An argumentation framework for merging conflicting knowledge bases

Leila AMGOUD[1] Simon PARSONS[2]

[1] Department of Electronic Engineering, Queen Mary and Westfield College
University of London, Mile End Road, E1 4NS, UK
E-mail: l.amgoud@elec.qmw.ac.uk

[2] Department of Computer Science, University of Liverpool
Chadwick Building. Liverpool L69 7ZF
E-mail: s.d.parsons@csc.liv.ac.uk

Abstract.

The problem of merging multiple sources of information is central in many information processing areas such as databases integrating problems, multiple criteria decision making, etc. Recently several approaches have been proposed to merge classical propositional bases. These approaches are in general semantically defined. They use priorities, generally based on Dalal's distance for merging classical conflicting bases and return a new classical base as a result. In this paper, we present an argumentation framework for solving conflicts which could be applied to conflicts arising between agents in a multi-agent system. We suppose that each agent is represented by a consistent knowledge base and that the different agents are conflicting. We show that by selecting an appropriate preference relation between arguments, that framework can be used for merging conflicting bases and recovers the results of the different approaches proposed for merging bases [6], [11], [12], [13], [14], [15].

1. Introduction

In many areas such as cooperative information systems, multi-databases, multi-agents reasoning systems, GroupWare, distributed expert systems and so on, the information comes from multiple sources. In these areas, information from different sources is often contradictory. For example, in a distributed medical expert system, different experts often disagree on the diagnoses of patients' diseases. In a multi-database system two component databases may record the same data item but give it different values because of incomplete updates, system error, or valid differences in underlying semantics. Some researchers claim that, on an abstract level, the above problem can be subsumed under the general problem of merging multiple bases that may contradict each other. Several different approaches have been proposed for that purpose [6], [11], [12], [13], [14], [15]. Starting from different bases (Σ_1, …, Σ_n) which are conflicting, these works return a unique consistent base. However, in the case of multi-agent reasoning systems where each agent is supposed to have its own knowledge base, merging the bases looks debatable, since the goal of retaining all available information is quite legitimate in that case. As a result, other authors have considered reasoning with such bases without merging them. Argumentation is one of the most promising approaches developed for that purpose. Argumentation is based on the construction of arguments and counter-arguments (defeaters) and the selection of the most acceptable of these arguments. Inspired by the work presented in [1], we present a preference-based argumentation framework for reasoning with conflicting knowledge bases where each base could be part of a separate agent. This framework uses preference relations between arguments in order to determine the acceptable ones. We show that by selecting an appropriate preference relation between arguments, the preference-based argumentation framework can be used to merging conflicting bases in the sense that it recovers the results of fusion operators defined in [6], [11], [12], [13], [14], [15]. Thus the approach could be used by an agent, engaged in the kind of dialogue we have described in [2], as a means of handling conflicts between different agents' views of the world.

This paper is organized as follows: section 2 introduces the preference-based argumentation framework developed. In section 3 we show that the new framework recovers the results of other approaches to fusion namely those proposed in [12], [13], and in section 4 we briefly discuss how this work ties in with our work on multi-agent dialogues. Section 5 is devoted to some concluding remarks and perspectives.

2. Basic definitions

Let's consider a propositional language L over a finite alphabet P of atoms. Ω denotes the set of all the interpretations. Logical equivalence is denoted by $\equiv$ and classical disjunction and conjunction are respectively denoted by $\vee$, $\wedge$. Let φ be a formula of L, $[\varphi]$ denotes the set of all models of φ. A literal is an atom or a negation of an atom. Σ_i represents a classical propositional base. Let $E = \{\Sigma_1, \ldots, \Sigma_n\}$ $(n \geq 1)$ be a multi-set of n consistent propositional bases. We denote by Σ the set $\Sigma_1 \cup \ldots \cup \Sigma_n$ for short ($\Sigma = \Sigma_1 \cup \ldots \cup \Sigma_n$). Note that Σ may be inconsistent.

