)$ and $\beta \Leftrightarrow \beta'$ then $\mathcal{R}(\alpha, \beta')$.

Observation 1 *$\mathcal{R}_s$ does not satisfy* **R6** *whereas the relatedness relation $\mathcal{R}_l$ does satisfy* **R6** *as well as conditions* **R1, R3-4** *(but not the conditions* **R2, R5***).*

Since condition **R6** is very natural, we wonder if conditions **R2, R5** were adopted out of a feeling that they are compatible with **R6**. Indeed, in general, though not always, we do have $L_{\beta\wedge\gamma} = L_{\beta\rightarrow\gamma} = L_\beta \cup L_\gamma$. In general α will be relevant to the composite formula iff it is relevant to at least one part. But this need not always happen.

For example, if we let $\beta = p, \gamma = \neg p$, then $\beta \rightarrow \gamma$ is equivalent to $\neg p$ and of course relevant to p. However, $\beta \wedge \gamma$ is $p \wedge \neg p$, a downright contradiction and not relevant (in our opinion) to *anything*. This fact casts some doubt on the intuition behind **R2**. Indeed it turns out that **R2,R5** are incompatible with the natural requirement **R6**.

Our R_l does satisfy *half* of **R5**:

R5a: If $\mathcal{R}_l(\alpha, \beta \to \gamma)$ then $\mathcal{R}_l(\alpha, \beta)$ or $\mathcal{R}_l(\alpha, \gamma)$.

With the discussion above as a guide, we now develop a *context-sensitive* measure for *relevance* amongst formulas in a belief sequence.

Definition 7 *Two formulas $\beta_1, \beta_2 \in \mathcal{L}$ are (logically)* disjoint *iff $L_\alpha \cap L_\beta = \emptyset$.*

Definition 8 1. *A pair of formulas, α, β are directly relevant if they are not logically disjoint, i.e., if $L_\alpha \cap L_\beta \neq \emptyset$.*

2. *Given a belief sequence σ, a pair of formulas α, β are k-relevant wrt (with respect to) σ if $\exists \chi_1, \chi_2, \ldots \chi_k \in \sigma$ such that:*
i) α, χ_1 are directly relevant
ii) χ_i, χ_{i+1} are directly relevant for $i = 1, \ldots k-1$
iii) χ_k, β are directly relevant.
We write $\mathcal{R}_k(\alpha, \beta, \sigma)$ to indicate that α, β are k-relevant wrt σ. If $k = 0$ above, the formulas are directly relevant.

3. *A pair of formulas are irrelevant if they are not k-relevant for any k.*

4. *$rel(\alpha, \beta, \sigma)$ is the the lowest k such that α, β are k-relevant wrt σ (we let it be ∞ if α, β are irrelevant).*

In the following observations, we omit the sequence σ when clear from the context.

Observation 2 1. *If a pair of formulas are k-relevant, then, $\forall m > k$, they are m-relevant as well.*

2. *Let $[\![\sigma]\!] = \{\alpha \mid \alpha$ occurs in $\sigma\}$. We say that $\sigma_1 \subseteq' \sigma_2$ iff $[\![\sigma_1]\!] \subseteq [\![\sigma_2]\!]$. If $\sigma_1 \subseteq' \sigma_2$, then $rel(\alpha, \beta, \sigma_2) \leq rel(\alpha, \beta, \sigma_1)$.*

3. *The relation $\mathcal{R}_k(\alpha, \beta, \sigma)$ is symmetric and reflexive in the first two arguments but obviously not transitive.*

4. *If $k = 0$ then the sequence σ is irrelevant to the question of k-relevance of formulas α, β; if $k \neq 0$, then σ is a parameter in determining relevance between α, β.*

$\mathcal{R}_k(\alpha, \beta, \sigma)$ depends only on L_α, L_β and $[\![\sigma]\!]$ and not on α, β and σ themselves.

So if two formulas are relevant to each other at one level, then they are relevant at all weaker (*higher* k) levels. The definition of relevance above explicitly brings in (as a third parameter) the sequence σ which can form a bridge between formulas which do not have a direct overlap. Normally, relevance has been thought of as a binary relation; our definition renders it a ternary one. Thus the sequence σ can play the role of connecting formulas which do not have any overlap in language but are connected *through* other beliefs. A fact about Taj Mahal and one about India will be connected because of our (true) belief that the Taj Mahal is in India. We can also have more distant - and less convincing - indirect connections. For example, we can think of beliefs linking the subject matter *European History*, to *European Music* and then to *Music* in general, to *Indian Music*. But it is unlikely that we have beliefs *directly* linking European History to Indian Music and this level ($k = 2$) may be too weak (too high) to be useful in most considerations.

