
Approximate Belief Revision - Preliminary Report

SAMIR CHOPRA, *CUNY Graduate Center, New York, USA. E-mail:*
schopra@gc.cuny.edu

ROHIT PARIKH, *CUNY Graduate Center, New York, USA. E-mail:*
ripbc@cunyvm.cuny.edu

RENATA WASSERMANN, *Department of Computer Science, Institute
of Mathematics and Statistics, University of São Paulo, Brazil. E-mail:*
renata@ime.usp.br

Abstract

The standard theory for belief revision [1] provides an elegant and powerful framework for reasoning about how a rational agent should change its beliefs when confronted with new information. However, the agents considered are extremely idealized. Some recent models attempt to tackle the problem of plausible belief revision by adding structure to the belief bases and using nonstandard inference operations. One of the key ideas is that not all of an agent's beliefs are relevant for an operation of belief change.

In this paper we incorporate the insights pertaining to local change and relevance sensitivity with the use of *approximate inference relations* [13]. These approximate inference relations offer us partial solutions at any stage of the revision process. The quality of the approximations improves as we allow for more and more resources to be used. We are provided with upper and lower bounds to what would be obtained with the use of classical inference.

We briefly summarize the concepts of local change and approximate inference and argue for heuristics based on relevance to guide the approximations. We show how approximate inference can be integrated with the local change paradigm.

Keywords: Belief Change, Approximate Reasoning

1 Introduction

Belief revision, the process of transforming a belief state on receipt of new information, is an integral part of any attempt to provide plausible architectures for automated agents. Conventional AGM-style revision [1] provides an elegant and powerful framework for reasoning about how a rational agent should change its beliefs when confronted with new information, but it tells us very little about *how* that agent could really perform such belief changes. Moreover, the rational agent described is a highly idealized one, a perfect reasoner with unbounded memory, logical ability, no inconsistent beliefs and no time constraints.

In the AGM paradigm, the belief state of an agent is represented by a *belief set*, i.e., a logically closed, consistent theory. Such an approach suffers from problems associated with representational infeasibility and computational intractability. To alleviate these problems, *belief base* revision offers open, finite representations. This

2 Approximate Belief Revision - Preliminary Report

alternative has been extensively studied and AGM-like operations have been defined for belief bases. Although the use of belief bases solves the problem of representing a belief state, belief bases are typically quite large and the belief change operations make use of computationally expensive consistency checks.

One way to attack the problem is to try to reduce the *size* of the set to be explored. Intuitively, not all of an agent's beliefs are relevant for deciding what to do with new information. There should be a way of isolating the subset of a belief base that contains the *relevant* beliefs for a query or an operation of belief change. Recent models such as [12], [8], [4], [16], attempt to tackle the problem of plausible belief revision by using nonstandard inference operations such as local change operators and offering structuring relations on belief bases via relevance sensitivity. In these frameworks for local reasoning and belief change, the key idea is that not all of an agent's beliefs are relevant for an operation of belief change. As an example, suppose that I wake up and see that the sun is shining in Amsterdam. This contradicts my previous beliefs about the weather, like the one that it always rains in Holland, and leads to revision. However, the revision process does not have to take my beliefs about mathematics into account. These beliefs are completely irrelevant for the change taking place and should not play any role in it.

In this paper we incorporate the insights pertaining to local change and relevance sensitivity with the use of *approximate inference relations* [13]. These approximate inference relations offer us *partial solutions* at any stage of the revision process. The quality of the approximations improves as we allow for more and more resources to be used. The process of belief revision is rendered an approximate one, with the understanding that at any given stage, we are provided with upper and lower bounds to what would be obtained with the use of classical inference.

The outline of this paper is as follows: we first briefly summarize the concepts of local change and approximate inference. We then argue for heuristics based on relevance to guide the approximations. We show how approximate inference can be integrated with the local change paradigm. We close the paper with a discussion about the gains of using the two different approximations proposed in [13]. Due to the lack of space, most proofs are left for the full version of the paper.

Notation: L is a finite propositional language with the usual logical connectives ($\neg, \vee, \wedge, \rightarrow, \leftrightarrow$). The constants *true*, *false* are in L ; $\alpha, \beta, \gamma \dots$ denote arbitrary formulae; $p, q, r \dots$ denote propositional atoms; $A, B \dots$ denote sets of formulae. For any formula α , $L(\alpha)$ is the set of propositional atoms in α . Cn stands for classical consequence, C stands for arbitrary inference operations.

2 Local Change

In this session, we briefly describe the paradigm for local change proposed in [8]. The main idea is that if we can isolate the part of a belief base which is relevant for a given change, we can apply the belief change operations to this subset of the base and keep the rest as it was.