Definition 1. An *argumentation framework* (AF) is a triplet <A(Σ), Undercut, Pref>. A(Σ) is the set of all the arguments constructed from Σ, Undercut is a binary relation representing defeat relationship between arguments. Pref is a (partial or complete) preordering on $A(\Sigma)\times A(\Sigma)$. $>>^{Pref}$ denotes the strict ordering associated with Pref.

Several definitions of argument and the notion of defeat exist. For our purpose, we will use the definitions proposed in [10].

Definition 2. An *argument* of A(Σ) is a pair (H, h), where h is a formula of the language L and H a subbase of Σ satisfying: i) H is consistent, ii) H $\vdash$ h, iii) H is minimal (no strict subset of H satisfies i and ii). H is called the *support* and h the *conclusion* of the argument.

Definition 3. Let (H, h) and (H', h') be two arguments of A(Σ). (H, h) *undercuts* (H', h') iff for some $k \in H'$, $h \equiv \neg k$. An argument is undercut if there exists an argument against one element of its support.

Different preference relations between arguments may be defined. These preference relations are induced by a preference relation defined on the supports of arguments. The preference relation on the supports may be is itself defined from a (total or partial) preordering on the knowledge base. So, for two arguments (H, h), (H', h'), (H, h) is preferred to (H', h') (denoted by (H, h) $>>^{Pref}$ (H', h')) iff H is preferred to H' w.r.t Pref. In the next section, we will show some examples of the origin of a preordering on the base.

An example of a preference relation is the one based on the elitism principle (ELI-preference [7]). Let ≥ be a total preordering on a knowledge base K and > be the associated strict ordering. In that case, the knowledge base K is supposed to be stratified into (K_1, …, K_n) such that K_1 is the set of ≥-maximal elements in K and K_{i+1} the set of ≥-maximal elements in $K \setminus (K_1 \cup \ldots \cup K_i)$.

Let H and H' be two subbases of K. H is *preferred* to H' according to ELI-preference iff $\forall k \in H \setminus H', \exists k' \in H' \setminus H$ such that $k > k'$.

Let (H_1, h_1), (H_2, h_2) be two arguments of A. (H_1, h_1) $>>^{ELI}$ (H_2, h_2) iff H_1 is preferred to H_2 according to ELI-preference.

Example 1. *$K = K_1 \cup K_2 \cup K_3$ such that $K_1 = \{a, \neg a\}$, $K_2 = \{a \rightarrow b\}$ and $K_3 = \{\neg b\}$. ($\{a, a \rightarrow b\}$, b) $>>^{ELI}$ ($\{\neg b\}$, $\neg b$).*

Using the defeat and the preference relations between arguments, the *acceptable arguments* among the elements of A(Σ) may be defined. Inspired by the work of Dung [9], in [1] several definitions of the notion of acceptability have been proposed and the acceptable arguments are gathered in a so-called *extension*. For our purpose, we are interested by the extension satisfying a stability and the following coherence requirements:

Definition 4. A set $S \subseteq A(\Sigma)$ of arguments is *conflict-free* iff there doesn't exist a set $A, B \in A$ such that A undercuts B and not(B $>>^{Pref}$ A).

Definition 5. Let <A(Σ), Undercut, Pref> be an AF. A conflict-free set of arguments S is a *stable extension* iff S is a fixed point of a function G defined as:

$G : 2^A \times 2^A$

$S \rightarrow G(S) = \{A \in A(\Sigma) \mid$ there does not exist $B \in S$ such that B undercuts A and not $(A >>^{Pref} B)\}$

From the set of arguments $A(\Sigma)$, several stable extensions can be found: $\Pi = \{S_1, \ldots, S_n\}$. These extensions represent the different sets of *acceptable arguments*. They include all the arguments defending themselves against any undercutting argument and those defended. These notions of defense have been defined in [1] as follows:

An argument A is *defended* by a set S of arguments (or S *defends* A) iff $\forall B \in A$, if B undercuts A and not$(A >>^{Pref} B)$ then $\exists C \in S$ such that C undercuts B and not$(B >>^{Pref} C)$. If B undercuts A and $A >>^{Pref} B$ then we say that A *defends itself* against B. C_{Pref} denotes the set of all the arguments defending themselves against their defeaters. All the proofs of the results presented in this paper can be found in [3].