The above definition extends the definition of relevance proposed in [Par96] and makes explicit the *contextual nature* of the relevance definition: two formulas α, β may have different degrees of relevance to each other in virtue of different belief sequences. A belief sequence defines a particular context for subject matters; pairs of formulas acquire different relationships to one another given differing contexts. The more basic beliefs (i.e., elements of σ) that an agent has, the more likely it is to connect two apparently unconnected subjects.

Revision and Inference on Belief Sequences

Revision on belief sequences is easily achieved: we simply concatenate the new formula γ to the sequence. The sequence $\beta_1, \ldots, \beta_n$ becomes $\beta_1 \ldots \beta_n, \beta_{n+1}$ where $\beta_{n+1} = \gamma$.

Definition 9 $\sigma * \gamma = \sigma; \gamma$ *where* $\sigma; \gamma$ *represents the concatenation of* γ *to a belief sequence* σ.

Note that $[\![\sigma * \gamma]\!] = [\![\sigma]\!] \cup \{\gamma\}$.
Example: Consider the sequence $\sigma = [p, q, p \wedge q, \neg r]$. If we receive the information $\beta = \neg p$, it is simply appended to the sequence to give us $\sigma * \beta = [p, q, p \wedge q, \neg r, \neg p]$. $[\![\sigma]\!]$ is now inconsistent, and we need a notion of inference which renders the agent's beliefs coherent.

Prioritized Inference

In the current study we settle on the following parameters: the context set Δ is L_α, the triggering function causes a reordering of the belief sequence σ (as defined below) and the triggering heuristic is logical language overlap i.e., a context-specific relevance relation.

Suppose that we have a sequence of formulas which is our current belief base and we are asked about some formula γ, whether we believe it or not. As we saw just above, the set $[\![\sigma]\!]$ may well be inconsistent, and hence to decide about the status of γ we will need to pick some consistent subset of $[\![\sigma]\!]$ which we bring into play. The choice of formulas to be in such a subset will be governed by two considerations. One is temporality under which more recently received formulas have priority over older formulas. The other is relevance according to which more relevant formulas have priority. Clearly we need to decide which order counts more. We have made the decision in this paper that relevance is more important than temporality but that between two formulas of equal relevance, the more recent formula has priority, We concede however, that the other procedure, to count temporality more, also has something to say in its favor. Ordinary human reasoning, in our opinion, is a pragmatic blend of the two techniques.

We use the *maxiconsistent* approach to define *prioritized* inference on a belief sequence σ. This method employs a consistent subset of $[\![\sigma]\!]$ obtained as follows. Consider a formula γ in its smallest language L_γ. We construct a maxiconsistent subset $\Gamma_{\langle\sigma,k,\gamma\rangle}$ (of k-relevant to γ formulas) of $[\![\sigma]\!]$. The construction of this set is regulated by the ordering $\prec$ that γ creates on σ, which arranges $\beta_1, \ldots, \beta_n$ into $\delta_1, \ldots, \delta_n$ as follows:

Definition 10 *Given a formula γ, a sequence σ, $\beta, \beta' \in \sigma$, $\beta' \prec \beta$ if either a) $rel(\gamma, \beta, \sigma) < rel(\gamma, \beta', \sigma)$ (i.e., β is more relevant to γ than β') or b) β, β' are equally relevant (i.e., $rel(\gamma, \beta, \sigma) = rel(\gamma, \beta', \sigma)$) but β is more recently received than β'.*

In other words, we are still respecting temporal order, but formulas which are more relevant are brought towards the tail of the sequence which we count more heavily[3]. The $\delta_1, \ldots, \delta_n$ are the $\beta_1, \ldots, \beta_n$ under this order. Let σ' be the re-ordered sequence $\delta_1, \ldots, \delta_n$. σ' of course depends on γ, and we will refer to it as $\sigma(\gamma)$ in case we ever need this (admittedly clumsy) notation.

Let Δ be the set of symbols occurring in the query formula γ. For some given level of relevance k, let the triggering function τ pick out formulas from σ' which are at least k-relevant to γ. Let $\sigma'' = \tau(\Delta, \sigma')$ and $\Gamma_{\langle\sigma,k,\gamma\rangle} = maxi(\sigma'')$. What we have done is to reorder σ so as to bring more relevant formulas forward, then selected formulas which are at least k-relevant to γ, and finally, created a maxi-consistent set from this reordered and culled sequence. This final maxi-consistent set is $\Gamma_{\langle\sigma,k,\gamma\rangle}$.