In [8], given two sets of formulas A and B , the relevant formulas in B for A are defined to be those that contribute to proving or disproving any formula of A . This is formalized using the notion of *kernel sets* - minimal subsets implying a given sentence:

DEFINITION 2.1

[7] Let C be an inference operation on \mathcal{L} . Then the kernel operation $\perp\!\!\!\perp_C$ is the operation such that for every set B and every formula α , X is an element of $B \perp\!\!\!\perp_C \alpha$ if and only if:

1. $X \subseteq B$
2. $\alpha \in C(X)$
3. for all Y , if $Y \subset X$ then $\alpha \notin C(Y)$

The elements of $B \perp\!\!\!\perp_C \alpha$ are called α -kernels.

We write $\perp\!\!\!\perp$ as an abbreviation of $\perp\!\!\!\perp_{Cn}$, where Cn is a classical consequence operation.

DEFINITION 2.2

[8] The A -compartment of B , where A and B are sets of sentences, is defined as:

$$c(A, B) = \bigcup_{\alpha \in A} (\bigcup ((B \perp\!\!\!\perp \alpha) \cup (B \perp\!\!\!\perp \neg\alpha) \setminus (B \perp\!\!\!\perp \perp)))$$

We call c the compartmentalization function for the sets A and B .

The compartments defined above are usually overlapping, and should not be seen as a partition of the belief base into different topics or subjects.

The *localization* of a given inference operation C to A is given by:

DEFINITION 2.3

[8] Let C be an inference operation on \mathcal{L} and let c be a compartmentalization function as in definition 2.2. Then for any set A , the A -localization of C is the inference operation C_A such that for all sets B of sentences: $C_A(B) = C(c(A, B))$.

A set B is *A -locally consistent* if and only if $\perp \notin C_A(B)$.

Local versions of the operations of contraction, revision, consolidation and semi-revision can be constructed and characterized by means of postulates. We reproduce below the characterization of *local kernel contraction* [8], which can be used for deriving the others. In kernel contraction, if we remove from the belief base at least one element of each α -kernel, we obtain a belief base that does not imply α . In the case of the local operation, we want a belief base that does not A -locally imply α , where A is a set of sentences. We consider only those kernel sets that A -locally imply α and use an *incision function* to select the formulas from the kernel sets to be removed.

DEFINITION 2.4

[7] An *incision function* is any function $\sigma : \mathcal{P}(\mathcal{P}(\mathcal{L})) \rightarrow \mathcal{P}(\mathcal{L})$ such that for any subset S of $\mathcal{P}(\mathcal{L})$:

1. $\sigma(S) \subseteq \bigcup S$
2. If $\emptyset \neq X \in S$, then $X \cap \sigma(S) \neq \emptyset$

Example: Let $S = \{\{a, b\}, \{b, c\}, \{a, d, e\}\}$, (one possible) $\sigma(S) = \{b, c, e\}$.

DEFINITION 2.5

[7] Let C be an inference operation on \mathcal{L} and σ an incision function. The local kernel contraction of B determined by C and σ is the operation $\dot{-}_{C, \sigma}$ such that for all sets of sentences B :

$$B \dot{-}_{C, \sigma} \alpha = B \setminus \sigma(B \perp\!\!\!\perp_C \alpha)$$

4 Approximate Belief Revision - Preliminary Report

The following theorem characterizes the operation of local kernel contraction and is a generalization of the result obtained in [7]. The intended interpretation is that C is a localization C_A of an inference operator. However, the formal result is more general:

THEOREM 2.6

[8] Let C be an inference operation satisfying monotony and compactness. An operation $\dot{-}$ is an operation of local kernel contraction determined by C and some incision function if and only if for all sets of sentences B :

- If $\alpha \notin C(\emptyset)$, then $\alpha \notin C(B \dot{-} \alpha)$ (success)
- $B \dot{-} \alpha \subseteq B$ (inclusion)
- If $x \in B \setminus B \dot{-} \alpha$, then there is some $B' \subseteq B$ such that $\alpha \notin C(B')$ and $\alpha \in C(B' \cup \{x\})$ (core-retainment)
- If for all subsets B' of B , $\alpha \in C(B')$ if and only if $\beta \in C(B')$, then $B \dot{-} \alpha = B \dot{-} \beta$ (uniformity)

Sets of postulates and representation results are given for operations of partial meet contraction, kernel revision and partial meet revision. Besides compactness and monotony, two other properties are needed for the characterization of revision:

- **consistency** $\perp \notin C(\emptyset)$.
- **α -local non-contravention** if $\neg\alpha \in C(B \cup \{\alpha\})$, then $\neg\alpha \in C(B)$.

The construction of the compartments required for the local change inference operations is computationally expensive. Local change operations are computationally as costly as the original versions, although intuitively more appropriate. However, the representation results proven in [8] do not rely on the way the compartments are defined, but only on properties of the local inference operation obtained, namely, monotony, compactness, consistency, and α -local non-contravention.

