

# Arguing over ontology alignments

L. Laera<sup>1</sup>, V. Tamma<sup>1</sup>, J. Euzenat<sup>2</sup>, T. Bench-Capon<sup>1</sup>, and T. Payne<sup>3</sup>

<sup>1</sup> Department of Computer Science, University of Liverpool, UK

<sup>2</sup> INRIA Rhône-Alpes, Montbonnot, France

<sup>3</sup> Computer Science, University of Southampton, UK

## 1 Introduction

Traditionally ontologies have been used to achieve semantic interoperability between software applications, as such applications provide the definitions of the vocabularies they use to describe the world [14], and they have proved especially effective when systems are embedded in open, dynamic environments, such as the Web and the Semantic Web [6]. Interoperability relies on the ability to reconcile the differences between heterogeneous ontologies [22]. This reconciliation usually relies on the existence of correspondences (or mappings) between different ontologies (*ontology alignment* [12]), and uses them in order to interpret or translate messages exchanged by applications. Such correspondences may be generated by a variety of different matching algorithms [17]<sup>4</sup>, and their production usually requires several steps. These can include the definition of an initial alignment, the training of some examples, and that invariably involves an some form of interpretation of preliminary results [11]. Therefore, approaches to ontology alignment can only be effective when used to support semantic interoperation at *design time* in closed or partially open environments, where the actors involved are often known, where ontology changes are controlled and thus the alignments can be established before the systems interact. However, these approaches are not sufficient to support semantic interoperation in open environments, where systems can dynamically join or leave and no prior assumption can be made on the ontologies to align. In such environments, the different systems involved need to agree on the semantics of the terms used during the interoperation, and reaching this agreement can only come through some sort of negotiation process [1].

This paper extends the notion of reaching agreement through automated negotiation (*i.e.* without human intervention) by considering the type of systems that need to interoperate, which can affect how the negotiation should proceed. Specifically, *autonomous agents* (within an open environment) may perform different tasks depending on their state and the service providers they interact with. Thus, such agents will differ in the domain ontologies they commit to [14]; and their perception of the world (and hence the choice of vocabulary used to represent concepts). Imposing a single, universally shared ontology on agents is not only impractical because it would result in assuming a standard communication vocabulary (and thus violate the dynamics of open environments) but it also does not take into account the conceptual requirements of services that could appear in future. Instead, every agent assumes its own heterogeneous private ontology, which may not be understandable by other agents. The availability of *Alignment Services* that are able to provide correspondences between two ontologies is only

---

<sup>4</sup> A comprehensive review can be found at <http://www.ontologymatching.org>

a potential solution to achieving interoperability between agents, as any given candidate set of alignments is only suitable to certain contexts. For a given context, agents might have different and inconsistent perspectives; *i.e.* *interests* and *preferences*, on the acceptability of a candidate mapping, each of which may be rationally acceptable. This may be due to the subjective nature of ontologies, to the context and the requirement of the alignments and so on. For example, an agent may be interested in accepting only those mappings that have linguistic similarities, since its ontology is too *structurally simple* to realise any other type of mismatch. In addition, any decision on the acceptability of these mappings has to be made dynamically (at run time), due to the fact that the agents have no prior knowledge of either the existence or constraints of other agents.

In order to address this problem, we present a framework to support agents negotiate agreement on the terminology they use in order to communicate, by allowing them to express their preferred choices over candidate correspondences. This is achieved by adapting argument-based negotiation to deal specifically with arguments that support or oppose the proposed correspondences between ontologies. The set of potential arguments are clearly identified and grounded on the underlying ontology languages, and the kinds of mapping that can be supported by any such argument are clearly specified. In order to compute the preferred ontology alignments for each agent, we use a value-based argumentation framework [5], allowing each agent to express its preferences between the categories of arguments that are clearly identified in the context of ontology alignment. Our approach is able to give a formal motivation for the selection of any correspondence, and enables consideration of an agents' interests and preferences that may influence the selection of a given correspondence. Therefore, this work provides a concrete instantiation of the "meaning negotiation" process that we would like agents to achieve. Moreover, in contrast to current ontology matching procedures, the choice of alignment is based on two clearly identified elements: (i) the argumentation framework, which is common to all agents, and (ii) the preference relations which are private to each agent.

The remainder of this paper is structured as follows. Section 2 presents the argumentation framework and how it can be used. Section 3 defines the various categories of arguments that can support or attack mappings, and defines the notion of agreed and agreeable alignments for agents, whereas Section 4 proposes a procedure to find them. An example illustrating the argumentation process is given in Section 5, followed concluding remarks in Section 6<sup>5</sup>.

