An Inconsistency-Tolerant Approach to Information Merging Based on Proposition Relaxation

Steven Schockaert*
Ghent University
Dept. of Applied Mathematics and Computer Science
Krijgslaan 281, 9000 Gent, Belgium
steven.schockaert@ugent.be

Henri Prade
Toulouse University, Université Paul Sabatier
IRIT, CNRS,118 route de Narbonne,
31062 Toulouse Cedex 09, France
prade@irit.fr

Abstract

Inconsistencies between different information sources may arise because of statements that are inaccurate, albeit not completely false. In such scenarios, the most natural way to restore consistency is often to interpret assertions in a more flexible way, i.e. to enlarge (or relax) their meaning. As this process inherently requires extra-logical information about the meaning of atoms, extensions of classical merging operators are needed. In this paper, we introduce syntactic merging operators, based on possibilistic logic, which employ background knowledge about the similarity of atomic propositions to appropriately relax propositional statements.

Introduction

Inconsistency is often encountered, especially in case of multiple-source information. Common to most approaches that have been proposed to cope with inconsistency is the idea that some pieces of information are wrong, and thus responsible for the inconsistency. In this view, consistency should be recovered by removing incorrect pieces of information. This may be done, for instance by dealing with maximal consistent subsets, or by ignoring less entrenched propositions whose departure reestablishes consistency. This intuition is at work in most of the approaches to reasoning under inconsistency (Besnard and Hunter 1998), or belief revision (Gärdenfors 1988).

Obviously it might be the case that certain information is just wrong. However, there are inconsistency situations where, in fact, all the pieces of information are (at least) approximately right. This means that if we interpret each statement in a sufficiently broad sense, inconsistency vanishes. Uncertainty in meaning is then the source of inconsistency, rather than the presence of information that would be "really" false.

Example 1 *Consider a source K_1 claiming that John lives in Gainesville FL, while another source K_2 claims that John lives in the city of Atlanta GA. While these sources are clearly in conflict, consistency can be restored by interpreting their assertions more liberally. In particular, it is plausible that the first source mistakingly assumed that John lives in Gainesville FL, whereas in fact John lives in Gainesville GA, Gainesville TX or Gainesville VA, e.g. because of a human error or because of a default assumption in K_1 that* Gainesville *means* Gainesville FL. *Similarly, the claim that John lives in the city of Atlanta may be weakened to a claim that John lives in the Atlanta Metropolitan Area, or even that he lives in Atlanta's combined statistical area (as defined by the US Census Bureau). As Gainesville GA is located in Atlanta's combined statistical area, the most natural repair would be to assume that John lives in Gainesville GA.*

Being liberal in the interpretation of propositions may be justified by slightly different reasons. Indeed a source may provide statements that are somewhat too precise w.r.t. the actual state of available information (e.g. some city called Gainesville vs. Gainesville FL). In addition, sources are often heterogeneous in their use of categories for describing reality, and they may make slightly different uses of the same label (e.g. city of Atlanta vs. Atlanta Metropolitan Area). Lastly, information may also evolve with time or space. In practice, this requires to have some background knowledge about similarities between interpretations, or between atoms of the language. This point of view is reminiscent of a proposal in (Wahlster 1980), where the author allows for vagueness in dialogues, and acknowledges the idea that two agents may use the same word with a slightly different meaning (e.g. "a room being quite large" does not, perhaps, refer exactly to the same space for a customer and for a hotel manager).

The idea of introducing a similarity relation between interpretations has been little studied, even if it received a beginning of attention in the early work of (Ruspini 1991), where approximate entailment is defined as "p approximately entails q" iff every model of p is similar to a model of q. This idea can be contrasted with non-monotonic reasoning, where "if p, generally q" iff the most plausible models of p are also models of q (Kraus, Lehmann, and Magidor 1990), i.e. in the former case the set of models of q is expanded, whereas in the latter case the set of models of p is restricted. The use of a similarity relation between interpretations, which was also discussed from a belief revision point of view in a short note by (Rodriguez, Garcia, and Godo 1995), has recently been reconsidered, in a preliminary manner, by (Schockaert and Prade 2009) for merging conflicting information.

*Postdoctoral fellow of the Research Foundation – Flanders.

This paper further explores the use of similarity information for merging conflicting sources. Our main aim is to show how practical merging operators can be defined in a syntactic way using the setting of possibilistic logic. The paper is organized as follows. In the next section some important concepts from propositional merging and possibility theory are recalled. Subsequently, a qualitative representation of similarity between atoms is proposed, after which we study merging operators based on weakening of propositions for knowledge bases in conjunctive-normal form. We then briefly describe a dual approach for propositions in disjunctive-normal form. Finally, an overview of related works is presented as well as our conclusions.