Property 1. Let <A(Σ), Undercut, Pref> be an AF. For any stable extension $S_I \in \Pi$, the following inclusion holds: $\underline{S} = C_{Pref} \cup [\cup F^{i\geq 1}(C_{Pref})] \subseteq S$. Let $T \subseteq A(\Sigma)$, $F(T) = \{A \in A(\Sigma) \mid$ A is defended by $T\}$.

This means that each stable extension contains the arguments which are not undercut, the arguments which can defend themselves against the undercutting arguments, and also the arguments defended by that extension.

Property 2. Let $\Sigma \neq \emptyset$ and <A(Σ), Undercut, Pref> be an AF. The set Π is not empty $(\Pi \neq \emptyset)$. This means that each argumentation framework has at least one stable extension.

Property 3. Let <A(Σ), Undercut, Pref> be an AF. $\underline{S} = \cap S_I$, i = 1, n. This means that each argument which is in one extension and not in another does not defend itself and it is not defended.

Acceptable arguments are defined in order to define the acceptable conclusions of an inconsistent knowledge base. So from the notion of acceptability, we define the following consequence relations.

Definition 5. Let <A(Σ), Undercut, Pref> be an AF. Let φ be a formula of the language L.

- φ is a *plausible consequence* of Σ iff there exists $H \subseteq \Sigma$ such that $(H, \varphi) \in A(\Sigma)$.
- φ is a *probable consequence* of Σ iff there exists a stable extension S_i and $\exists (H, \varphi) \in A(\Sigma)$ such that $(H, \varphi) \in S_i$.
- φ is a *certain consequence* of Σ iff there exists a stable extension S and $\exists (H, \varphi) \in A(\Sigma)$ such that $(H, \varphi) \in \underline{S}$.

The terms “plausible”, “probable” and “certain” are taken from [10]. Let’s denote by C_{pl}, C_{pr}, C_{ce} respectively, the set of all plausible consequences, all probable consequences and all certain consequences. The following inclusions hold.

Property 3. Let <A(Σ), Undercut, Pref> be an AF. $C_{ce} \subseteq C_{pr} \subseteq C_{pl}$.

Note that when the different bases are not conflicting, the relation Undercut is empty. Formally:

Property 4. Let Σ be a consistent base and <A(Σ), Undercut, Pref> be an AF.

- Undercut = ∅.
- There exists a unique stable extensions S = A(Σ) = the set of arguments which are not undercut.
- $C_{ce} = C_{pr} = C_{pl}$ = Th(Σ) (Th(Σ) denotes the deductive closure of Σ)

Example 2. *Let's consider three bases $\Sigma_1 = \{a\}$, $\Sigma_2 = \{a \to b\}$, $\Sigma_3 = \{\neg b\}$. We suppose that Σ_1 is more reliable than Σ_2 and Σ_2 is more reliable than Σ_3. Then $a > a \to b > \neg b$. In the framework <A(Σ), Undercut, ELI>, the set A(Σ) is {A = ({a}, a), B = ({a→b}, a→b), C = ({¬b}, ¬b), D = ({a, a→b}, b), E = ({¬b, a→b}, ¬a), F = ({a, ¬b},¬ (a→b))}. According to ELI-preference, A $>>^{ELI}$ E, D $>>^{ELI}$ C, E, F, B $>>^{ELI}$ F. A, B and C are preferred to their defeaters, so A, B, D ∈ S. In this example, there is a unique stable extension which is S.*

3. Connection with works on merging conflicting bases

Recently, several approaches have been proposed to merge classical propositional bases. These approaches can be divided into two categories; those approaches in which (explicit or implicit) priorities are used, and those in which priorities are not used. These approaches define a merging operator Λ which is a function associating to a set E = $\{\Sigma_1, \ldots, \Sigma_n\}$ a consistent classical propositional base, denoted by Λ(E).