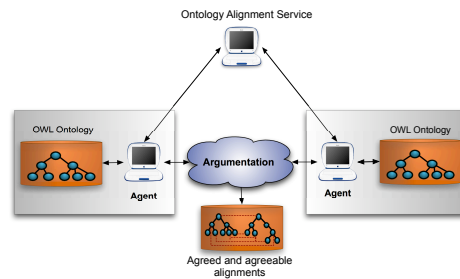
## 2 Argumentation framework

This paper focuses on autonomous agents situated within an open system. Each agent has a knowledge base, expressed using one of several possible ontologies. The *mental attitudes* of an agent towards correspondences are represented in terms of *interests* and *preferences*, which represent the motivations of the agent, and thus determine whether a mapping is accepted or rejected. The preferences are represented as a (partial or total) pre-ordering of preferences over different types of ontology mismatches (*Pref*)<sup>6</sup>.

<sup>5</sup> A survey of related work is given in an extended version of this paper [15].

<sup>6</sup> Although the agents' ontologies may differ, we eliminate the problem of integrating different ontology languages by assuming that ontologies are encoded in the same language, *i.e.* OWL.

For agents to communicate, they first need to establish a mutually acceptable set of alignments between their ontologies. Potential alignments are generated at design time (by a variety of different ontology-matching approaches [17]), and provided at runtime by a dedicated agent, called an *Ontology Alignment Service (OAS)* (Figure 1). An alignment consists of a set of all possible correspondences between the two ontologies. A correspondence (or a mapping) can be described as a tuple:  $m = \langle e, e', n, R \rangle$ , where  $e$  and  $e'$  are the entities (concepts, relations or individuals) between which a relation is asserted by the correspondence;  $n$  is a degree of confidence in that correspondence; and  $R$  is the relation (e.g., equivalence, more general, etc.) holding between  $e$  and  $e'$  asserted by the correspondence [17]. A *candidate mapping* is a correspondence (provided by an OAS) that could be used by the agents to align their ontologies. Each correspondence  $m$  is accompanied by a set of justifications  $G$ , which provide an explanation as to why the correspondence was generated<sup>7</sup>. This information is used by the agents when generating and exchanging arguments, for and against a candidate mapping. In addition, every agent has a private threshold value  $\varepsilon$  which will be compared to the degree of confidence,  $n$ , of a mapping, to decide whether it should be considered.



**Fig. 1.** Reaching agreement over ontology alignments

In order for the agents to consider potential mappings and the reasons for and against accepting them, we use an argumentation framework based on *Value-based Argument Frameworks (VAFs)* [5], that extends Dong’s classical argument system [9]<sup>8</sup>.

**Definition 1.** An *Argumentation Framework (AF)* is a pair  $AF = \langle AR, A \rangle$ , where  $AR$  is a set of arguments and  $A \subset AR \times AR$  is the attack relationship for  $AF$ .  $A$  comprises a set of ordered pairs of distinct arguments in  $AR$ . A pair  $\langle x, y \rangle$  is referred to as “ $x$  attacks  $y$ ”. We also say that a set of arguments  $S$  attacks an argument  $y$  if  $y$  is attacked by an argument in  $S$ .

An argumentation framework can be simply represented as a directed graph whose vertices are the arguments and whose edges correspond to the elements of  $A$ . In this paper, we are concerned only with arguments about mappings. We can therefore define arguments as follows:

**Definition 2.** An argument  $x \in AF$  is a triple  $x = \langle G, m, \sigma \rangle$  where  $m$  is a correspondence  $\langle e, e', n, R \rangle$ ;  $G$  is the grounds justifying a *prima facie* belief that the correspondence does, or

<sup>7</sup> Although few approaches for ontology alignment provide justifications [18, 7], tools such as [10] combine different similarity metrics which can be used to provide necessary justifications.

<sup>8</sup> More details can be found in an extended version of this paper [15].

does not hold;  $\sigma$  is one of  $\{+, -\}$  depending on whether the argument is that  $m$  does or does not hold.

An argument  $x$  is attacked by the assertion of its negation  $\neg x$ , namely the *counter-argument*, defined as follows:

**Definition 3.** An argument  $y \in AF$  rebuts an argument  $x \in AF$  if  $x$  and  $y$  are arguments for the same mapping but with different signs, e.g. if  $x$  and  $y$  are in the form  $x = \langle G_1, m, + \rangle$  and  $y = \langle G_2, m, - \rangle$ ,  $x$  counter-argues  $y$  and vice-versa.

Moreover, if an argument  $x$  supports an argument  $y$ , they form the argument  $(x \rightarrow y)$  that attacks an argument  $\neg y$  and is attacked by argument  $\neg x$ .

When the set of such arguments and counter arguments have been produced, it is necessary for the agents to consider which of them they should accept.