We now turn to more efficient ways in which to isolate a meaningful part of a belief base for a change operation.

3 Approximate entailment

In this section, we summarize the main results concerning approximate inferences.

In [13], Schaerf and Cadoli define two approximations of classical entailment: \models_S^1 which is complete but not sound and \models_S^3 which is sound and incomplete. These approximations are carried out over a set of atoms $S \subseteq L$ which determines their closeness to classical entailment. In the trivial extreme of approximate entailment, i.e., when $S = L$, classical entailment is obtained. At the other extreme, when $S = \emptyset$, \models_S^1 holds for any two formulas and \models_S^3 corresponds to Levesque's logic for explicit beliefs [11], which bears a strong resemblance to Anderson and Belnap's relevance logic.

In an S_1 assignment, if $x \in S$, then $x, \neg x$ are given opposite truth values; if $x \notin S$, then $x, \neg x$ both get the value 0. In an S_3 assignment, if $x \in S$, then $x, \neg x$ get opposite truth values, while if $x \notin S$, $x, \neg x$ do not both get 0, but may both get 1. The belief base B is assumed to be in clausal form.¹

¹All results can be extended for negation normal form cf. [2], but further generalization implies missing the

The following examples illustrate the use of approximate entailment. Since \models_S^3 is sound but incomplete, it can be used to approximate \models , i.e., if for some S we have that $B \models_S^3 \alpha$, then $B \models \alpha$. On the other hand, since \models_S^1 is unsound but complete, it can be used for approximating $\not\models$, i.e., if for some S we have that $B \not\models_S^1 \alpha$, then $B \not\models \alpha$.

EXAMPLE 3.1

[13]

$L = \{\text{cow, grass-eater, dog, carnivore, has-canine-teeth, mammal, has-molar-teeth, vertebrate, animal}\}$.

$$B = \{ \neg \text{cow} \vee \text{grass-eater}, \\ \neg \text{dog} \vee \text{carnivore}, \\ \neg \text{grass-eater} \vee \neg \text{has-canine-teeth}, \\ \neg \text{grass-eater} \vee \text{mammal}, \\ \neg \text{carnivore} \vee \text{mammal}, \\ \neg \text{mammal} \vee \text{has-canine-teeth} \vee \text{has-molar-teeth}, \\ \neg \text{mammal} \vee \text{vertebrate}, \\ \neg \text{vertebrate} \vee \text{animal} \}.$$

We want to check whether $B \models \alpha$, where $\alpha = \neg \text{cow} \vee \text{has-molar-teeth}$.

For $S = \{\text{grass-eater, mammal, has-canine-teeth}\}$, we have that $B \models_S^3 \alpha$, and hence $B \models \alpha$.

EXAMPLE 3.2

[13] $L = \{\text{person, child, youngster, adult, senior, student, pensioner, worker, unemployed}\}$.

$$B = \{ \neg \text{person} \vee \text{child} \vee \text{youngster} \vee \text{adult} \vee \text{senior}, \\ \neg \text{youngster} \vee \text{student} \vee \text{worker}, \\ \neg \text{adult} \vee \text{student} \vee \text{worker} \vee \text{unemployed}, \\ \neg \text{senior} \vee \text{pensioner} \vee \text{worker}, \\ \neg \text{student} \vee \text{child} \vee \text{youngster} \vee \text{adult}, \\ \neg \text{pensioner} \vee \text{senior}, \\ \neg \text{pensioner} \vee \neg \text{student}, \\ \neg \text{pensioner} \vee \neg \text{worker} \}.$$

We want to check whether $B \not\models \beta$, where $\beta = \neg \text{child} \vee \text{pensioner}$.

For $S = \{\text{child, worker, pensioner}\}$, we have that $B \not\models_S^1 \beta$, and hence $B \not\models \beta$.

Note that in both examples above, S is a small part of the total language L . The approximation of classical inference is made via a simplification of the belief base B as follows (for a given conclusion α and context set S):

LEMMA 3.3

[13] Let $\text{simplify-}\mathcal{B}(B, S)$ be the result of deleting all clauses of B which contain an atom outside S . Then B is S_3 -satisfiable if and only if $\text{simplify-}\mathcal{B}(B, S)$ is classically satisfiable.

good complexity upper bounds, as shown in [3]. Besides the increase in complexity, the standard translation of formulas into NNF makes use of De Morgan's law and does not preserve truth values under Schaerf and Cadoli's non-standard semantics.

6 Approximate Belief Revision - Preliminary Report

LEMMA 3.4

[13] Let $simplify-1(B, S)$ be the result of deleting all literals of B which mention atoms outside S . B is S_1 -satisfiable if and only if $simplify-1(B, S)$ is classically satisfiable.

LEMMA 3.5

Let i be 1 or 3. If $B \subseteq B'$, then $simplify-i(B', S) = simplify-i(B, S) \cup simplify-i(B' \setminus B, S)$.