Definition 11 $\sigma \vdash_k \gamma$ *iff* $\Gamma_{\langle\sigma,k,\gamma\rangle} \vdash \gamma$

Once $\Gamma_{\langle\sigma,k,\gamma\rangle}$ has been constructed, the inference operation defined above enables a query answering scheme for the agent with definite responses:

If $\Gamma_{\langle\sigma,k,\gamma\rangle} \vdash \gamma$, then answer 'yes'.
If $\Gamma_{\langle\sigma,k,\gamma\rangle} \vdash \neg\gamma$ then answer 'no'.
Otherwise, answer 'no information'.

Even if $[\![\sigma]\!]$ is inconsistent, the agent is able to give consistent answers to every individual query.

[3] If $\alpha_1, \alpha_2, \alpha_3, \alpha_4$ are a given sequence with α_4 the most recent, but α_2, α_3 are more relevant than α_1, α_4, then the sequence will be reordered to be $\alpha_1, \alpha_4, \alpha_2, \alpha_3$.

Discussion

The notion of inference thus defined has some desirable features. As an example, suppose our belief sequence is initiated by first being told p and then $\neg p \wedge \neg q$. $\neg p \wedge \neg q$ overrides p and we will no longer answer 'yes' to p. However, if we are now told $p \vee q$, this new information overrides $\neg p \wedge \neg q$. Thus, our maxiconsistent set $\Gamma_{<\sigma,0,p>}$ is $\{p \vee q, p\}$ and the query p? will now be answered in the affirmative. This is plausible since the latest information decreases the reliability of $\neg p \wedge \neg q$ and the original information p regains its original standing. Such accomodations are not easily made within traditional, AGM-based frameworks for belief revision.

Derived conclusions are justified on the basis of their relevance relations with formulas present in the sequence *and* their order of reception. For example, consider $\sigma = [s, s \to r, p, p \to q, q \to \neg r]$, and $\gamma = r$. Now $s \to r, q \to \neg r]$ are directly relevant to γ. $s, p \to q]$ are 1-relevant, and finally p is 2-relevant. Moving the more relvant formulas forward we get $\sigma' = [p, s, p \to q, s \to r, q \to \neg r]$. At level 0 we get no answer to γ. At levels 1 and higher the last four formulas constitute our maxi-consistent set $ch(sigma)$ and yield "yes" to γ.

The effect of some new information can depend on whether it came "in several pieces" or as a single formula. Receiving two pieces of information individually and together often has different effects on an agent's epistemic state. If we receive α and β seperately, then a later information which undermines α need not undermine β as well. But if we received the *conjunction* $\alpha \wedge \beta$ then undermining one will undermine both. As an example, consider the arguments against the AGM Axioms 7,8 in [Par96] where it is argued that revision by conjunctions is not the same as revising by conjuncts individually. (The following two examples pertain to the framework of [Par96].)

Example: Suppose our current beliefs are $\neg p, \neg q$ and we receive the information that $p \vee q$. If p is much more believable than q we might now decide that p holds. Suppose we next hear that $\neg p \vee q$. Since this is consistent with our current state, we accept it, ending with a state generated by p, q. However, if we had received the *conjunction* $(p \vee q) \wedge (\neg p \vee q)$, this would be equivalent to receiving just q and we would never have believed p at all. Note that this is also a situation where we found Epstein's postulate **R2** to be implausible, for the language $L_{(p \vee q) \wedge (\neg p \vee q)}$ does not equal $L_{(p \vee q)} \cup L_{(\neg p \vee q)}$ in this case.

For a concrete example let p be "The stove is smoking" and q be "The house is on fire". If I am told that either the stove is smoking or that the house is on fire, I will choose to believe that the stove is smoking. If I later find that either the stove is not smoking or the house is on fire, since this belief is consistent, I must now add it, and conclude that the stove is smoking *and* the house is on fire! But if I had been told the conjunction, which is equivalent to "The house is on fire", I would never have thought about the stove at all.

Returning to our current framework, consider the sequence $\sigma_2 = [p \wedge q, r \wedge \neg q]$. Both p and r are derivable although their respective derivation is based on *incompatible* information. This is due to the fact that the formulas which are preferred in each case depend on the query. This is plausible however; agents often think about unrelated pieces of information in isolation from each other. Here p and r are not directly related and therefore may be thought about by the agent separately.

Prioritization of Directly relevant formulas

Further depth could be introduced into the method provided above by introducing a prioritization amongst formulas based on the *amount* by which formulas β of σ overlap with the query formula γ. Thus such a prioritization $\preceq_\gamma$ might be defined as follows: let $\beta_1 \preceq_\gamma \beta_2$ if $|L_{\beta_1} \cap L_\gamma| \geq |L_{\beta_2} \cap L_\gamma|$. Under this prioritization scheme, we would use the size of overlap of languages as a measure.