Background

We first review some standard approaches to merging propositional knowledge bases, before providing a short background on possibilistic logic.

Propositional merging

In the following, we consider a propositional language built from a finite set of atoms $\mathcal{A}$ and the connectives $\vee, \wedge, \rightarrow, \neg$ in the usual way. An interpretation I is defined as a subset of $\mathcal{A}$, where $I \models a$ for an atom a iff $a \in I$. An interpretation is said to be a model of a formula (resp. set of formulas) if it satisfies that formula (resp. every formula in the set) in the usual sense. We write $[\![\phi]\!]$ to denote the set of all models of a formula ϕ. The set of all interpretations will be written as $\mathcal{W}$.

Let $K_1, ..., K_n$ be propositional knowledge bases that are individually consistent, and let C be a set of integrity constraints. The purpose of a merging process is to find one knowledge base $\mathcal{K}$ which is consistent with the integrity constraints (i.e. $[\![\mathcal{K}]\!] \subseteq [\![C]\!]$), and which integrates the information from the knowledge bases $K_1, ..., K_n$. When these knowledge bases are in conflict with each other or with the integrity constraints, i.e. $[\![K_1]\!] \cap ... \cap [\![K_n]\!] \cap [\![C]\!] = \emptyset$, some of the assertions in the knowledge bases need to be ignored or weakened to this end.

A common strategy is to define a pseudo-metric d on $\mathcal{W}$, and to define $\mathcal{K}$ semantically by

$$[\![\mathcal{K}]\!] = Min(\leq_{(K_1,...,K_n)}, [\![C]\!]) \qquad (1)$$

where $Min(\leq, X)$ denotes the set of elements from X that are minimal w.r.t. the preorder $\leq$, and $I_1 \leq_{(K_1,...,K_n)} I_2$ for $I_1, I_2 \in \mathcal{W}$ iff

$$f(\min_{J\in K_1} d(I_1, J), ..., \min_{J\in K_n} d(I_1, J)) \leq f(\min_{J\in K_1} d(I_2, J), ..., \min_{J\in K_n} d(I_2, J)) \qquad (2)$$

with f an appropriate aggregation function, such as (refinements of) the maximum or a (weighted) sum. This approach to propositional merging is called the distance-based framework, and is usually implemented using the Hamming distance $d_{Ham}(I_1, I_2) = |(I_1 \setminus I_2) \cup (I_2 \setminus I_1)|$ (Dalal 1988; Revesz 1993; Lin 1996; Konieczny and Pino Pérez 2002).

Another interesting view on propositional merging, proposed in (Bloch and Lang 2002), is to weaken propositions using a logical counterpart of mathematical morphology (Serra 1982). In particular, the dilation $D_B(\phi)$ of a formula ϕ is defined by

$$[\![D_B(\phi)]\!] = \{I \in \mathcal{W} | B(I) \cap [\![\phi]\!] \neq \emptyset\}$$

where $B(I) \subseteq \mathcal{W}$ is a set of interpretations that are related to I in some way. Note that B, which is called the structuring element, can be regarded as a relation on $\mathcal{W}$. Again the Hamming distance is most commonly used, choosing $B(I) = \{I' \in \mathcal{W} | d_{Ham}(I, I') \leq 1\}$. This particular choice makes it straightforward to implement this operator (see (Bloch and Lang 2002)). Unfortunately, it also reduces dilation to a purely syntactical operation. An analogous remark applies to the distance-based framework, noting that distances and similarities are essentially equivalent. Again, using the Hamming distance has important computational and practical advantages, but does not allow to incorporate available extra-logical information.

The two aforementioned approaches are very general, and, from a formal point of view, our proposal can be related to this setting, as we show below. Still our approach has a particular ingredient in the form of extra-logical knowledge, which provides it with a quite different flavor.

Possibilistic logic

Given a possibility distribution[1] π on the set of all possible interpretations $\mathcal{W}$, the possibility $\Pi(\phi)$, necessity $N(\phi)$ and guaranteed possibility $\Delta(\phi)$ of a formula ϕ are defined as

$$\Pi(\phi) = \max_{I\in[\![\phi]\!]} \pi(I)$$
$$N(\phi) = 1 - \Pi(\neg\phi) = \min_{I\in[\![\neg\phi]\!]} 1 - \pi(I)$$
$$\Delta(\phi) = \min_{I\in[\![\phi]\!]} \pi(I)$$