Let B be a subset of A(Σ), Supp(B) is a function which returns the union of the supports of all the elements of B. Let T be subset of Σ, Arg(T) is a function which returns the arguments having their support in T.

Case 1. Non use of priorities.

There are two straightforward ways for defining Λ(E) depending on whether the bases are conflicting or not, namely:

- Classical conjunctive merging: $\Lambda(E) = \wedge\Sigma_i$ i = 1, n. In this case, Λ(E) = Σ.
- Classical disjunctive merging: $\Lambda(E) = \vee\Sigma_i$ i = 1, n. If we have two bases $\Sigma_1 = \{a\}$ and $\Sigma_2 = \{\neg a\}$, the result of merging is Λ(E) = {a ∨ ¬a} a tautology, which does supports neither a nor ¬a. Let's see what is the result provided by the argumentation framework. We consider then the base Σ = {a, ¬a}. A(Σ) = {({a}, a), ({¬a}, ¬a)}. Since no preferences are used in this approach, the preference relation between arguments Pref is

empty. Two stable extensions are computed: S_1 = {({a}, a)} and S_2 = {({¬a}, ¬a)}, consequently, a and ¬a are two probable consequences of the base Σ.

Other approaches, developed in [14], consist of computing the maximal (for set inclusion) consistent subsets of the union of the knowledge bases and then take as the result the intersection of the deductive closure of these subsets. The result of that merging operator is captured by our argumentation framework. Since no preferences exist, the relation Pref is empty (Pref = ∅). In this particular case, the set $\underline{S}$ contains the arguments which are not undercut. Let's denote by T = {T_1, …,T_n} the set of maximal consistent subsets of Σ. In that case, the following result has been proved in [5].

Proposition 1. Let <A(Σ), Undercut, Pref = ∅> be an AF. Π = {S_1, …, S_n} is the set of the corresponding stable extensions.

- ∀ S_i ∈ Π, Supp(S_i) ∈ T.
- ∀ T_i ∈ T, Arg(T_i) ∈ Π.
- ∀ S_I ∈ Π, Supp(S_I) is consistent.

As a direct consequence of this result we have:

Proposition 2. Let <A(Σ), Undercut, Pref> be an AF. Λ(E) = C_{ce}.

In other words, the result of merging the knowledge-bases is the set of certain consequences. The above results show that the approach developed in [14] is captured by our preference-based argumentation framework.

Example 3. *Let's extend the above example to three bases Σ_1 = {a}, Σ_2 = {¬a}, Σ_3 = {a}. There are two maximal (for set-inclusion) subsets of $\Sigma_1 \cup \Sigma_2 \cup \Sigma_3$, T_1 = {a} and T_2 = {¬a}. The result of merging these three bases is the empty set since the intersection between T_1 and T_2 is empty. In the argumentation framework, a and ¬a are two probable consequences, but there are no certain consequences.*

Case 2. Use of priorities.

Two kinds of approach which use priorities can be distinguished depending on whether they use implicit or explicit priorities. In [6] for example, the different bases are supposed weighted, the base having a higher weight is more reliable than the others. In presence of such weights, the maximal consistent subsets are computed by taking as many sentences as possible from the knowledge bases of higher weights (or priorities). In this case, the result Λ(E) is also captured by the argumentation framework presented in section 2, provided that we choose an appropriate preference relation between arguments and strengthen the definition of "conflict-free". The most appropriate relation is the one based on certainty level and defined in [4] in a possibilistic context. In that case, a knowledge base E is supposed to be stratified in $E_1, \dots, E_n$ such that the beliefs in E_i have the same certainty level and are more certain than the elements in E_j where $i < j$. (This notion of "certainty" corresponds to the degree of belief that an agent has in given propositions, and can be combined with a notion of its belief in other agents when that proposition is from another agent's knowledge base.) For our purpose, we suppose that $\Sigma = \Sigma'_1 \cup \dots \cup \Sigma'_k$ such that Σ'_i is the union of the bases having the same weight. Σ'_1 is the union of the bases having the highest weight and Σ'_k is the union of the bases having the smallest weight.