Such a measure however, is open to objections that propositions that share fewer symbols might sometimes be more relevant than those that share more symbols *depending upon the symbols shared.* As an example, consider our intuition that if two formulas both mention *aardvarks*, they are (most likely) more relevant to each other than if they mentioned *cats*. Or, the fact that the authors of this paper are logicians is a closer relatedness relation than the fact that they are all human beings. The ordering $\beta_1 \preceq_\gamma \beta_2$ can be *further* refined by the following heuristic: if symbols shared by formulas occur frequently in σ, then there is a smaller likelihood that the formulas are relevant to each other whereas if they share symbols that occur with less frequency in the sequence σ. then the relevance is greater. This provides for a notion of *degree of relevance* amongst directly (and only for directly) relevant formulas.

Definition 12 *Consider formulas α, β, γ such that $L_\alpha \cap L_\beta = \Delta_1, L_\alpha \cap L_\gamma = \Delta_2$ i.e., α, β and α, γ are directly relevant. For a belief sequence σ, define integers m, n as follows:*

$$|\{\psi | \psi \in \sigma \wedge \Delta_1 \subseteq L_\psi\}| = m$$

$$|\{\phi | \phi \in \sigma \wedge \Delta_2 \subseteq L_\phi\}| = n$$

Definition 13 *The* degree *of direct relevance of α, β wrt σ,*

$$drel(\alpha, \beta, \sigma) = \frac{|\sigma|}{m}$$

A similar definition holds for the degree of direct relevance of α, γ w.r.t σ. Now, if $m > n$, then $drel(\alpha, \beta, \sigma) < drel(\alpha, \gamma, \sigma)$

In keeping with the ternary notion of relevance developed in this paper, we note that the degree of direct relevance between two formulas is dependent on the sequence σ and is not just a matter of syntactic or logical relations between the formulas in question. The effect of the definition above on the ordering created in the sequence by the query α will affect the definition of the inference operator $\vdash_k$. However, a fuller discussion would take us too far afield and we postpone such refinements for future work.

Belief Sets

We now define an attendant notion of a belief set $C_k(\sigma)$ at a given level of relevance, k:

Definition 14 $C_k(\sigma) = \{\gamma | \sigma \vdash_k \gamma\}$

Definition 15 $C(\sigma) = \{\gamma | \exists k, \sigma \vdash_k \gamma\} = \bigcup C_k(\sigma)$

Proposition 1 *The inference procedure defined above is monotonic in k, the degree of relevance, i.e., $\sigma \vdash_k \gamma \Rightarrow \sigma \vdash_{k+1} \gamma$.*

The above follows immediately from the fact that if $k < k'$ then $\Gamma_{<\sigma,k,p>} \subseteq \Gamma_{<\sigma,k',p>}$. Note also that if two formulas α, β have the same language, $L_\alpha = L_\beta$, then $\Gamma_{<\sigma,k,\alpha>} = \Gamma_{<\sigma,k,\beta>}$.

Since $C_k(\sigma) \subseteq C_{k+1}(\sigma)$, if we were not worried about computational costs then $C(\sigma)$ would be the only notion which would interest us and so we would only stop at a smaller k to conserve resources.

Definition 16 *If Δ is a set of propositional atoms, then* $C_\Delta = \{\alpha | L_\alpha = \Delta \wedge \alpha \in C(\sigma)\}$

Observation 3 *If Δ is a set of propositional atoms, then C_Δ will be consistent.*

The agent's responses to a particular subject matter are guaranteed to be consistent by the query answering scheme.

Remark: The inference procedure is of course *non-monotonic* in expansions of a belief sequence, i.e., if $\sigma \sqsubseteq \sigma'$ and $\sigma \vdash_k \gamma$, it is not necessarily the case that $\sigma' \vdash_k \gamma$. For example p is derivable from the sequence p, but not from, say $p; \neg(p \vee q)$. But if we revise by formulas that are 'irrelevant' to the query γ in question then they will have no effect on the derivation of γ and γ will still be derived from the longer sequence.

Remark: Our inference method corresponds to the *liberal* inference defined on a linearly prioritized sequence of formulas in [Geo96]. [Geo96] defines a *strict* notion of inference as well, which in our case would correspond to stopping the construction of $\Gamma_{\langle\sigma,k,\gamma\rangle}$ upon encountering the first formula that would make the set inconsistent.