Schaerf and Cadoli then obtain the following results for tractable approximated inference:

THEOREM 3.6

[13] There exists an algorithm for deciding if $B \models_S^3 \alpha$ and deciding $B \models_S^1 \alpha$ which runs in $O(|B| \cdot |\alpha| \cdot 2^{|S|})$ time.

The result above depends on a poly-time satisfiability algorithm for belief bases and formulas in clausal form alone. This result has been extended in [2] for formulas in negation normal form, but is not extendable to formulas in arbitrary forms [3].

4 Using Relevance to Guide the Approximations

The main problem with the use of approximate entailment, as noted in [13] and in [15], is that there is no general strategy or heuristic for choosing the right set S for the approximation. In this section, we argue for the use of a notion of *relevance* in order to guide the choice of the context set S .

It has been shown in [5] and [16] that one way in which relevance in a database can be defined is by looking at the syntactical structure of the formulas. In these studies, two formulas were considered relevant to each other if they shared an atom; the definition was generalized to the case when formulas did not share an atom with each other but did so with other formulas that in turn shared atoms. This enabled the definition of a notion of degrees or levels of relevance. In the present paper, we will use a notion of relevance between propositional variables: two atoms are relevant to each other if and only if they occur in the same formula in the belief base.²

DEFINITION 4.1

Given a belief base B , two atoms p, q are directly relevant (denoted by $R(p, q, B)$) if and only if there is a formula $\alpha \in B$ such that $p, q \in L(\alpha)$. A pair of atoms, p, q are k -relevant w.r.t B if $\exists p_1, p_2, \dots, p_k \in L$ such that:

- i) p, p_1 are directly relevant
- ii) p_i, p_{i+1} are directly relevant, $i = 1, \dots, k - 1$
- iii) p_k, q are directly relevant.

We first look at some necessary conditions that the context set S must satisfy. As shown in [14]³, if we want to use S_1 interpretations to approximate $B \models \alpha$, S must contain at least $L(\alpha)$ and at least one propositional variable occurring in each of the clauses of B . For S_3 interpretations, no such requirement is needed.

The following strategy may then be used in the construction of the set S . If we are considering S_1 entailments, we start with S being such that $L(\alpha) \subseteq S$ and for every

²This is used in RABIT [6] and cited in [16].

³Theorems 2 and 3.

formula $\gamma \in B$, $L(\gamma) \cap S \neq \emptyset$. For S_3 entailment, lemma 5.1 from [13] says that $L(\alpha)$ can be left out of S . Thus, we start with S being the set of propositional variables that are directly relevant to the elements of $L(\alpha)$. We calculate the approximations and check whether the result is satisfactory, i.e., whether $B \not\models_S^1 \alpha$ or $B \models_S^3 \alpha$. If it is not and we still have resources available to proceed the computation, we repeat the calculation adding to S the set of atoms that are directly relevant to the elements of S .

Construction strategy for S :

1. For S_1 entailment, start with S containing $L(\alpha)$ and at least one atom from each formula in B ; for S_3 , start with $S = \{p|p \text{ directly relevant to } q \text{ and } q \in L(\alpha)\}$.
2. Test entailments.
3. If the result is not satisfactory and resources are available, add to S the atoms relevant to elements in S and go to 2.
4. Return approximations.

If we look back to Example 3.1, we see that we start with S being the set of atoms directly relevant to `cow` and `has-molar-teeth`, i.e., $S = \{\text{grass-eater, mammal, has-canine-teeth}\}$.

Applying this technique to example 3.2, we start with an initial $S = L(\alpha)$. Since $L(\alpha)$ does not share atoms with all clauses in B , we have a choice of adding either `worker`, `student` or both `youngster` and `unemployed`, or both `youngster` and `adult` to S . If we start with 1-relevant, S is just the entire language L except for `unemployed`. Clearly, such a result is not satisfactory since our strategy has not provided us with a useful heuristic for minimizing the size of the set S . In session 6, we suggest an improvement to this strategy.

The construction of the set S and the level of relevance of formulas checked in the construction of S is a function of the length of proofs required for the conclusion in question. A query that requires a longer proof will require a deeper search in terms of levels of relevance for the members of S .

EXAMPLE 4.2

Take $B = \{p, p \rightarrow q, q \rightarrow r, r \rightarrow s, s \rightarrow t\}, \alpha = t$. Now, $S = \{p, q, r, s\}$, is constructed by collecting all atoms up to 3-degrees of relevance and the degree of relevance which we need to consider is a function of the length of proof required to derive α from B .

Note, however, that belief bases of the form above are likely to be extremely rare in practice and therefore, we may expect that in most cases, S will be constructed fairly quickly i.e., with small k . In any case, the fact that it takes a larger S to find a long proof is in line with our intuition that long proofs are only found if the agent has enough resources and can, therefore, reason with a larger context set S .