**Definition 4.** Let  $\langle AR, A \rangle$  be an argumentation framework. Let  $R, S$ , subsets of  $AR$ . An argument  $s \in S$  is attacked by  $R$  if there is some  $r \in R$  such that  $\langle r, s \rangle \in A$ . An argument  $x \in AR$  is acceptable with respect to  $S$  if for every  $y \in AR$  that attacks  $x$  there is some  $z \in S$  that attacks  $y$ .  $S$  is conflict free if no argument in  $S$  is attacked by any other argument in  $S$ . A conflict free set  $S$  is admissible if every argument in  $S$  is acceptable with respect to  $S$ .  $S$  is a preferred extension if it is a maximal (with respect to set inclusion) admissible subset of  $AR$ .

In addition, an argument  $x$  is *credulously accepted* if there is some preferred extension containing it; whereas  $x$  is *sceptically accepted* if it is a member of every preferred extension. The key notion here is the *preferred extension* which represents a consistent position within  $AF$ , which is defensible against all attacks and which cannot be further extended without becoming inconsistent or open to attack.

In Dung's framework, attacks always succeed. This is reasonable when dealing with deductive arguments, but in many domains, including the one under consideration, arguments lack this coercive force: they provide reasons which may be more or less persuasive. Moreover, their persuasiveness may vary according to their audience. To handle such defeasible reasons giving arguments, we need to be able to distinguish attacks from successful attacks, those which defeat the attacked argument, therefore we use a Value-based Argumentation Framework, which prescribes different strengths to arguments on the basis of the values they promote and the ranking given to these values by the audience for the argument. This allows us to systematically relate strengths of arguments to their motivations, and to accommodate different audiences with different interests and preferences.

**Definition 5.** A Value-Based Argumentation Framework ( $VAF$ ) is defined as  $\langle AR, A, \mathcal{V}, \eta \rangle$ , where  $(AR, A)$  is an argumentation framework,  $\mathcal{V}$  is a set of  $k$  values which represent the types of arguments and  $\eta: AR \rightarrow \mathcal{V}$  is a mapping that associates a value  $\eta(x) \in \mathcal{V}$  with each argument  $x \in AR$

In section 3, the set of values  $\mathcal{V}$  will be defined as the different types of ontology mismatch, which we use to define the categories of arguments and to assign to each argument one category.

**Definition 6.** An audience for a  $VAF$  is a binary relation  $\mathcal{R} \subset \mathcal{V} \times \mathcal{V}$  whose (irreflexive) transitive closure,  $\mathcal{R}^*$ , is asymmetric, i.e. at most one of  $(v, v')$ ,  $(v', v)$  are members of  $\mathcal{R}^*$  for

any distinct  $v, v' \in \mathcal{V}$ . We say that  $v_i$  is preferred to  $v_j$  in the audience  $\mathcal{R}$ , denoted  $v_i \succ_{\mathcal{R}} v_j$ , if  $(v_i, v_j) \in \mathcal{R}^*$ .

Let  $\mathcal{R}$  be an audience,  $\alpha$  is a specific audience (compatible with  $\mathcal{R}$ ) if  $\alpha$  is a total ordering of  $\mathcal{V}$  and  $\forall v, v' \in \mathcal{V}, (v, v') \in \alpha \Rightarrow (v', v) \notin \mathcal{R}^*$

In this way, we take into account that different agents (represented by different audiences) can have different perspectives on the same candidate mapping. Acceptability of an argument is defined in the following way: <sup>9</sup>

**Definition 7.** Let  $\langle AR, A, \mathcal{V}, \eta \rangle$  be a VAF and  $\mathcal{R}$  an audience.

- a. For arguments  $x, y$  in  $AR$ ,  $x$  is a successful attack on  $y$  (or  $x$  defeats  $y$ ) with respect to the audience  $\mathcal{R}$  if:  $(x, y) \in \mathcal{A}$  and it is not the case that  $\eta(y) \succ_{\mathcal{R}} \eta(x)$ .
- b. An argument  $x$  is acceptable to the subset  $S$  with respect to an audience  $\mathcal{R}$  if: for every  $y \in AR$  that successfully attacks  $x$  with respect to  $\mathcal{R}$ , there is some  $z \in S$  that successfully attacks  $y$  with respect to  $\mathcal{R}$ .
- c. A subset  $S$  of  $AR$  is conflict-free with respect to the audience  $\mathcal{R}$  if: for each  $(x, y) \in S \times S$ , either  $(x, y) \notin \mathcal{A}$  or  $\eta(y) \succ_{\mathcal{R}} \eta(x)$ .
- d. A subset  $S$  of  $AR$  is admissible with respect to the audience  $\mathcal{R}$  if:  $S$  is conflict free with respect to  $\mathcal{R}$  and every  $x \in S$  is acceptable to  $S$  with respect to  $\mathcal{R}$ .
- e. A subset  $S$  is a preferred extension for the audience  $\mathcal{R}$  if it is a maximal admissible set with respect to  $\mathcal{R}$ .
- f. A subset  $S$  is a stable extension for the audience  $\mathcal{R}$  if  $S$  is admissible with respect to  $\mathcal{R}$  and for all  $y \notin S$  there is some  $x \in S$  which successfully attacks  $y$  with respect to  $\mathcal{R}$ .