Properties of $\vdash_k$

Proposition 2 *The following properties hold for the process of revision defined on sequences.*

- *Weak Inclusion: If $\gamma \neq \perp$ then $\sigma * \gamma \vdash_k \gamma$*

The following additional properties hold under the condition *that $L_\alpha = L_\beta$:*

- *Weak (or Cautious) Monotonicity:*

$$\frac{\sigma \vdash_k \alpha, \sigma \vdash_k \beta}{\sigma * \alpha \vdash_k \beta}$$

- *Rational Monotonicity:*

$$\frac{\sigma \not\vdash_k \neg\alpha, \sigma \vdash_k \beta}{\sigma * \alpha \vdash_k \beta}$$

- *Weak Cut:*

$$\frac{\sigma * \alpha \vdash_k \beta, \sigma \vdash_k \alpha}{\sigma \vdash_k \beta}$$

- *Adjunction:*

$$\frac{\sigma \vdash_k \alpha, \sigma \vdash_k \beta}{\sigma \vdash_k \alpha \wedge \beta}$$

- *Right Weakening:*

$$\frac{\sigma \vdash_k \alpha, \alpha \vdash \beta}{\sigma \vdash_k \beta}$$

The condition $L_\alpha = L_\beta$ is important for the following reason. As we have remarked, $[\![\sigma]\!]$ may be, and usually is, inconsistent. To get consistent answers to a query γ we restrict ourselves to a certain portion of $[\![\sigma]\!]$ and this portion will depend on the language L_γ. It is natural then that different γ, γ' with different languages will evoke different subsets of $[\![\sigma]\!]$. We cannot then require that these subsets will always produce coherent answers which obey our rules. However, in practice there will be *some* coherence.

Note too, that a property like Right Weakening is actually undesirable without the restriction $L_\alpha = L_\beta$. This is because in actual practice, if $L_\beta \subset L_\alpha$, then we do not want β to trigger the same part of the sequence that α does. The inference procedure depends crucially on the fact that we are using information which appears relevant to the query, and ignoring what is probably not relevant. To allow β and α to trigger the same formulas when their languages are not equal would vitiate this insight. In other words, we want to make use of the fact that while the general problem of classical inference is **NP**-complete, a particular instance of the inference problem can be solved faster by ignoring side information.

As an example, the general problem of determining a Hamiltonian circuit for a graph is **NP**-complete, but the problem of determining whether a particular circuit in a graph is Hamiltonian is solvable in polynomial time. Thus even though β is stronger than α, it may happen that β is k-derivable from σ and α is not. This is a particular instance of the general fact that an instance of an **NP**-complete problem is polynomially checkable when the witness is present.

For just one example that Weak Monotonicity can fail in general unless $L_\alpha = L_\beta$: let $\sigma = [p, \neg p \wedge \neg q, q]$, $\alpha = p \wedge q$, and $\beta = \neg p$. Now, $\forall k, \sigma \vdash_k p \wedge q$, but also, $\sigma \vdash_k \neg p$. But clearly, $\sigma * (p \wedge q) \not\vdash_k \neg p$. Thus weak monotonicity fails.

Similar examples can be constructed for the other rules mentioned above. From a practical point of view, if two queries α, β are asked on the *same* occasion, a smarter query answering procedure should use $L = L_\alpha \cup L_\beta$ as the language for determining direct relevance. The answers thus generated will be compatible with each other. *The decision to use the language is $L_\alpha \cup L_\beta$ is motivated by the fact that when queries α, β are asked on the same occasion we are in fact employing the context set $\Delta = L_\alpha \cup L_\beta$.*

'Loss' of Beliefs

New formulas can block the derivation of formulas which were derivable before and so provide a simple modeling for loss of belief in a proposition. After all, agents do not lose beliefs without a reason: to drop the belief that α is to *revise* by some information that changes our reasoning. Still, it is possible in our model to lose α without acquiring $\neg\alpha$. For example, consider the sequences $\sigma = p \wedge q$ and $\sigma * (\neg p \vee \neg q)$. The revised sequence no longer answers 'yes' to p but neither does it answer 'yes' to $\neg p$. $(\neg p \vee \neg q)$ has undermined $\sigma = p \wedge q$ without actually making $\neg p$ derivable. Suppose someone tells us that his name is Adam and that he is married. We later find out that his name is not Adam. We will drop both the belief that his name is Adam and that he is married. But we will not acquire the belief that he is unmarried.

Equivalence

Given the notion of inference defined above, we say that belief sequences which yield the same answers to all queries are *equivalent*. However, equivalence is not always preserved under extensions. Consider the following example: $\sigma_1 = [\neg p \wedge \neg q], \sigma_2 = [p, \neg p \wedge \neg q]$. σ_1, σ_2 are equivalent, and neither σ_1 nor σ_2 implies p as a conclusion. Revising by the formula $p \vee q$ yields p as a conclusion for σ_2 though not for σ_1. Given this observation, we would like to define a strong notion of equivalence which is unaffected by revision.