PROPOSITION 4.3

Using the above strategy, S_1 and S_3 entailment converge in a finite number of steps, i.e., after a finite number of iterations, we have that either $B \not\models_S^1 \alpha$, and hence $B \not\models \alpha$ or $B \models_S^3 \alpha$ and hence $B \models \alpha$.

It is not hard to see that in most cases, our strategy will return a set S which is larger than what is needed. As noted in [15], if we have domain information about

the particular problem we are trying to solve, specific heuristics can be used to limit the context set S . The particular notion of relevance that we use in this paper is to be used in the general case, when no domain information is available. In the case of particular applications, usually there is a better notion of relevance that can be used (see for example [17]). If B is in clausal form, we can use dependence links as used in logic programming [10].

5 Approximating Belief Change Operators

We now show that the notion of approximate entailment described in the previous section can be used for implementing local change efficiently. As we have noted above in section 2, if an inference operation satisfies monotony, compactness, consistency, and α -local non-contravention, it can be used together with the axiomatizations given in [8]. We now show that approximate entailment satisfies these requirements.

DEFINITION 5.1

$C_S^i(B) = \{\alpha \mid B \models_S^i \alpha\}$, where S is a set of propositional variables and $i \in \{1, 3\}$.

PROPOSITION 5.2

Let S be a fixed set of propositional variables. Then C_S^3 satisfies monotony, compactness, consistency, and α -local non-contravention for every formula α .

Proof: Monotony follows from the monotony of classical entailment together with lemmas 3.3 and 3.5. Compactness holds since we only consider finite sets. Consistency of C_S^3 follows from consistency of classical consequence, since $C_S^3(\emptyset) = Cn(\emptyset)$. If $L(\alpha) \subseteq S$, then α -local non-contravention follows directly from the non-contravention property of Cn . Suppose $L(\alpha) \not\subseteq S$. If all S_3 interpretations that assign 1 to B also assign 1 to $\neg\alpha$, then α -local non-contravention holds trivially. So let t be an S_3 interpretation such that $v(t, B) = 1$ and $v(t, \neg\alpha) = 0$. Then $v(t, \alpha) = 1$ and the property holds. \square

PROPOSITION 5.3

Let S be a fixed set of propositional variables. Then C_S^1 satisfies monotony, compactness, consistency, and α -local non-contravention iff $L(\alpha) \subseteq S$.

Proof: Monotony follows from the monotony of classical entailment together with lemmas 3.4 and 3.5. Compactness holds since we only consider finite sets. Consistency of C_S^1 follows from consistency of classical consequence, since $C_S^1(\emptyset) = Cn(\emptyset)$. If $L(\alpha) \subseteq S$, then α -local non-contravention follows directly from the non-contravention property of Cn . Suppose $L(\alpha) \not\subseteq S$. If all S_1 interpretation that assign 1 to B also assign 1 to $\neg\alpha$, then α -local non-contravention holds trivially. So let t be an S_1 interpretation such that $v(t, B) = 1$ and $v(t, \neg\alpha) = 0$. Then $v(t, \alpha) = 0$ and the property fails. This shows that $L(\alpha) \subseteq S$ is a necessary condition. \square

As an example of how C_S^1 can fail to satisfy the conditions above when $L(\alpha) \not\subseteq S$, take $\alpha = p$, $B = \emptyset$, and $p \notin S$. Then any S_1 interpretation assigns both α and $\neg\alpha$ the value 0. We have that $\neg\alpha \in C_S^1(B \cup \{\alpha\}) = Cn(\perp)$, but $\neg\alpha \notin C_S^1(B) = Cn(\emptyset)$.

The propositions above imply that both inference operations can be used to define local operations of belief change. This allows us to combine the computational

efficiency of Schaerf and Cadoli's method for approximate entailment with the logical characterization of belief change operations given in [8]. The idea is that these approximate inferences give us an approximation of the revised belief base and the larger the set S grows, the closer we get to the classical definition. At any point in the revision process, we can stop to check our results and depending upon resource availability, we can choose to carry on the revision process further or stop. To this end, we define the following contraction operations:

DEFINITION 5.4

Let C_S^i be an inference operation on \mathcal{L} and σ an incision function. The approximate local kernel contraction of B determined by C_S^i and σ is the operation $\dot{-}_{C_S^i, \sigma}$ such that for all sets of sentences B : $B \dot{-}_{C_S^i, \sigma} \alpha = B \setminus \sigma(B \perp_{C_S^i} \alpha)$

Note that in keeping with the approximated inference operation used above, we will obtain two kinds of contraction operations. Operations based on S_1 entailment will model more radical contraction operations (leading to a greater loss of beliefs) and those based on S_3 will model cautious operations (fewer beliefs than warranted are dropped). The choice of which contraction operation to use will be a context-sensitive one. At any stage of the approximated revision process, we know that contraction based on classical consequence lies between contraction based on approximated consequence relations, and that adding more propositional variables to the context S provides a better approximation. If $S = L$, then $B \dot{-}_{C_S^i, \sigma} \alpha = B \dot{-}_{C_n, \sigma} \alpha$. Usually, the approximations will converge for a proper subset S of L .