Definition 6. The weight of a non-empty subset H of Σ is the highest number of a layer (i.e. the lower layer) met by H, so $weight(H) = max \{j \mid 1 \leq j \leq k \text{ and } H_j \neq \emptyset\}$, where H_i denotes $H \cap \Sigma'_i$.

Definition 7. Let H, H' be two subsets of Σ. H is preferred to H' (denoted H $Pref_1$ H') iff weight(H) < weight(H'). Consequently, (H, h) is preferred to (H', h'), (denoted $(H, h) >>^{Pref1} (H', h')$), iff H $Pref_1$ H'.

The new definition of conflict-free is the following one:

Definition 8. A set $S \subseteq A(\Sigma)$ of arguments is *conflict-free* iff there doesnot exist a set A, B ∈ A such that A undercuts B.

Let's denote by $T = \{T_1, \dots, T_n\}$ the set of maximal consistent subsets of Σ. As in the case without preferences, these subsets can be computed from the stable extensions of the framework $<A(\Sigma), Undercut, Pref_1>$.

Proposition 3. Let $<A(\Sigma), Undercut, Pref_1>$ be an AF. $\Pi = \{S_1, \dots, S_n\}$ is the set of the corresponding stable extensions.

- $\forall S_i \in \Pi$, $Supp(S_i) \in T$.
- $\forall T_i \in T$, $Arg(T_i) \in \Pi$.
- $\forall S_I \in \Pi$, $Supp(S_I)$ is consistent.

Proposition 4. Let $<A(\Sigma), Undercut, Pref_1>$ be an AF. $\Lambda(E) = C_{ce}$.

Propositions 3 and 4 show that the approach developed in [6] is captured by our preference-based argumentation framework as well.

Some recent approaches are proposed for merging conflicting knowledge bases using implicit priorities. These priorities are extracted from the different interpretations. The three basic steps followed in [11], [12], [13], [15] for the semantic of a merging operator Λ are:

a) Rank-order the set of interpretations Ω w.r.t each propositional base Σ_i by computing a local distance, denoted $d(\omega, \Sigma_i)$, between ω and each Σ_i. The local distance is based on Dalal's distance [8]. The distance between an interpretation ω and a propositional base Σ_i

is the number of atoms on which this interpretation differs from some model of the propositional base. Formally, $d(\omega, \Sigma_i) = \min dist(\omega,\omega')$, $\omega' \in [\Sigma_i]$ where $dist(\omega,\omega')$ is the number of atoms whose valuations differ in the two interpretations.

b) Rank-order the set of interpretations Ω w.r.t all the propositional bases. This leads to the overall distance denoted $d(\omega, E)$. This later, computed from local distances $d(\omega, \Sigma_i)$, defines an ordering relation between the interpretations defined as follows: $\omega \leq \omega'$ iff $d(\omega, E) \leq d(\omega', E)$.

c) $\Lambda(E)$ is defined by being such that its models are minimal with respect to $\leq$, namely: $[\Lambda(E)] = \min(\Omega, \leq)$.

Example 3. (continued) *Let's consider the three following bases* $\Sigma_1 = \{a\}$, $\Sigma_2 = \{\neg a\}$, $\Sigma_3 = \{a\}$. *There exist two interpretations* $\omega_0 = \{a\}$ *and* $\omega_1 = \{\neg a\}$.

- $d(\omega_0, \Sigma_1) = d(\omega_0, \Sigma_3) = 0$, $d(\omega_0, \Sigma_2) = 1$.
- $d(\omega_1, \Sigma_1) = d(\omega_1, \Sigma_3) = 1$, $d(\omega_1, \Sigma_2) = 0$.