Definition 17

- *Sequences σ_1, σ_2 are equivalent if $C(\sigma_1) = C(\sigma_2)$.*
- *Sequences σ_1, σ_2 are strongly equivalent if $C(\sigma_1 * \alpha_1 * \ldots \alpha_n) = C(\sigma_1 * \alpha_1 * \ldots \alpha_n)$ for all sequences of revisions $\alpha_1, \ldots, \alpha_n$.*
- *σ_1 subsumes σ_2 if $C(\sigma_2) \subseteq C(\sigma_1)$.*

We suspect that the notion of strong equivalence is testable via a single revision, i.e., if σ and σ' are not strongly equivalent, then there is a γ such that $\sigma * \gamma$ and $\sigma' * \gamma$ are not equivalent.

An obvious, related, question is, How can all sequences be trimmed or reduced to their shortest equivalent form? The task of reducing sequences to their simplest equivalent form is most likely, computationally non-trivial.

Complexity of the Inference Procedure

In the methods defined above, there are two sources of complexity. For any query γ the first one involves the calculation of the smallest language L_γ, which is a co-**NP** complete problem [HR99], but only in the (relatively short) length of the individual formula γ and not in the size of the entire belief base σ. The second one involves checking the consistency of the set Γ_i at each step of the construction of the set Γ of the maxiconsistent set related to the query formula γ. However, for normal bases σ and small k it is likely to be the case that symbols which occur in formulas which are k-relevant to γ are few in number. In other words, there is a small language L' with $|L'| \ll |L|$ such that for any formula $\alpha \in \sigma$, with α k-relevant to γ, $L_\alpha \subseteq L'$. The consistency check then will be quite feasible.

Our answering method first calculates the relevance relation exhaustively for the entire sequence. Then the set Γ is constructed formula by formula. Since each stage of the construction involves a consistency check, the complexity of the procedure is polynomial with an **NP** oracle, but only in $|L'|$, which will be small if we keep k small. Moreover, in checking for k-derivability, the costs would be reduced sharply as most formulas in the sequence would not be k-relevant and the size of $\Gamma_{\langle\sigma,k,\gamma\rangle}$ could be quite small. This is indeed how we reason in practice. While the logical problem of deriving consequences from our (possibly inconsistent) beliefs is daunting, the notion of relevance cuts down sharply on the effort involved.

Comparison with existing frameworks

In the following sections we compare and contrast our method with existing frameworks for belief revision and inference from belief states.

Comparison with B-structures

In [CP99] a method for answering queries from B-structures is presented that allowed the answer $\top$, i.e., (*over-defined* or *inconsistent*). In our method of inference, we ensure that $\Gamma_{\langle\sigma,k,\gamma\rangle}$ is always consistent, thus preventing an inconsistent answer. The B-structure query answering method corresponds to $\vdash_0$, i.e., to $\Gamma_{<\sigma,0,\gamma>}$ constructed only with directly relevant formulas. In that case, the procedure of [CP99] will agree with our method when the former gave 'yes', 'no', or 'no information' answers.

Conformance with AGM axioms

It is natural to ask how our inference conforms to AGM-like postulates. We cannot apply the results in [Geo96] that show that maxiconsistent inference on a belief sequence is rational; our inference procedure uses relevance relations to bring about a reordering of our sequence, depending on the formula whose inference we are testing. This reordering can cause some logical complexities. However, some AGM-like properties do hold.

(σ^* 1) $\sigma * \gamma$ is a belief sequence
(σ^* 2) $\gamma \in C_k(\sigma * \gamma)$ provided that γ is consistent.
(σ^* 3) If $\gamma \Leftrightarrow \delta$ then $\forall \alpha, \sigma, \sigma * \gamma \vdash_k \alpha$ iff $\sigma * \delta \vdash_k \alpha$.

Note too, that there is no counterpart to the AGM expansion operation $\dot{+}$ in our model; revision just is concatenation and the content of the new epistemic state upon revision is given by the inference operation defined above. There is no counterpart to the AGM contraction operator either: contraction in belief sequences as noted previously is modeled by the blocking of a conclusion.