DEFINITION 5.5

Let \mathcal{X} and \mathcal{Y} be two families of sets. We say that $\mathcal{X} \sqsubseteq \mathcal{Y}$ if and only if for all $Y \in \mathcal{Y}$ there is $X \in \mathcal{X}$ such that $X \subseteq Y$.

PROPOSITION 5.6

$(B \perp_{C_S^1} \alpha) \sqsubseteq (B \perp_{C_n} \alpha) \sqsubseteq (B \perp_{C_S^3} \alpha)$.

It is not difficult to see that since C_S^1 is complete but not sound, there will be at least as many C_S^1 kernels as there are classical ones. However, some of the sets which are kernels according to C_S^1 may not be subsets of any classical kernel. A similar relation holds for classical and C_S^3 kernels. An example of these inclusion relations is provided in Section 5.1 below. The following result, a special case of theorem 2.6, characterizes the approximations thus obtained:

THEOREM 5.7

An operation $\dot{-}_S^i$ is an operation of local kernel contraction determined by C_S^i and some incision function if and only if for all sets of sentences B :

- If $\alpha \notin C_S^i(\emptyset)$, then $\alpha \notin C_S^i(B \dot{-}_S^i \alpha)$ (success)
- $B \dot{-}_S^i \alpha \subseteq B$ (inclusion)
- If $x \in B \setminus B \dot{-}_S^i \alpha$, then there is some $B' \subseteq B$ such that $\alpha \notin C_S^i(B')$ and $\alpha \in C_S^i(B' \cup \{x\})$ (core-retainment)
- If for all subsets B' of B $\alpha \in C_S^i(B')$ if and only if $\beta \in C_S^i(B')$, then $B \dot{-}_S^i \alpha = B \dot{-}_S^i \beta$ (uniformity)

In the model that we have presented above, belief revision is viewed as an approximated process, arrived at by a combination of the revision schemes presented above. Given that the approximate inference operations satisfy the conditions required for the general results obtained in [8], we note that similar results will obtain for the operations of partial meet contraction, partial meet and kernel revision.

5.1 Example

We now present an example that demonstrates the use of the relevance based heuristic for constructing a context-set S and the application of approximated inference to local kernel contraction. The following represents our belief base about a young student, Hans.

$L = \{\text{student, pensioner, young, six-feet-tall, likes-dancing, blue-eyes, worker}\}.$

$B = \{\text{student, } \neg\text{student} \vee \text{young, } \neg\text{young} \vee \neg\text{pensioner, worker, } \neg\text{worker} \vee \neg\text{pensioner, blue-eyes, likes-dancing, six-feet-tall}\}.$

We now find out that Hans is a young genius who made millions of dollars in the stock market and is now happily retired and drawing a pension. As a result, we want to contract B by $\neg\text{pensioner}$. The two tables below show how B gets simplified as S grows.

B	$S = \{\text{pensioner}\}$	$S = \{\text{pensioner, young, worker}\}$	$S = \{\text{pensioner, young, worker, student}\}$
student	\perp	\perp	student
$\neg\text{student} \vee \text{young}$	\perp	young	$\neg\text{student} \vee \text{young}$
$\neg\text{young} \vee \neg\text{pensioner}$	$\neg\text{pensioner}$	$\neg\text{young} \vee \neg\text{pensioner}$	$\neg\text{young} \vee \neg\text{pensioner}$
worker	\perp	worker	worker
$\neg\text{worker} \vee \neg\text{pensioner}$	$\neg\text{pensioner}$	$\neg\text{worker} \vee \neg\text{pensioner}$	$\neg\text{worker} \vee \neg\text{pensioner}$
blue-eyes	\perp	\perp	\perp
likes-dancing	\perp	\perp	\perp
six-feet-tall	\perp	\perp	\perp

TABLE 1. Simplify-1 applied to B with different values of S .