Once $d(\omega, \Sigma_i)$ is defined for each knowledge base Σi, several methods have been proposed in order to aggregate the local distances $d(\omega, \Sigma_i)$ according to whether the bases have the same weight or not. In particular the following operators have been proposed respectively in [12] and [13]: $d(\omega, E) = \sum_{i=1}^{n} d(\omega, \Sigma_i)$ and $d(\omega, E) = \sum_{i=1}^{n} d(\omega, \Sigma_i) \times \alpha_i$ where α_i is the weight associated with the base Σ_i. We denote by Λ_1 the first operator and by Λ_2 the second one.

Example 3. (continued) *Let's consider the three following bases* $\Sigma_1 = \{a\}$, $\Sigma_2 = \{\neg a\}$, $\Sigma_3 = \{a\}$. *According to the first operator* $\Lambda 1$, $d(\omega_0, E) = 1$ *and* $d(\omega_1, E) = 2$. *Then the generated base has* ω_0 *as an interpretation.*

Let's suppose now that Σ_2 *is more reliable than the two others and it has a weight 3.* Σ_1, Σ_3 *have weight 1. Using the operator* Λ_2, $d(\omega_0, E) = d(\omega_0, \Sigma_1) \times 1 + d(\omega_0, \Sigma_2) \times 3 + d(\omega_0, \Sigma_3) \times 1 = 3$ *and* $d(\omega_1, E) = 2$. *Then the generated base has* ω_1 *as an interpretation.*

To capture the results of these two merging operators, we consider the new definition of conflict-free (given in definition 8) and we define two new preference relations between arguments $Pref_2$ and $Pref_3$. These relations are based on Dalal's distance. The basic idea is to associate to the support of each argument a weight. This last corresponds to the minimal distance between the support and the different bases. The distance between a support H and a base Σ_i is computed as follows: $dist(H, \Sigma_i) = \min dist(\omega, \omega')$, $\omega \in [H]$ and $\omega \in [\Sigma_i]$.

Example 2. (continued) $\Sigma_1 = \{a\}$, $\Sigma_2 = \{a \rightarrow b\}$, $\Sigma_3 = \{\neg b\}$ *are three bases.* $H = \{a, a \rightarrow b\}$, $H' = \{\neg b\}$ *are two subsets of* Σ. $dist(H, \Sigma_1) = dist(H, \Sigma_2) = 0$, $dist(H, \Sigma_3) = 1$, $dist(H', \Sigma_1) = 0$, $dist(H', \Sigma_2) = 0$, $dist(H, \Sigma_3) = 0$.

To capture the results of the merging operator Λ_1, the weight of a support is defined as follows:

Definition 9. Let H be a subset of Σ. $\text{Weight}(H) = \sum_{i=1}^{n} \text{dist}(H, \Sigma_i)$.

Pursuing definition 7, a subset H of Σ is preferred to another subset H' (denoted H Pref_2 H') iff weight(H) < weight(H'). Consequently, (H, h) is preferred to (H', h'), (denoted (H, h) $>>^{\text{Pref2}}$ (H', h')), iff H Pref_2 H'.

Example 2. (continued) *H = {a, a→b}, H' = {¬b} are two subsets of Σ. Weight(H) = 1 and weight(H') = 0, then H' is preferred to H. Consequently, ({¬b},¬b) $>>^{Pref2}$ ({a, a→b}, b).*

Proposition 5. Let S_1, …, S_n be the stable extensions of the framework <A(Σ), Undercut, Pref_2>.

- $\underline{S}$ is not necessarily included in each S_i.
- $\forall S_i$, $\text{Supp}(S_i)$ is consistent.
- $[\text{Supp}(S_1)]$, …, $[\text{Supp}(S_n)]$ are the models obtained by the merging operator Λ_1.