In order to determine whether the dispute is resolvable, and if it is, to determine the preferred extension with respect to a value ordering promoted by distinct audiences, [5] introduces the notion of objective and subjective acceptance as follows:

**Definition 8.** Given a VAF,  $\langle AR, A, \mathcal{V}, \eta \rangle$ , an argument  $x \in AR$  is subjectively acceptable if and only if,  $x$  appears in the preferred extension for some specific audiences but not all. An argument  $x \in AR$  is objectively acceptable if and only if,  $x$  appears in the preferred extension for every specific audience. An argument which is neither objectively nor subjectively acceptable is said to be indefensible.

### 3 Arguments for correspondences

Potential arguments are clearly identified and grounded on the underlying ontology language OWL. Therefore, the grounds justifying correspondences can be extracted from the knowledge in ontologies<sup>10</sup>. Our classification of the grounds justifying correspondences is the following:

**semantic ( $M$ ):** the sets of models of two entities do or do not compare;

**internal structural ( $IS$ ):** two entities share more or less internal structure (e.g., the value range or cardinality of their attributes);

**external structural ( $ES$ ):** the set of relations, each of two entities have, with other entities do or do not compare;

**terminological ( $T$ ):** the names of two entities share more or less lexical features;

**extensional ( $E$ ):** the known extension of two entities do or do not compare.

<sup>9</sup> Note that all these notions are now relative to some audience.

<sup>10</sup> This knowledge includes both the extensional and intensional OWL ontology definitions.

These categories correspond to the type of categorizations underlying ontology matching algorithms [22]. In our framework, we will use the types of arguments described above as types for the *VAF*; hence  $\mathcal{V} = \{M, IS, ES, T, E\}$ . For example, an audience may specify that terminological arguments are preferred to semantic arguments, or vice versa. Note that this may vary according to the nature of the ontologies being aligned. Semantic arguments will be given more weight in a fully axiomatised ontology, compared to that in a lightweight ontology where there is very little reliable semantic information on which to base such arguments.

Table 1 presents a sample set of argument schemes, instantiations of which will comprise *AR*. Attacks between these arguments will arise when we have arguments for the same mapping but with conflicting values of  $\sigma$ , thus yielding attacks that can be considered symmetric. Moreover, the relations in the mappings can also give rise to attacks: if relations are not deemed exclusive, an argument against inclusion is a fortiori an argument against equivalence (which is more general).

*Example 1.* Consider a candidate mapping  $m = \langle c, c', -, \equiv \rangle$  between two OWL ontologies  $O_1$  and  $O_2$ , with concepts  $c$  and  $c'$  respectively. An argument for accepting the mapping  $m$  may be that the labels of  $c$  and  $c'$  are synonymous. An argument against may be that some of their super-concepts are not mapped.

In *VAFs*, arguments against or in favour of a candidate mapping are seen as grounded on their type. In this way, we are able to motivate the choice between preferred extensions by reference to the type ordering of the audience concerned. Moreover, the pre-ordering of preferences *Pref* for each agent will be over  $\mathcal{V}$ , that corresponds to the determination of an audience. Specifically, for each candidate mapping  $m$ , if there exist justification(s)  $G$  for  $m$  that corresponds to the highest preferences *Pref* (with the respect of the pre-ordering), assuming  $n$  is greater than its private threshold  $\varepsilon$ , an agent will generate arguments  $x = \langle G, m, + \rangle$  (or  $x = \langle G, m, - \rangle$  otherwise), by instantiating the argumentation schema.

**Table 1.** Argument scheme for OWL ontological alignments

Mapping	$\sigma$	Grounds	Comment
$\langle e, e', n, \equiv \rangle$	+	$\exists m_i = \langle ES(e), ES(e'), n', \equiv \rangle$	$e$ and $e'$ have mapped neighbours (e.g., super-entities, sibling-entities, etc.) of $e$ are mapped in those of $e'$
$\langle e, e', n, \sqsubseteq \rangle$	+	$\exists m_i = \langle ES(e), ES(e'), n', \equiv \rangle$	(some or all) Neighbours (e.g., super-entities, sibling-entities, etc.) of $e$ are mapped in those of $e'$
$\langle c, c', n, \sqsubseteq \rangle$	+	$\exists m_i = \langle IS(c), IS(c'), n', \equiv \rangle$	(some or all) Properties of concept $c$ are mapped to those of concept $c'$
$\langle c, c', n, \sqsupset \rangle$	-	$\bar{\exists} m_i = \langle IS(c), IS(c'), n', \equiv \rangle$	No properties of $c$ are mapped to those of $c'$
$\langle e, e', n, \equiv \rangle$	+	$\exists m_i = \langle E(e), E(e'), n', \equiv \rangle$	(some or all) Instances of $e$ and $e'$ are mapped
$\langle e, e', n, \sqsubseteq \rangle$	+	$\exists m_i = \langle E(e), E(e'), n', \equiv \rangle$	(some or all) Instances of $e$ are mapped to those of $e'$
$\langle e, e', n, \equiv \rangle$	+	$label(e) \approx_T label(e')$	Entities' labels share lexical features (e.g., synonyms and lexical variants)
$\langle e, e', n, \sqsubseteq \rangle$	-	$label(e) \not\approx_T label(e')$	Entities' labels do not share lexical features (e.g., homonyms)