Table 1 shows the result of simplifying the belief base B according to S_1 using different context sets. The second column of table 1 shows the result of *simplify-1*(B , $\{\text{pensioner}\}$). Since all formulas which are reduced to inconsistency are minimal sets S_1 -implying anything, this gives us the following set of kernels:

$B \perp_{C_5^1} \neg\text{pensioner} = \{\{\text{student}\}, \{\neg\text{student} \vee \text{young}\}, \{\neg\text{young} \vee \neg\text{pensioner}\}, \{\text{worker}\}, \{\neg\text{worker} \vee \neg\text{pensioner}\}, \{\text{blue-eyes}\}, \{\text{likes-dancing}\}, \{\text{six-feet-tall}\}\}.$

As S grows, some S_1 -kernels are eliminated. The third column, where $S = \{\text{pensioner, young, worker}\}$, gives:

$$B \perp\!\!\!\perp_{C_S^1} \neg \text{pensioner} = \{\{\text{student}\}, \{\neg \text{student} \vee \text{young}, \neg \text{young} \vee \neg \text{pensioner}\}, \{\text{worker}, \neg \text{worker} \vee \neg \text{pensioner}\}, \{\text{blue-eyes}\}, \{\text{likes-dancing}\}, \{\text{six-feet-tall}\}\}.$$

Finally, the fourth column, where $S = \{\text{pensioner, young, worker, student}\}$ gives:

$$B \perp\!\!\!\perp_{C_S^1} \neg \text{pensioner} = \{\{\text{student}, \neg \text{student} \vee \text{young}, \neg \text{young} \vee \neg \text{pensioner}\}, \{\text{worker}, \neg \text{worker} \vee \neg \text{pensioner}\}, \{\text{blue-eyes}\}, \{\text{likes-dancing}\}, \{\text{six-feet-tall}\}\}.$$

B	$S = \{\text{pensioner}\}$	$S = \{\text{pensioner, young, worker}\}$	$S = \{\text{pensioner, young, worker, student}\}$
student	-	-	student
$\neg \text{student} \vee \text{young}$	-	-	$\neg \text{student} \vee \text{young}$
$\neg \text{young} \vee \neg \text{pensioner}$	-	$\neg \text{young} \vee \neg \text{pensioner}$	$\neg \text{young} \vee \neg \text{pensioner}$
worker	-	worker	worker
$\neg \text{worker} \vee \neg \text{pensioner}$	-	$\neg \text{worker} \vee \neg \text{pensioner}$	$\neg \text{worker} \vee \neg \text{pensioner}$
blue-eyes	-	-	-
likes-dancing	-	-	-
six-feet-tall	-	-	-

TABLE 2. Simplify-3 applied to B with different values of S .

Table 2 shows the result of simplifying B according to S_3 . The second column of table 2 shows the result of $\text{simplify-}\beta(B, \{\text{pensioner}\})$. This gives us:

$$B \perp\!\!\!\perp_{C_S^3} \neg \text{pensioner} = \emptyset.$$

The third column, where $S = \{\text{pensioner, young, worker}\}$ gives:

$$B \perp\!\!\!\perp_{C_S^3} \neg \text{pensioner} = \{\{\text{worker}, \neg \text{worker} \vee \neg \text{pensioner}\}\}.$$

Finally, the fourth column, where $S = \{\text{pensioner, young, worker, student}\}$ gives:

$$B \perp\!\!\!\perp_{C_S^3} \neg \text{pensioner} = \{\{\text{student}, \neg \text{student} \vee \text{young}, \neg \text{young} \vee \neg \text{pensioner}\}, \{\text{worker}, \neg \text{worker} \vee \neg \text{pensioner}\}\}.$$

Our example above serves as an illustration of proposition 5.6 to obtain upper and lower bounds for classical kernels via approximated kernels. Using an appropriate incision function would provide us with similar upper and lower bounds for the contracted belief base.

6 Alternative Heuristics

As we have seen in Section 4, for S_1 , we have to start the approximation with a set S which intersects all formulas of the belief base in order to guarantee convergence. There are many possible such sets and it is not trivial how to select one. For a special sort of belief base, this is not necessary:

DEFINITION 6.1

A belief base B is connected iff any two atoms $p, q \in L(B)$ are k -relevant for some k .

PROPOSITION 6.2

If B is connected, then our strategy for S_1 approximation converges to the classical result for the initial context set $S = L(\alpha)$.

We illustrate the necessity of the connectedness condition via the example in Section 5.1. It is easy to see that the initial base is not connected. The last three formulas (**blue-eyes**, **likes-dancing**, **six-feet-tall**) are completely disconnected from the rest. They are treated in the appropriate way by S_3 approximations, i.e., they are ignored. On the other hand, *simplify-1*(B, S) renders them inconsistent whenever S does not mention their atoms. This means that these irrelevant formulas will always generate inconsistent kernels which are singletons and therefore will be needlessly deleted by a contraction operation.

We can choose to pre-process the belief base eliminating irrelevant formulas as suggested in [4] and [16]: we only revise the relevant component and keep the rest of the belief base unchanged.