Example 2. (continued) *Σ_1 = {a}, Σ_2 = {a→b}, Σ_3 = {¬b}. P = {a, b}, so the possible models are: ω_0 = {a, b}, ω_1 = {a, ¬b}, ω_2 = {¬a, b} and ω_3 = {¬a, ¬b}.*

d(ω_0, E) = d(ω_1, E) = d(ω_3, E) = 1 and d(ω_2, E) = 2, then the result of merging is the three models ω_0, ω_1, ω_3. Let's consider now the framework <A(Σ), Undercut, $Pref_2$> where A(Σ) = {A = ({a}, a), B = ({a→b}, a→b), C = ({¬b}, ¬b), D = ({a, a→b}, b), E = ({¬b, a→b}, ¬a), F = ({a, ¬b},¬ (a→b))}. weight({a}) = weight({a→b}) = weight({¬b}) = 0 and weight({a, a→b}) = weight({¬b, a→b}) = weight({a, ¬b}) = 1. The conflicts (in the sense of the relation "undercut") are represented by the figure below:

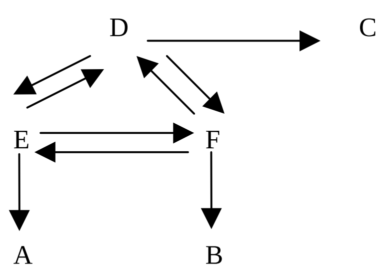

Three stable extensions can be computed: S_1 = {B, C, E}, S_2 = {A, B, D}, S_3 = {A, C, F}. [Supp(S_1)] = [{¬b, a→b}] = {¬a, ¬b} = ω_3, [Supp(S_2)] = [{a, a→b}] = {a, b} = ω_0, [Supp(S_3)] = [{a, ¬b}] = {a, ¬b} = ω_1.

To capture the results of the merging operator Λ_2, we suppose that each base Σ_i is equipped with a weight (priority) α_i. The definition of a support weight is:

Definition 10. Let H be a subset of Σ. $\text{Weight}(H) = \sum_{i=1}^{n} \text{dist}(H, \Sigma_i) \times \alpha_i$.

This new definition of weight leads to a new preference relation denoted by $Pref_3$.

Proposition 6. Let S_1, …, S_n be the stable extensions of the framework <A(Σ), Undercut, $Pref_3$>.

- ∀ S_i, Supp(S_i) is consistent.
- [Supp(S_1)], …, [Supp(S_n)] are the models obtained by the merging operator Λ_2.

Propositions 5 and 6 show that the approach developed in [11], [12], [13], [15] is captured by our preference-based argumentation framework too.

Overall, then, we can conclude that the approach outlined in this paper can capture a wide range of different approaches to merging information from inconsistent knowledge-bases.

4. Conflicts in multi-agent dialogues

Previous work on merging conflicting knowledge bases has been largely from the perspective of knowledge fusion. The idea is that there are a number of different repositories of information about the world which need to be merged in order to discover the true picture. From this perspective, it makes perfect sense to think of taking the separate knowledge bases and building one large coherent knowledge base from them. We, however, have a rather different view. Our work concentrates on dialogue in multi-agent systems. Individual agents have access to private knowledge bases that they use as a basis for constructing arguments justifying requests that they make to one another [2]. Here, too, we have conflict, when agents disagree about the truth of propositions, but it is just not practical for the agents to pool their knowledge in order to resolve them. They will not want to share all of their knowledge, and, even if they did, the overheads in establishing a coherent knowledge base would be too high. Instead, what the agents require is a mechanism for resolving, on an "as required" basis, the conflicts that arise between the arguments that they make in a way which corresponds to the sound principles for merging information that have previously been proposed. Such a mechanism is what we have discussed in this paper. As a result, we know that if we adopt the mechanisms suggested here as part of our approach to inter-agent dialogue [2] (and doing this is simple since we already make use of the underlying argumentation framework), we have the choice of a range of conflict resolution mechanisms each of which relates exactly to one of the options argued for elsewhere in the literature. So, the work discussed in this paper allows us to resolve conflicts in multi-agent dialogues.