A formula in possibilistic logic (Dubois and Prade 2004) is a pair (ϕ, λ) consisting of a classical logic formula ϕ and a weight $\lambda \in (0, 1]$ expressing either a lower bound on the necessity $N(\phi)$, or on the guaranteed possibility $\Delta(\phi)$ of ϕ. The pair (ϕ, λ) is called an N-formula in the first case and a Δ-formula in the second case. To avoid confusion, we will write $(\!|\phi, \lambda|\!)$ in the case of Δ-formulas. Interpreted as an N-formula, (ϕ, λ) means that ϕ should be satisfied with priority (or certainty) λ. Thus, if ϕ is not satisfied by some interpretation I, the desirability (or possibility) of I is upper bounded by $1 - \lambda$. Similarly, the Δ-formula $(\!|\phi, \lambda|\!)$ means that the desirability (or possibility) of interpretations satisfying ϕ is at least λ. A set of N-formulas is called an N-possibilistic logic base (or N-base), and a set of Δ-formulas a Δ-possibilistic logic base (or Δ-base). Because of the decomposability properties $N(\phi \wedge \psi) = \min(N(\phi), N(\psi))$ and $\Delta(\phi \vee \psi) = \min(\Delta(\phi), \Delta(\psi))$, it is always possible to

[1] Recall that a possibility distribution in a universe X is an $X - [0, 1]$ mapping, used as a convenient way to encode a complete ordering that models plausibility or preference.

represent N- and Δ-bases as weighted sets of clauses, and weighted sets of terms[2] respectively.

An important feature of possibilistic logic is its ability to cope with inconsistency. Given an N-base B, the λ-cut B_λ and strict λ-cut $B_{\underline{\lambda}}$ are defined as

$$B_\lambda = \{\phi_i | (\phi_i, \lambda_i) \in B, \lambda_i \geq \lambda\}$$
$$B_{\underline{\lambda}} = \{\phi_i | (\phi_i, \lambda_i) \in B, \lambda_i > \lambda\}$$

Then, the inconsistency level of B is defined as

$$inc(B) = \max\{\lambda | B_\lambda \text{ is inconsistent}\}$$

where $\max \emptyset = 0$ is assumed. Clearly, formulas in $B_{\underline{inc(B)}}$ are safe from inconsistency. For Δ-bases, we can define the dual notion of the λ-cut of N-bases

$$B^\lambda = \bigvee \{\phi_i | (\!|\phi_i, \lambda_i|\!) \in B, \lambda_i \geq \lambda\}$$

The satisfiability level of B is then defined as

$$sat(B) = \max\{\lambda | B^\lambda \text{ is consistent}\}$$

Representing Similarity

Although conceptually any type of similarity relation between interpretations could be useful, in practice we will focus on similarity relations that are induced by similarities between atoms. This makes it feasible to derive interesting syntactic operators, and has computational advantages, as explicitly enumerating similarity relations in the universe of interpretations $\mathcal{W} = 2^{\mathcal{A}}$ would not be tractable.

Because in most application scenarios the exact strength of such similarities cannot be quantified, we focus on a qualitative notion of similarity. More in particular, we assume that for each atom a, a collection of $k+1$ sets of atoms $X_a^0, X_a^1, ..., X_a^k$ is available such that $\{a\} = X_a^0 \subseteq X_a^1 \subseteq ... \subseteq X_a^k$. Intuitively, the sets X_a^i contain those atoms that are considered similar to a, adopting an increasingly more tolerant notion of similarity. Note that we do not require that X_a^k is equal to the set of all atoms $\mathcal{A}$, as there can be many pairs of atoms that cannot be considered similar, even when a very weak notion of similarity is used. Similarity information can be encoded explicitly by a domain expert. In many applications, however, such information can be obtained automatically from readily available resources such as taxonomies or, in the case of Example 1, geographic gazetteers.

In the following, we tacitly assume that similarity is symmetric in the sense that $b \in X_a^i$ iff $a \in X_b^i$, but we make no further assumptions. In particular note that similarity is not required to be transitive ($b \in X_a^u$ and $c \in X_b^v$ does not require that $c \in X_a^w$ for any w).

Given this encoding of similarity between atoms, we can define the (l–)expansion $\langle A \rangle^l$ and (l–)contraction $[A]^l$ of a set of atoms A, as

$$\langle A \rangle^l = \bigcup_{a \in A} X_a^l \qquad [A]^l = co\langle coA \rangle^l$$

where the complement co is understood as set complement w.r.t. $\mathcal{A}$. The set $\langle A \rangle^l$ contains all atoms that are similar to some atom in A, while $[A]^l$ contains those atoms that are only similar to atoms in A.