Conformance with Lehmann Axioms

In Lehmann's sequences model [Leh95], a belief state results from a sequence of revisions—the individual revisions themselves are just those by consistent formulas. A concatenation of these sequences is denoted by $\circ$ and sequences of length one are formulas[4]. For example, we can think of the sequence $\sigma \circ \alpha \circ \beta \circ \tau$ as denoting, first, the formulas of the sequence σ, then the formula α, then β and then the formulas of the sequence τ. The belief set that results from a sequence σ of individual revisions is denoted $[\sigma]$; $\circ$ denotes a revision procedure. The belief set $[\sigma \circ \alpha]$ denotes the result of revising $[\sigma]$ by the formula α. We do not however, identify $[\sigma \circ \alpha]$ with $[\sigma] * \alpha$ ($*$ is an AGM operator). The AGM postulates require that if $[\sigma] = [\tau]$, then $[\sigma] * \alpha = [\tau] * \alpha$, but in this framework $[\sigma] = [\tau]$ does not imply $[\tau \circ \alpha] = [\sigma \circ \alpha]$. Lehmann's framework allows a reasoning agent to base its revision *not just on the belief set* $[\sigma]$ *but also on the sequence itself i.e., on the history of the agent.* In this way, the temporal ordering of the beliefs plays a part in the revisions that the agent carries out. Given the definitions above, a set of axioms is presented for revisions of belief sequences:
(**I1**) $[\sigma]$ is a consistent theory
(**I2**) $\alpha \in [\sigma \circ \alpha]$
(**I3**) If $\beta \in [\sigma \circ \alpha]$, then $\alpha \rightarrow \beta \in [\sigma]$
(**I4**) If $\alpha \in [\sigma]$, then $[\sigma \circ \tau] = [\sigma \circ \alpha \circ \tau]$
(**I5**) If $\beta \models \alpha$, then $[\sigma \circ \alpha \circ \beta \circ \tau] = [\sigma \circ \beta \circ \tau]$
(**I6**) If $\neg\beta \notin [\sigma \circ \alpha]$, then $[\sigma \circ \alpha \circ \beta \circ \tau] = [\sigma \circ \alpha \circ \alpha \wedge \beta \circ \tau]$
(**I7**) $[\sigma \circ \neg\beta \circ \beta] \subseteq Cn([\sigma], \beta)$

We provide counterparts of (some of) the Lehmann axioms into our framework as follows:

- $C_\Delta(\sigma)$ is consistent for every Δ.
- $\alpha \in C(\sigma * \alpha)$ provided that α is consistent.
- If $\beta \in C(\sigma * \alpha)$, then $\alpha \rightarrow \beta \in C(\sigma)$ provided that $L_{\alpha\rightarrow\beta} = L_\beta$.

[4] We have modified Lehmann's original notation for greater readability.

- If $\beta \models \alpha$, then $C_\Delta(\sigma * \alpha * \beta * \gamma) = C_\Delta(\sigma * \beta * \gamma)$ provided that both L_α and L_β overlap Δ.
- $C_\Delta(\sigma * \neg\beta * \beta) \subseteq Cn(C_\Delta(\sigma), \beta)$.

Conformance with the Darwiche-Pearl Axioms

In [DP97] a framework for iterated belief revision is presented via the four axioms below ($\circ$ is the the update operator, α, μ represent new epistemic inputs, ψ represents a belief state).

- **(C1)** If $\alpha \models \mu$, then $(\psi \circ \mu) \circ \alpha \leftrightarrow \psi \circ \alpha$.
- **(C2)** If $\alpha \models \neg\mu$, then $(\psi \circ \mu) \circ \alpha \leftrightarrow \psi \circ \alpha$.
- **(C3)** If $\psi \circ \alpha \models \mu$, then $(\psi \circ \mu) \circ \alpha \models \mu$.
- **(C4)** If $\psi \circ \alpha \not\models \neg\mu$, then $(\psi \circ \mu) \circ \alpha \models \mu$.

We provide translations of the axioms above as follows:

- $C1$: If $\alpha \models \mu$ then $C_\Delta(\sigma * \mu * \alpha) = C_\Delta(\sigma * \alpha)$.
- $C2$: If $\alpha \models \neg\mu$ then $C_\Delta(\sigma * \alpha) = C_\Delta(\sigma * \mu * \alpha)$.
- $C3$: If $\mu \in C_\Delta(\sigma * \alpha)$ the $\mu \in C_\Delta(\sigma * \mu * \alpha)$.
- $C4$: If $\neg\mu \notin C_\Delta(\sigma * \alpha)$ then $\mu \in C_\Delta(\sigma * \mu * \alpha)$.

These translations all hold provided that L_α intersects Δ.

Conclusion

In the study above, we have presented a model for inference from belief states that relies upon a particular context set. The method of inference (and belief revision) described is easily implemented via a finite representation of the sequence. A full revision history is maintained and iterated revision is routine. Explicit attention is paid to temporal ordering and relevance relations. Our method extends intuitions found in the work on belief bases, sequences and inference relations based on prioritizations of beliefs and goes beyond by making clear the importance of a plausible notion of relevance to guide inference. The method presented is (for small k) resource conscious; in future work, we plan to implement this method and further investigate its properties.