Although in *VAFs* there is always a unique non-empty preferred extension with respect to a specific audience, provided the *AF* does not contain any cycles in a single argument type, an agent may have multiple preferred extensions either because no preference between two values in a cycle has been expressed, or because a cycle in a single value exists. The first may be eliminated by committing to a specific audience,

but the second cannot be eliminated in this way. In our domain, where many attacks are symmetric, two cycles will be frequent and in general an audience may have multiple preferred extensions.

Thus, given a set of arguments justifying mappings organised into an argumentation framework, an agent will be able to determine which mappings are acceptable by computing the preferred extensions with respect to its preferences. If there are multiple preferred extensions, the agent must commit to the arguments present in all preferred extensions, but it has some freedom of choice with respect to those in some but not all of them. This will partition arguments into three sets: *desired arguments*, present in all preferred extensions; *optional arguments*, present in some but not all; and *rejected arguments*, present in none. If we have two agents belonging to different audiences, these sets may differ.

Based on the above considerations, we thus define an *agreed alignment* and an *agreeable alignment* as follows. An *agreed alignment* is the set of correspondences supported<sup>11</sup> by those arguments which are in every preferred extension of every agent. An *agreeable alignment* extends the agreed alignment with those correspondences supported by arguments which are in some preferred extension of every agent. Whilst the mappings included in the agreed alignments can be considered valid and consensual for all agents, the agreeable alignments have a uncertain background, due to the different alternative positions that each agent can take. However, given our context of agent communication, we seek to accept as many candidate mappings as possible. We will therefore take into consideration both set of alignments - agreed and agreeable.

## 4 Instantiating argumentation frameworks

In order to reach agent consensus about ontology alignments, first we have to build the argumentation frameworks and evaluate them to find which arguments are agreed and agreeable. There are four main steps in applying our argumentation approach:

1. Given a single agent, and for each candidate mapping, we construct an argumentation framework by considering the repertoire of argument schemes available to the agent, and constructing a set of arguments by instantiating these schemes with respect to the interests of the agent. Each argument either supports or rejects the conclusion that the mapping is valid. Internally, an argument is represented by a simple identifier (letter A,B,C, etc.), the type of value which it promoted, and optionally, the agent(s) introducing the argument. Having established the set of arguments, we then determine the attacks between them by considering their mappings and signs, and the other factors discussed above. The formulation of suitable attacks is a key part of representing the different point of views of agents. Arguments may have different strength, which depends on the values they promote. Therefore, an attack can fail, since the attacked argument may be stronger than its attacker.
2. Given multiple agents, we simply merge their individual frameworks by forming the union of their individual argument sets and individual attack relations, and then extend the attack relation by computing attacks between the arguments present in the framework of one, but not both, agents.
3. Then, for each *VAF*, we determine which of the arguments are undefeated by attacks from other arguments. We employ the algorithm in [4] for computing the preferred extensions of a

---

<sup>11</sup> Note that a correspondence  $m$  is *supported* by an argument  $x$  if  $x$  is  $\langle G, m, + \rangle$

value-based argumentation framework given a value ordering. The global view is considered by taking the union of these preferred extensions for each audience.

4. Finally, we consider which arguments are in every preferred extension of every audience. The mappings that have only arguments *for* will be included in the agreed alignments, and those *against* will be rejected. For those mappings where we cannot establish their acceptability, we extend our search space to consider those arguments which are in some preferred extension of every audience. The mappings supported by those arguments are part of the set of agreeable alignments. Figure 2 shows how to find such alignments.

The dialogue between agents can thus consist simply of the exchange of individual argumentation frameworks, from which they can individually compute acceptable mappings. If necessary and desirable, these can then be reconciled into a mutually acceptable position through a process of negotiation [8]. In the course of constructing a position, an ordering of values best able to satisfy the joint interests of the agents concerned is determined. These issues are the subject of ongoing research.