Relaxing propositions in CNF

If the knowledge bases $K_1, ..., K_n$ are in conjunctive-normal form (CNF), their clauses can be relaxed by adding additional disjuncts. Specifically, to weaken a positive literal a (i.e. an atom) we can replace it by the disjunction of all atoms in X_a^l for some $l \in \{1, ..., k\}$, i.e. a is considered "almost true" if an atom similar to a is true. To weaken a negative literal $\neg a$, we can replace it by the negation of the conjunction of atoms in X_a^l, i.e. $\neg a$ is considered "almost true" if an atom similar to a is false. For the ease of presentation, we write $a^{(l)}$ to denote $\bigvee X_a^l$ and $a_{(l)}$ to denote $\bigwedge X_a^l$. We also write $(\neg a)^{(l)}$ for $\neg(a_{(l)})$.

We can extend the notation $.^{(l)}$ to clauses by defining

$$(\alpha_1 \vee ... \vee \alpha_s)^{(l)} = \alpha_1^{(l)} \vee ... \vee \alpha_s^{(l)}$$

and to sets K of clauses by $K^{(l)} = \{\alpha^{(l)} | \alpha \in K\}$. Notice that the result $K^{(l)}$ is not independent of the syntactic encoding of K.

Example 2 *Let* $K_1 = \{a \vee b, \neg b\}$, $K_2 = \{a, \neg b\}$, $X_a^1 = \{a, a'\}$ *and* $X_b^1 = \{b, b'\}$. *Then clearly* K_1 *and* K_2 *are semantically equivalent, while*

$$K_1^{(1)} = \{a \vee a' \vee b \vee b', \neg b \vee \neg b'\}$$
$$K_2^{(1)} = \{a \vee a', \neg b \vee \neg b'\}$$

which are not semantically equivalent.

If desired, syntax-independence can always be obtained by replacing a propositional knowledge base K by its prime implicates, before applying the weakening operator $.^{(l)}$. Recall that a clause γ is an implicate of K iff $K \models \gamma$, and a prime implicate if furthermore for every other implicate γ', if $\gamma' \models \gamma$ then also $\gamma \models \gamma'$. Clearly, if γ is a prime implicant, clauses that are equivalent to γ are prime implicates too. Therefore, let $\Gamma(K)$ be the (finite) set of prime implicates of K in which no literals are repeated (e.g. clauses such as $a \vee a \vee b$ are excluded). This idea of enforcing syntax-independence was proposed in (Bienvenu, Herzig, and Qi 2008), in the context of belief revision.

Given a propositional knowledge base K, we define the associated N-possibilistic logic base K^N as

$$K^N = \{(\alpha^{(l)}, \lambda_l) | \alpha \in \Gamma(K) \text{ and } l \in \{0, ..., k\}\} \qquad (3)$$

Where $1 = \lambda_k > \lambda_{k-1} > ... > \lambda_0 > 0$, i.e. the unit interval is used to attach certainty values to formulas in a purely ordinal way, such that the more we weaken a formula, the higher its priority or certainty becomes. In other words, although we do not assume that the formulas in K are completely accurate, we assume that we can make them accurate (albeit less informative) by sufficiently weakening them. As the integrity constraints C should not be weakened, the associated N-possibilistic logic base C^{N^*} is given by

$$C^{N^*} = \{(\alpha, \lambda_k) | \alpha \in \Gamma(C)\} \qquad (4)$$

[2]We use *term* to denote the conjunctive counterpart of a clause, i.e. a conjunction of literals.

Example 3 *Consider again the scenario from Example 1, where we write aCity(x), aMA(x) and aCSA(x) to denote that person x lives in Atlanta city, metropolitan area, and combined statistical area respectively. We write nyCity(x), nyMA(x) and nyCSA(x) to denote that x lives in New York city, metropolitan area and combined statistical area, and gFL(x), gGA(x), gTX(x) and gVA(x) to denote that x lives in Gainesville FL, GA, TX, and VA respectively. Assume that for all x*

$$X^1_{aCity(x)} = \{aCity(x), aMA(x)\}$$
$$X^2_{aCity(x))} = X^1_{aCity(x)} \cup \{aCSA(x)\}$$
$$X^1_{gFL(x)} = X^2_{gFL(x)} = \{gFL(x), gGA(x), gTX(x), gVA(x)\}$$