We now present a slightly different strategy which uses \models_S^1 in a more sophisticated way. Recall from Definition 4.1 that $R(p, q, B)$ iff there is a clause in B which has literals based on both p, q . Now define R_i as the set of atoms which are i relevant to atoms in $L(\alpha)$ according to Definition 4.1. Then for any k we choose, divide all atoms into three classes: $V = R_k$, $T = R_{k+1} \setminus R_k$, and $U = L - R_{k+1}$. V are the atoms most relevant to α , T are the atoms of medium relevance and U contains those which, we hope, are not relevant to the truth of α . Now eliminate all literals based on atoms from T . The remaining clauses either contain literals only from R_k or else literals not from R_{k+1} , and these two varieties never occur together in a clause. Let B_1 be the first (relatively small) set of clauses and B_2 the second. Assuming that B_2 is not inconsistent, we get, $B_1 \not\models \alpha$ iff $B \not\models_S^1 \alpha$ and of course if $B \not\models_S^1 \alpha$, then $B \not\models \alpha$. This allows us to approximate $B \models \alpha$ from the negative direction. We postpone details and proofs to the full paper.

7 Conclusion

In this paper, we have presented an application of the techniques of approximate entailment to belief revision. We have shown how approximate belief change fits the local change paradigm given in [8]. This provides us a clear logical axiomatization of the approximate operations. It also provides us an efficient method for implementing operations of belief change as performed by resource-bounded agents. In addition, we have provided heuristics for the construction of the context set S , without which any implementation of approximated reasoning systems would be very difficult. We note that domain specific knowledge will play a key role in any heuristic for constructing the context set. In future work, we plan to implement such a system and further investigate its properties.

References

- [1] Carlos Alchourron, Peter Gärdenfors and David Makinson. On the logic of theory change: Partial meet functions for contraction and revision. In *Journal of Symbolic Logic*, Volume 50, pages 510-530, 1985.
- [2] Marco Cadoli and Marco Schaerf. Approximate inference in default logic and circumscription. In *Fundamenta Informaticae*, Vol. 23, 1995, 123-143.

- [3] Marco Cadoli and Marco Schaerf. The Complexity of Entailment in Propositional Multivalued Logics, *Annals of Mathematics and Artificial Intelligence*, Vol. 18 (1), 1996, 29-50.
- [4] Samir Chopra and Rohit Parikh. An inconsistency tolerant model for belief representation and belief revision. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence-IJCAI 99*, Morgan Kaufmann.
- [5] Samir Chopra, Konstantinos Georgatos, and Rohit Parikh. Non-monotonic inference on belief sequences (preliminary report). In Phokion Kolaitis and George Koletsos, editors, *Proceedings of the Second Pan-Hellenic Symposium on Logic*, 1999.
- [6] Kevin Gary. RABIT – A spreading activation approach to real-time commonsense reasoning. Master's thesis, Arizona State University, Tempe, Arizona, 1993.
- [7] Sven Ove Hansson. Kernel contraction. *Journal of Symbolic Logic*, 59:845–859, 1994.
- [8] Sven Ove Hansson and Renata Wassermann. Local change: A preliminary report. In *Fourth Symposium on Logical Formalizations of Commonsense Reasoning, London*, 1998.
- [9] Andreas Herzig. How to change factual beliefs using laws and independence information. In Dov M. Gabbay, Rudolf Kruse, Andreas Nonnengart, and Hans Juergen Ohlbach, editors, *Qualitative and Quantitative Practical Reasoning - First Int. Joint Conf. on Qual. and Quant. Practical Reasoning; ECSQARU-FAPR'97*, number 1244 in LNCS, pages 311–321. Springer, jun 1997.
- [10] Robert Kowalski. *Logic for Problem Solving*. North Holland, 1979.
- [11] Hector Levesque A logic of explicit and implicit belief. In *Proceedings of AAAI-84*.
- [12] Rohit Parikh. Beliefs, belief revision, and splitting languages (presented at ITALLIC-96) In *Logic, Language, and Computation, Volume 2*, eds. Lawrence S. Moss, Jonathan Ginzburg, and Maarten de Rijke, CSLI Lecture Notes No. 96, CSLI Publications, 1999, pp. 266–268.
- [13] Marco Schaerf and Marco Cadoli. Tractable reasoning via approximation. *Artificial Intelligence*, 74(2):249–310, 1995.
- [14] Annette ten Teije and Frank van Harmelen. Computing approximate diagnoses by using approximate entailment. In *Proceedings of the Fifth International Conference on Principles of Knowledge Representation and Reasoning (KR'96)*, pages 256–265, Boston, Massachusetts, November 1996.
- [15] Annette ten Teije and Frank van Harmelen. Exploiting domain knowledge for approximate diagnosis. In M. Pollack, editor, *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence (IJCAI'97)*, pages 454–459, Nagoya, Japan, August 1997.
- [16] Renata Wassermann. On structured belief bases. In Hans Rott and Mary-Anne Williams, editors, *Frontiers in Belief Revision*. Kluwer, 2000. to appear.
- [17] Renata Wassermann. Local Diagnosis. In Proceedings of the 8th International Workshop on Non-Monotonic Reasoning, Breckenridge, Colorado, 2000.

Received