5. Conclusions and future work

The work reported here concerns reasoning with conflicting knowledge-bases, that is knowledge bases which are mutually inconsistent in a classical sense. Our first main contribution is to propose a preference-based argumentation framework which resolves the conflicts. This approach is different from the ones existing in the literature. The existing approaches consist of first merging the various knowledge bases to obtain a unique one and then draw conslusions from the new base. In contrast, our approach allows arguments to be built from separate knowledge bases, and the arguments to then be merged. This method of obatining conclusions is much more practical in the context of multi-agent systems. The second contribution of this paper is to show that the preference-based argumentation framework we introduce is general enough to capture the results of some of the merging operators which have been developed. To cover the works proposed in [12] and [13], we have proposed two new preference relations between the arguments, relations which are based on Dalal's distance. Thus we can obtain all the advantages of the approaches for merging conflicting information, we can draw the same conclusions, but without having to actually construct the merged knowledge base.

An extension of this work would be the study of the properties of the new preference relations. Another immediate extension would be to consider several inconsistent knowledge bases instead of consistent ones, thus we can develop a distributed argumentation framework. This looks likely to be very useful in multi-agent systems where each agent is supposed to have its own (consistent / inconsistent) knowledge base.

References

[1] L. Amgoud, C. Cayrol. A model of reasoning based on the production of acceptable arguments. *Accepted for publication in the proceedings of the 8th International Workshop on Non-Monotonic Reasoning NMR'2000, session Uncertainty frameworks in NMR*.

[2] L. Amgoud, S. Parsons, N. Maudet. Modeling dialogues using argumentation. Accepted for publication *in: Proc. of the 4th International conference on multi-agent systems. ICMAS'2000*.

[3] L. Amgoud, S. Parsons. An argumentation framework for reasoning with conflicting bases. *Technical Report, Department of Electronic Engineering, Queen Mary and Westfield College, 2000*.

[4] S. Benferhat, D. Dubois, H. Prade. Argumentative Inference in Uncertain and Inconsistent Knowledge Bases. *In: Proc. of the 9th Conference on Uncertainty in Artificial Intelligence.* pp. 411- 419, 1993.

[5] C. Cayrol. On the relation between Argumentation and Non-monotonic Coherence-based Entailment. *In: Proc. of the 14th International Joint Conference on Artificial Intelligence, IJCAI'95*. pp. 1443 - 1448, 1995.

[6] L. Cholvy. A general framework for reasoning about contradictory information and

some of its applications. *In: Proc. of ECAI workshop "Conflicts among Agents", ECAI'98, 1998.*

[7] C. Cayrol, V. Royer, C. Saurel 1993. Management of preferences in Assumption-Based Reasoning. *Lecture Notes in Computer Science (B. Bouchon-Meunier, L. Valverde, R.Y. Yager Eds.)*, Vol. 682. pp. 13 - 22.

[8] M. Dalal. Investigations into a theory of knowledge base revision: preliminary report. *In: Proc. of AAAI'88*. pp. 475-479, 1988.

[9] P. M. Dung. On the acceptability of arguments and its fundamental role in non-monotonic reasoning, logic programming and *n*-person games. *Artificial Intelligence*. Vol.77, pp. 321-357, 1995.

[10] M. Elvang-Goransson, A. Hunter. 1995. Argumentative logics: Reasoning with classically inconsistent information. *Data & Knowledge Engineering.* Vol. 16, pp. 125-145.

[11] S. Konieczny, R. Pino Perez. On the logic of merging. *In: Proc. of KR'98*. pp. 488-498.

[12] J. Lin. Integration of weighted knowledge bases. *Artificial Intelligence*. Vol.83, pp. 3631-378, 1996.

[13] J. Lin, A. O. Mendelzon. Merging databases under constraints. *International journal of Coop. Inf. Sys*. Vol. 7, pp. 55-76, 1998.

[14] N. Rescher, R. Manor. On inference from inconsistent premises. *Theory and decision*. Vol. 1. pp. 179 - 219, 1970.

[15] P. Z. Revescz. On the semantics of arbitration. *International journal of Algebra and Computation*. Vol. 7, pp. 133-160, 1997.