Now consider the following two knowledge bases:

$$K_1 = \{aCity(john) \vee aCity(mary)\}$$
$$K_2 = \{gFL(john), nyCity(mary)\}$$

and assume that the integrity constraints C encode available geographic knowledge, e.g. $gGA(x) \rightarrow aCSA(x)$ and $\neg(aCity(x) \wedge nyCity(x))$. We obtain:

$$K_1^N = \{(aCity(john) \vee aCity(mary), \lambda_0), (aCity(john) \vee aMA(john) \vee aCity(mary) \vee aMA(mary), \lambda_1), (aCity(john) \vee aMA(john) \vee aCSA(john) \vee aCity(mary) \vee aMA(mary) \vee aCSA(mary), \lambda_2)\}$$

$$K_2^N = \{(gFL(john), \lambda_0), (nyCity(mary), \lambda_0), (gFL(john) \vee gGA(john) \vee gTX(john) \vee gFL(john), \lambda_1), (nyCity(mary) \vee nyMA(mary), \lambda_1), (gFL(john) \vee gGA(john) \vee gTX(john) \vee gFL(john), \lambda_2), (nyCity(mary) \vee nyMA(mary) \vee nyCSA(mary), \lambda_2)\}$$

Note that $(K_1^N \cup K_2^N)_{\lambda_0}$ is logically equivalent to $K_1 \cup K_2$, hence it is sufficiently informative to cover all we know, but it will not be consistent with the integrity constraints. On the other hand, $(K_1^N \cup K_2^N)_{\lambda_2}$ is logically much weaker than what is expressed by the sources, but it is in agreement with the integrity constraints.

Restoring consistency using λ-cuts

The possibilistic representation of the initial knowledge bases, defined by (3), enables us to adopt standard techniques for merging inconsistent possibilistic knowledge bases (e.g. (Benferhat and Sossai 2006; Hunter and Liu 2008)). In particular, we may define the result of the merging process in terms of λ-cuts of the possibilistic knowledge bases. The following proposition provides a semantic characterization of such merging operators.

Proposition 1 *Let K be a propositional knowledge base. For an interpretation I, the following two claims are equivalent ($0 \leq l \leq k$):*

1. $I \in [\![(K^N)_{\lambda_l}]\!]$.
2. *There exists a $J \in [\![K]\!]$ such that $[I]^l \subseteq J \subseteq \langle I \rangle^l$.*

The condition $[I]^l \subseteq J \subseteq \langle I \rangle^l$ may be interpreted as J being similar to I. In this view, the propositional knowledge base $(K^N)_{\lambda_l}$ is such that its models are exactly those interpretations that are similar to some model of K. Thus, restoring consistency by taking appropriate λ-cuts of the possibilistic bases K_i^N is effectively in line with the idea of similarity-based merging. Note that the procedure of forcing syntax-independence using prime implicates is essential in this result.

To see the relationship between relaxing propositions, as defined above, and the morpho-logical approach from (Bloch and Lang 2002), let us define B^l, for $I \in \mathcal{W}$ and $l \in \{0, ..., k\}$, as

$$B^l(I) = \{J \in \mathcal{W} | [I]^l \subseteq J \subseteq \langle I \rangle^l\} \quad (5)$$

Note that the relation B^l can be interpreted as a similarity relation on the set of interpretations, where l acts as a tolerance parameter. Proposition 1 teaches us that for every set of clauses K, viewed as a conjunction:

$$\left(\bigwedge K^N_{\lambda_l}\right) \equiv \left(D_{B^l}(\bigwedge K)\right)$$

In other words, the weakening operator $.^{(l)}$ introduced above essentially corresponds to logical dilation D_{B^l}, in the specific case that the structuring element is defined as (5). The approach can also be linked to the distance-based framework, using $d = d_1$ in (2), with

$$d_1(I, J) = \min\{l \in \mathbb{N} | [I]^l \subseteq J \subseteq \langle I \rangle^l\} \quad (6)$$

if $[I]^k \subseteq J \subseteq \langle I \rangle^k$, and $d_1(I, J) = k + 1$ otherwise. It is easy to show that $d_1(I, J) = 0$ iff $I = J$. Note, however, that d_1 is not symmetric (which is not essential in the distance-based framework). In the particular case where $f = \max$ in (2), and where $(K_1^N \cup ... \cup K_n^N \cup C^{N^*})_{\lambda_k}$ is consistent (i.e. inconsistencies can actually be repaired by sufficiently weakening each of the knowledge bases), it holds that

$$Min(\leq_{(K_1,...,K_n)}, [\![C]\!]) = (K_1^N \cup ... \cup K_n^N \cup C^{N^*})_{\underline{\alpha}}$$

where $\alpha = inc(K_1^N \cup ... \cup K_n^N \cup C^{N^*})$. Note that when $(K_1^N \cup ... \cup K_n^N \cup C^{N^*})_{\lambda_k}$ is not consistent, we should conclude that inconsistencies were not caused by inaccuracies but by "real errors", and fall back on standard techniques, e.g. based on the Hamming distance.