---

**Algorithm 1** Find agreed and agreeable alignments

---

**Require:** a set of *VAFs*  $\langle AR, A, \mathcal{V}, \eta \rangle$ , a set of audiences  $\mathcal{R}_i$ , a set of candidate mappings  $M$

**Ensure:** Agreed alignments  $AG$  and agreeable alignments  $AG_{ext}$

```

1:  $AG := \emptyset$ 
2:  $AG_{ext} := \emptyset$ 
3: for all audience  $\mathcal{R}_i$  do
4:   for all VAF do
5:     compute the preferred extensions for  $\mathcal{R}_i$ ,  $P_j(\langle AR, A, \mathcal{V}, \eta \rangle, \mathcal{R}_i)$ ,  $j \geq 1$ 
6:   end for
7:    $P_k(\mathcal{R}_i) := \bigcup_j P_j(\langle AR, A, \mathcal{V}, \eta \rangle, \mathcal{R}_i)$ ,  $k \geq 1$ 
8:   end for
9:    $AGArg := x \in \bigcap_{k,i} P_k(\mathcal{R}_i)$ ,  $\forall k \geq 1, \forall i \geq 0$ 
10:  for all  $x \in AGArg$  do
11:    if  $x$  is  $\langle G, m, + \rangle$  then
12:       $AG := AG \cup \{m\}$ 
13:    else
14:      reject mapping  $m$ 
15:    end if
16:  end for
17:  if  $\exists m \in M$  such that  $m$  is neither in  $AG$  and rejected then
18:     $AG_{ext} := x \in \bigcap_i P_k(\mathcal{R}_i)$ ,  $\forall i \geq 0, k \geq 1$ 
19:    for all  $x \in AG_{ext}$  do
20:      if  $x$  is  $\langle G, m, + \rangle$  then
21:         $AG_{ext} := AG_{ext} \cup \{m\}$ 
22:      end if
23:    end for
24:  end if

```

---

**Fig. 2.** Find agreed and agreeable alignments

## 5 A walk through example

Let us assume that some agents or services need to interact with each other using two independent but overlapping ontologies. The first agent,  $Ag_1$  uses the bibliographic ontology<sup>12</sup> from the University of Toronto, based on bibTeX; whereas the second agent,  $Ag_2$ , uses the General University Ontology<sup>13</sup> from Mondeca<sup>14</sup>. For space reasons, we only consider a subset of these ontologies, shown in Table 2, where the first and second ontologies are represented by  $O_1$  and  $O_2$  respectively.

We will reason about the following candidate mappings, provided by the *Ontology Alignment Service* (OAS):

<sup>12</sup> <http://www.cs.toronto.edu/semanticweb/maponto/ontologies/BibTex.owl>

<sup>13</sup> <http://www.mondeca.com/owl/mondes/univ.owl>

<sup>14</sup> Note that ontology  $O_2$  has been slightly modified for the purposes of this example.



$$\begin{aligned}
m_1 &= \langle O_1: \textit{Press}, O_2: \textit{Periodical}, n, = \rangle;^{15} \\
m_2 &= \langle O_1: \textit{publication}, O_2: \textit{Publication}, n, = \rangle; \\
m_3 &= \langle O_1: \textit{hasPublisher}, O_2: \textit{publishedBy}, n, = \rangle; \\
m_4 &= \langle O_1: \textit{Magazine}, O_2: \textit{Magazine}, n, = \rangle; \\
m_5 &= \langle O_1: \textit{Newspaper}, O_2: \textit{Newspaper}, n, = \rangle; \\
m_6 &= \langle O_1: \textit{Organization}, O_2: \textit{Organization}, n, = \rangle.
\end{aligned}$$

The generation of the arguments and counter-arguments is based on the agent's preferences and threshold. However, here we assume a degree of confidence  $n$  that is above the threshold of each agent, and so will not influence their acceptability.

Assume now that there are two possible audiences,  $\mathcal{R}_1$ , which prefers terminology to external structure, ( $T \succ_{\mathcal{R}_1} ES$ ), and  $\mathcal{R}_2$ , which prefers external structure to terminology ( $ES \succ_{\mathcal{R}_2} T$ ). The pre-ordering of preference  $Pref$  will correspond to the agents's audience.

**Table 2.** Excerpts of  $O_1$  and  $O_2$  ontologies

$O_1$ Ontology	$O_2$ Ontology
$Artifact \sqsubseteq \top$	$Document \sqsubseteq \top$
$Print\_Media \sqsubseteq Artifact$	$Publication \sqsubseteq Document$
$Press \sqsubseteq Print\_Media$	$Periodical \sqsubseteq Publication$
$Magazine \sqsubseteq Press$	$Magazine \sqsubseteq Periodical$
$Newspaper \sqsubseteq Press$	$Newspaper \sqsubseteq Periodical$
$publication \sqsubseteq \forall hasPublisher.Publisher$	$Newsletter \sqsubseteq Periodical$
$publication \sqsubseteq Print\_Media$	$Journal \sqsubseteq Periodical$
$Publisher \sqsubseteq Organization$	$Publication \sqsubseteq Document$
	$Publication \sqsubseteq \forall publishedBy.Organization$