References

[AGM85] C. Alchourron, P. Gärdenfors, and D. Makinson. On the logic of theory change: Partial meet functions for contraction and revision. *Journal of Symbolic Logic*, 50:510-530, 1985.

[AP99] Horacio Arlo-Costa and Rohit Parikh. Two place probabilities, beliefs and belief revision: on the foundations of iterative belief kinematics.. In *Proceedings of the Twelfth Amsterdam Colloquium*, 1999.

[Bre91] Gerhard Brewka. Belief revision in a framework for default reasoning. In Andre Fuhrmann and M. Morreau, editors *The Logic of Theory Change*, Lecture Notes in Artificial Intelligence, Number 465, pages 206-222, Springer, Berlin, 1991.

[CP99] Samir Chopra and Rohit Parikh. An inconsistency tolerant model for belief representation and belief revision. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence*, pages 192-197, Morgan-Kaufmann, 1999.

[DP97] Adnan Darwiche and Judea Pearl. On the logic of iterated belief revision. Artificial Intelligence, 89(1-2):1-29, January, 1997

[Doy79] Jon Doyle. A truth maintenance system. *Artificial Intelligence*, 12:231-272, 1979.

[Eps95] Richard Epstein. *The Semantical Foundations of Logic: Propositional Logics*. Oxford University Press, 1995.

[FL94] Michael Freund and Daniel Lehmann. Belief revision and rational inference Technical Report TR-ARP-94-16, The Leibniz Center for Research in Computer Science, Insititute of Computer Science, Hebrew University, July 1994, 1994.

[GM88] P. Gärdenfors and D. Makinson. Revisions of knowledge systems using epistemic entrenchment. In Moshe Vardi, editor, *Proceedings of Theoretical Aspects of Reasoning about Knowledge*, pages 83-96, 661-672, Morgan-Kaufmann, 1988.

[GR95] P. Gärdenfors and Hans Rott. Belief revision. In Dov M. Gabbay, Christpher J. Hogger and John A. Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming, Volume IV: Epistemic and Temporal Reasoning* , pages 35-132. Oxford University Press, Cambridge, 1995.

[Geo96] Konstantinos Georgatos. Ordering based representations of rational inference. In J.J Alferes, L.M. Pereira, and E. Orlowska, editors, *Logics in Artificial Intelligence (JELIA '96)*, Lecture Notes in Artifcial Intelligence, page s 176-191, Springer, Berlin, 1996.

[Geo97] Konstantinos Georgatos. Entrenchment relations: A uniform approach to non-monotonicity. In *Proceedings of the International Joint Conference on Qualitative and Quantitative Practical Reasoning (ESCQARU/FAPR 97)*, Lecture Notes in Computer Science, Number 1244, pages 282-297, Springer, Berlin, 1997.

[Han92] Sven Ove Hansson. In defense of base contraction. *Synthese*, 91:239-245, 1992.

[HR99] Andreas Herzig and Omar Rifi. Propositional belief base update and minimal change. *Artificial Intelligence*, 115(1):107-138, 1999.

[Leh95] Daniel Lehmann. Belief revision, revised. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, pages 1534-1540, 1995.

[NFPS96] Abhaya Nayak, Norman Foo, Maurice Pagnucco and Abdul Sattar. Changing conditional beliefs unconditionally. In *Proceedings of the Sixth Conference on Theoretical Aspects of Rationality and Knowledge*, pages 119-135, De Zeeuwse Stromen, The Netherlands, Morgan-Kaufmann, 1996.

[Neb92] Bernhard Nebel. Syntax based approaches to belief revision. In Peter Gärdenfors, editor, *Belief revision*, Cambridge Tracts in Theoretical Computer Science, Cambridge University Press, Cambridge, 1992.

[Par96] Rohit Parikh. Beliefs, belief revision and splitting languages. In Lawrence Moss, Jonathan Ginzburg and Maarten de Rijke, editors, *Logic, Language and Computation, Volume 2*, CSLI Lecture Notes No. 96, pages 266-278, CSLI Publi cations, 1999. Initially presented in *Preliminary Proceedings of Information Theoretic Approaches to Logic, Language and Computation*, 1996.

[Rod97] *A methodology for iterated information change*, Ph.D thesis. Imperial College, University of London, 1997.

[Ry91] Mark D. Ryan. Ordered theory Presentations. In *Proceedings of the 8th Amsterdam Colloquium*, 1991.

[Was99] Renata Wassermann. Resource bounded belief revision. *Erkenntnis*, 1999. To appear.