Restoring consistency using preferred subtheories

Selecting from an N-base B those formulas whose certainty degree is above $inc(B)$ is a standard approach to deal with conflicts. However, it is also rather coarse because it does not discriminate between interpretations that satisfy most of the remaining formulas (i.e. those with certainty at most $inc(B)$) and interpretations that satisfy almost none of the remaining formulas. A similar observation was made in (Dubois and Prade 1997) for belief revision using possibilistic logic, where it was proposed, as a refinement, to only consider models of preferred subtheories of B in the sense of (Brewka 1989). Specifically, $B^* = B_k \cup ... \cup B_0$ is a preferred subtheory of an N-base $B = K^N$ iff $B_k \cup ... \cup B_l$ is a maximal consistent subset of B_{λ_l} for all $l \in \{0, ..., k\}$.

This boils down to selecting as many formulas with certainty λ_k as possible (without getting inconsistency), subsequently adding as many formulas with certainty λ_{k-1} as possible, etc. Hence when I is a model of a preferred subtheory of B then also $I \in [\![B_{\underline{inc(B)}}]\!]$, but the converse does not hold in general.

At the semantic level, we can refine the preference relation on interpretations induced by (6). Specifically, if neither $[I_1]^l \subseteq J \subseteq \langle I_1 \rangle^l$ nor $[I_2]^l \subseteq J \subseteq \langle I_2 \rangle^l$ for any $J \in [\![K_i]\!]$, we may still prefer I_1 over I_2 if some model $J \in [\![K_i]\!]$ contains most of the atoms in $[I_1]^l$ and only few atoms outside $\langle I_2 \rangle^l$. Formally, this intuition leads to the following preorder $\leq_i^l$ on interpretations: $I \leq_i^l I'$ iff

$$\forall J' \in [\![K_i]\!] \,.\, \exists J \in [\![K_i]\!] \,.\, [I]^l \setminus J \subseteq [I']^l \setminus J' \text{ and } J \setminus \langle I \rangle^l \subseteq J' \setminus \langle I' \rangle^l$$

Furthermore, let $I \leq^l I' \equiv \forall i \in \{1, ..., n\} \,.\, I \leq_i^l I'$ and $I <^l I' \equiv l \leq^l I' \wedge \neg(I' \leq I)$. The following proposition allows us to relate this semantic refinement to maximal consistent subsets, and thus to preferred subtheories.

Proposition 2 *Let K_i be a propositional knowledge base, $l \in \{0, ..., k\}$ and let I and I' be two interpretations. Furthermore, let $F = \{\alpha | (\alpha, \lambda_l) \in K_i^N$ and $I \models \alpha\}$ and $F' = \{\alpha | (\alpha, \lambda_l) \in K_i^N$ and $I' \models \alpha\}$. It holds that $F \supseteq F'$ iff $I \leq_i^l I'$.*

Let $lex(\leq^k, ..., \leq^0)$ be the lexicographic extension of the preorders $\leq^l$, i.e. $(I, I') \in lex(\leq^k, ..., \leq^0)$ iff either $\forall l \,.\, I \leq^l I'$ or $\exists l_1 \,.\, I <^{l_1} I' \wedge \forall l_2 > l_1 \,.\, I \leq^{l_1} I'$. The following characterization, which follows easily from Proposition 2, allows us again to relate our similarity-based approach to the dilation operation and to the distance-based framework; we omit the details.

Corollary 1 *Assume that $(K_1^N \cup ... \cup K_n^N \cup C^{N^*})_{\lambda_k}$ is consistent. It holds that I is a model of a preferred subtheory of $K_1^N \cup ... \cup K_n^N \cup C^{N^*}$ iff $I \in Min(lex(\leq^k, ..., \leq^0), [\![C]\!])$.*

Example 4 *Consider again the knowledge bases K_1 and K_2 from Example 3. Let $\mathcal{K} = (K_1^N \cup K_2^N \cup C^{N^*})_{\lambda_2} = (K_1^N \cup K_2^N \cup C^{N^*})_{\underline{inc(K_1^N \cup K_2^N \cup C^{N^*})}}$. Then $\mathcal{K}$ is logically equivalent to $C \cup \{$gGA(john), nyCity(mary)$\vee$nyMA(mary)$\vee$ nyCSA(mary)$\}$ (noting among others that $gGA(john) \rightarrow aCSA(john)$). However, the most natural result would be $\mathcal{K}' = C \cup \{$gGA(john), nyCity(mary)$\}$ as there is no reason to doubt that Mary lives in New York city. If we define the result of the merging process as the union of the preferred subtheories of $K_1^N \cup K_2^N \cup C$, we indeed find a knowledge base which is equivalent to $\mathcal{K}'$, since the only preferred subtheory of $K_1^N \cup K_2^N \cup C^{N^*}$ is given by*