We can identify a set of arguments and the attacks between them. We assume that a set of arguments is generated by instantiating the argumentation schemes (Table 1), with respect to the interests and preferences  $Pref$  of the agents and taking into consideration the justifications  $G$ . Table 3 shows each argument, labeled with an identifier  $Id$ , its type  $\mathcal{V}$ , and the attacks  $A$  that can be made on it by opposing arguments. Based upon these arguments and the attacks, we can construct the argumentation frameworks which bring the arguments together so that they can be evaluated. These are shown in Figure 3, where nodes represent arguments (labelled with their  $Id$ ) with the respective type value  $\mathcal{V}$ . The arcs represent the attacks  $A$ , whereas the direction of the arcs represents the direction of the attack. By instantiating the general VAF according to their own preferences,  $Ag_1$  and  $Ag_2$  obtain two possible argumentation frameworks, (a) and (b). In the argumentation framework (a), we have two arguments against  $m_1$ , and one for it:


- $A$  is against the correspondence  $m_1$ , since none of the super-concepts of the  $O_1: Press$  are mapped to any super-concept of  $O_2: Periodical$ .
- $B$  argues for  $m_1$  because two sub-concepts of  $O_1: Press$ , ( $O_1: Magazine$  and  $O_1: Newspaper$ ), are mapped to two sub-concepts of  $O_2: Periodical$ , ( $O_2: Magazine$  and  $O_2: Newspaper$ ), as established by  $m_4$  and  $m_5$ .
- $C$  argues against  $m_1$ , because  $Press$  and  $Periodical$  do not have any lexical similarity.

<sup>15</sup>  $m_1$  states an equivalence correspondence with confidence  $n$  between the concept  $Press$  in the ontology  $O_1$  and the concept  $Periodical$  in the ontology  $O_2$

**Table 3.** Arguments for and against the correspondences  $m_1, m_2, m_3, m_4, m_5$  and  $m_6$

Id	Argument	A	V
A	$\langle \bar{A}m = \langle \text{superconcept}(\text{Press}), \text{superconcept}(\text{Periodical}), n, \equiv, \rangle, m_1, - \rangle$	B,L,O	ES
B	$\langle \exists m = \langle \text{subconcept}(\text{Press}), \text{subconcept}(\text{Periodical}), n, \equiv, \rangle, m_1, + \rangle$	A,C	ES
C	$\langle \text{Label}(\text{Press}) \not\approx_T \text{Label}(\text{Periodical}), m_1, - \rangle$	B	T
D	$\langle \text{Label}(\text{publication}) \approx_T \text{Label}(\text{Publication}), m_2, + \rangle$	E	T
E	$\langle \bar{A}m = \langle \text{superconcept}(\text{publication}), \text{superconcept}(\text{Publication}), n, \equiv, \rangle, m_2, - \rangle$	D,F	ES
F	$\langle \exists m = \langle \text{property}(\text{publication}), \text{property}(\text{Publication}), n, \equiv, \rangle, m_2, + \rangle$	E	IS
G	$\langle \bar{A}m = \langle \text{range}(\text{hasPublisher}), \text{range}(\text{publishedBy}), n, \equiv, \rangle, m_3, - \rangle$	F,H	IS
H	$\langle \text{Label}(\text{hasPublisher}) \approx_T \text{Label}(\text{publishedBy}), m_3, + \rangle$	G	T
I	$\langle \exists m = \langle \text{superconcept}(\text{Publisher}), \text{Organization}, n, \equiv, \rangle, m_7, + \rangle$	G	ES
J	$\langle \text{Label}(\text{Magazine}) \approx_T \text{Label}(\text{Magazine}), m_4, + \rangle$		T
K	$\langle \exists m = \langle \text{siblingConcept}(\text{Magazine}), \text{siblingConcept}(\text{Magazine}), n, \equiv, \rangle, m_4, + \rangle$		ES
L	$\langle \exists m = \langle \text{superconcept}(\text{Magazine}), \text{superconcept}(\text{Magazine}), n, \equiv, \rangle, m_4, + \rangle$		ES
M	$\langle \text{Label}(\text{Newspaper}) \approx_T \text{Label}(\text{Newspaper}), m_5, + \rangle$		T
N	$\langle \exists m = \langle \text{siblingConcept}(\text{Newspaper}), \text{siblingConcept}(\text{Newspaper}), m_5, + \rangle$		ES
O	$\langle \exists m = \langle \text{superconcept}(\text{Newspaper}), \text{superconcept}(\text{Newspaper}), n, \equiv, \rangle, m_5, + \rangle$		ES
P	$\langle \text{Label}(\text{Organization}) \approx_T \text{Label}(\text{Organization}), m_6, + \rangle$		T

Moreover, we have six arguments supporting the correspondences  $m_4, m_5$  and  $m_6$ .  $K, L$  and  $M$  justify the mapping  $m_4$ , since, respectively, the labels of  $O_1$ : *Magazine* and  $O_2$ : *Magazine* are lexically similar; their siblings are mapped, as established by  $m_5$  and their super-concepts;  $O_1$ : *Press* and  $O_2$ : *Periodical* are mapped by  $m_1$ . There is a similar situation for the arguments  $M, N$  and  $O$ . Clearly, argument  $A$  attacks the arguments  $L$  and  $O$ .