$\{$gFL(john) $\vee$ gGA(john) $\vee$ gTX(john) $\vee$ gFL(john),
aCity(john) $\vee$ aMA(john) $\vee$ aCSA(john)
$\vee$ aCity(mary) $\vee$ aMA(mary) $\vee$ aCSA(mary),
nyCity(mary), nyCity(mary) $\vee$ nyMA(mary),
nyCity(mary) $\vee$ nyMA(mary) $\vee$ nyCSA(mary)$\} \cup C$

Relaxing propositions in DNF

A dual approach to relaxing propositions can be developed for propositions in disjunctive-normal form (DNF). Due to reasons of space, and the analogy with the CNF-case, we only present the main ideas. A conjunction of literals can be weakened by removing all literals which are similar to literals that do not appear in that conjunction. Formally, if A and B are sets of atoms and $\gamma = \bigwedge_{a \in A} a \wedge \bigwedge_{b \in B} \neg b$, we define $\gamma^{\textcircled{l}}$ as

$$\gamma^{\textcircled{l}} = \bigwedge_{a \in [A]^l} a \wedge \bigwedge_{b \in [B]^l} \neg b$$

Furthermore, we define $(\gamma_1 \vee ... \vee \gamma_s)^{\textcircled{l}}$ as $\gamma_1^{\textcircled{l}} \vee ... \vee \gamma_s^{\textcircled{l}}$. The intuition is that literals which are only similar to literals that appear in the conjunction cannot correspond to borderline cases.

Example 5 *Assume that we have information about different snapshots of a given knowledge base, and let atom $P(a_i)$ denote that object a had property P at time instant i. Moreover, let $X^l_{P(a_i)} = \{P(a_j) | abs(j - i) \leq l\}$. Intuitively, this means that when $P(a_i)$ is asserted, it is plausible that $P(a_{i-1})$ or $P(a_{i+1})$ also holds, even if $\neg P(a_{i-1})$ and $\neg P(a_{i+1})$ are believed, resp. because the knowledge base was not up-to-date at time instant $i - 1$, or because information was added to the knowledge base at time instant $i + 1$, which was not yet valid. In this view, when we know $\gamma = P(a_1) \wedge P(a_2) \wedge P(a_3)$, the most certain conjunct is $\gamma^{\textcircled{1}} = P(a_2)$.*

Let γ' be a CNF-formula which is equivalent to the DNF-formula γ, then in general γ'^l is not equivalent to $\gamma^{\textcircled{l}}$. Indeed, the syntactic weakening operators $.^l$ and $.^{\textcircled{l}}$ are based on a different, albeit related intuition: while γ'^l is obtained by stretching the meaning of what is asserted by γ', $\gamma^{\textcircled{l}}$ is obtained by removing all pieces of information that may be falsified by stretching the meaning of what is not asserted by γ (i.e. what may possibly be false).

We associate a Δ-possibilistic logic base K^Δ with K as follows:

$$K^\Delta = \{(\!|\alpha^{\textcircled{l}}, \mu_l|\!) | \alpha \in K \text{ and } l \in \{0, ..., k\}\}$$

where K is interpreted as a set of terms, and $\mu_0 = 1 > \mu_1 > ... > \mu_k$. Note that the less we weaken a formula, the better to have it satisfied. Counterparts of Proposition 1 and Proposition 2 can be derived, provided that care is taken to enforce syntax-independence; we omit the details.

Related Work

As already discussed above, there are formal links between our approach and the distance-based framework from (Konieczny and Pino Pérez 2002) and the morphological operators from (Bloch and Lang 2002). However, due to their generality, these existing works provide no special guidance on how to use knowledge about the similarity of atoms.

There is almost no work that looks at restoring consistency in logical settings by enlarging the set of models of propositions using an explicit notion of similarity, apart from the early note by (Rodriguez, Garcia, and Godo 1995) already mentioned in the introduction. In the setting of description logics, (Ovchinnikova, Wandmacher, and Kühnberger 2007) looks for overgeneralized concepts that create inconsistency problems, and tighten or enlarge such concepts in order to restore subconcept relations in a consistent way. In (Condotta, Kaci, and Schwind 2008), the idea of using similarities between atoms is used for the specific case of merging networks of qualitative temporal and spatial relations. If, e.g., a source claims that spatial regions a and b are disjoint while another asserts that in fact a is a part of b, the result of merging might be that a overlaps with b.

The present paper substantially advances the results provided in a recent short paper by the same authors (Schockaert and Prade 2009), where the central ideas were discussed at the semantic level. In particular, the development of the two dual views on syntactically encoding merging operators based on similarity, and the corresponding characterization results are new.

Conclusion

Our work originates from the observation that inconsistencies often arise because propositions are understood too strictly. In such a case, suitable merging operators can only be implemented by using appropriate extra-logical information about the meaning of atoms. We have focused specifically on using information about the similarity of atoms to this end, which turns out to be sufficiently expressive to allow interesting merging operators, while we can still realistically assume that such information is available in many applications (e.g. derived from taxonomies, geographic gazetteers, or domain experts). Using similarity information, we can associate a particular possibilistic knowledge base with each information source, in which increasingly weakened propositions receive an increasingly higher certainty (or preference). This reflects the intuition that we cannot always expect available information to be accurate, while we insist that the truth is never far from this information. This idea can be implemented in necessity and in guaranteed-possibility based possibilistic knowledge bases, which naturally corresponds to two dual forms of weakening. We have shown that well-known concepts in possibilistic knowledge bases, viz. λ-cuts and preferred subtheories, then play a central role in implementing merging operators. Still our proposal fully departs from existing merging approaches in the possibilistic setting.

References

Benferhat, S., and Sossai, C. 2006. Reasoning with multiple-source information in a possibilistic logic framework. *Information Fusion* 7(1):80–96.

Besnard, P., and Hunter, A. 1998. Reasoning with actual and potential contradictions. In *Vol. II of the Handbook of Defeasible Reasoning and Uncertainty Management Systems*. Kluwer Acad. Publ.

Bienvenu, M.; Herzig, A.; and Qi, G. 2008. Prime implicate-based belief revision operators. In *Proc. of the 18th European Conf. on Artificial Intelligence*, 741–742.

Bloch, I., and Lang, J. 2002. Towards mathematical morpho-logics. In *Technologies for constructing intelligent systems*, 380. Physica-Verlag GmbH.

Brewka, G. 1989. Preferred subtheories: An extended logical framework for default reasoning. In *Proc. of the 11th Int. Joint Conf. on Artificial Intelligence*, 1043–1048.

Condotta, J.-F.; Kaci, S.; and Schwind, N. 2008. A framework for merging qualitative constraint networks. In *Proc. of the 21st International FLAIRS Conf.*, 586–591.

Dalal, M. 1988. Investigations into a theory of knowledge base revision: preliminary report. In *Proceedings of the 7th National Conference on Artificial Intelligence*, 475–479.

Dubois, D., and Prade, H. 1997. A synthetic view of belief revision with uncertain inputs in the framework of possibility theory. *International Journal of Approximate Reasoning* 17(2-3):295 – 324.

Dubois, D., and Prade, H. 2004. Possibilistic logic: a retrospective and prospective view. *Fuzzy Sets and Systems* 144(1):3 – 23.

Gärdenfors, P. 1988. *Knowledge in Flux*. MIT Press.

Hunter, A., and Liu, W. 2008. A context-dependent algorithm for merging uncertain information in possibility theory. *IEEE Transactions on Systems, Man and Cybernetics, Part A* 38(6):1385–1397.

Konieczny, S., and Pino Pérez, R. 2002. Merging information under constraints: a logical framework. *Journal of Logic and Computation* 12(5):773–808.

Kraus, S.; Lehmann, D.; and Magidor, M. 1990. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence* 44(1-2):167–207.

Lin, J. 1996. Integration of weighted knowledge bases. *Artificial Intelligence* 83(2):363 – 378.

Ovchinnikova, E.; Wandmacher, T.; and Kühnberger, K.-U. 2007. Solving terminological inconsistency problems in ontology design. *Interoperability in Business Information Systems* 2(1):65–80.

Revesz, P. 1993. On the semantics of theory change: arbitration between old and new information. In *Proceedings of the 12th ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, 71–82.

Rodriguez, R.; Garcia, P.; and Godo, L. 1995. Relating similarity-based models, counterfactuals and theory change. In *Proc. EUFIT'95*, 230–234.

Ruspini, E. 1991. On the semantics of fuzzy logic. *International Journal of Approximate Reasoning* 5:45–88.

Schockaert, S., and Prade, H. 2009. Merging Conflicting Propositional Knowledge by Similarity. In *Proceedings of the 2009 21st IEEE International Conference on Tools with Artificial Intelligence*, 224–228.

Serra, J. 1982. *Image Analysis and Mathematical Morphology*. Academic Press.

Wahlster, W. 1980. Implementing fuzziness in dialogue systems. In Rieger, B., ed., *Empirical Semantics*. Brockmeyer, Bochum.