**Fig. 3.** Value-Based Argumentation Frameworks

In the second argumentation framework (b) we relate the following arguments:  $D$  justifies the mapping  $m_2$ , since the labels of  $O_1$ : *publication* and  $O_2$ : *Publication* are lexically similar. Their super-concepts, however, are not mapped (argument  $E$ ). Argument  $F$  is based on the fact that  $O_1$ : *publication* and  $O_2$ : *Publication* have mapped properties,  $O_1$ : *hasPublisher* and  $O_2$ : *publishedBy*, as defined in  $m_3$ .  $F$  is then attacked by  $G$ , which states that the range of these properties, respectively  $O_1$ : *Publisher* and  $O_2$ : *Organization*, are not mapped. This is in turn counter-attacked by the arguments  $H$  and  $I$ .  $H$  states the mapping  $m_3$  is correct, since  $O_1$ : *hasPublisher* and  $O_2$ : *publishedBy* are lexically similar; whereas  $I$  attacks the justification of  $G$  stating that the ranges of these properties are similar, since a super-concept of  $O_1$ : *Publisher*, ( $O_1$ : *Organization*), is already mapped to  $O_2$ : *Organization*.  $P$  states that  $O_1$ : *Organization* and  $O_2$ : *Organization* are mapped since their labels are lexically similar.

The above analysis gives different, but sometimes overlapping reasons to argue for and against several candidate mappings. Given the audiences,  $\mathcal{R}_1$  and  $\mathcal{R}_2$ , the preferred extensions for the union of the argumentation frameworks (a) and (b) are:

Preferred Extensions for the union of (a) and (b)	Audience
$\{A, C, J, K, M, N, D, F, I, H, P\}$	$\mathcal{R}_1$
$\{A, C, J, K, M, N, D, F, I, H, P\}, \{B, O, L, J, K, M, N, D, F, I, H, P\}$	$\mathcal{R}_2$
$\{A, C, J, K, M, N, E, I, H, P\}, \{B, O, L, J, K, M, N, E, I, H, P\}$	

Therefore, the arguments accepted by both audiences are  $\{I, H, J, K, M, N, P\}$ . Arguments  $A, C, D, E$ , and  $F$  are, however, all potentially acceptable, since both audiences can choose to accept them, as they appear in some preferred extension for each audience. This means that the mapping  $m_1$  will be rejected (since  $B$  is unacceptable to  $\mathcal{R}_1$ ), while the mappings  $m_3, m_4, m_5$  and  $m_6$  will be all accepted (they are all accepted by  $\mathcal{R}_1$  and all acceptable to  $\mathcal{R}_2$ ).  $m_2$  will also be acceptable as the arguments supporting it are in some preferred extension for these audiences. The *agreed alignments* are then  $m_3, m_4, m_5$  and  $m_6$ , while the *agreeable alignment* is  $m_2$ . Interestingly, should an agent wish to reject the mappings  $m_2$  and  $m_3$ , it can achieve this by considering a new audience  $\mathcal{R}_3$ , whose internal structure is valued more than external structure, which is valued more than terminology ( $IS \succ_{\mathcal{R}_3} ES \succ_{\mathcal{R}_3} T$ ). In this case, the preferred extension from framework (b) is  $\{E, G, I, P\}$ , since the new preference allows  $G$  to defeat  $H$  and resist  $I$ .  $G$  will also defeat  $F$  leaving  $E$  available to defeat  $D$ . This clearly shows how the acceptability of an argument crucially depends on the audience to which it is addressed.

## 6 Summary and Outlook

In this paper we have outlined a framework that provides a novel way for agents, who use different ontologies, to come to agreement on an alignment. This is achieved using an argumentation process in which candidate correspondences are accepted or rejected, based on the ontological knowledge and the agent's preferences. Argumentation is based on the exchange of arguments, against or in favour of a correspondence, that interact with each other using an *attack* relation. Each argument instantiates an argumentation schema, and utilises domain knowledge, extracted from extensional and intensional ontology definitions. When the full set of arguments and counter-arguments has been produced, the agents consider which of them should be accepted. As we have seen, the acceptability of an argument depends on the ranking - represented by a particular preference ordering on the type of arguments. Our approach is able to give a formal motivation for the selection of a correspondence, and enables consideration of an agent's interests and preferences that may influence the selection of a correspondence. We believe that this approach will aim at reaching mutual understanding and communicative work in agents system more sound and effective. In future work we intend to investigate the use of a negotiation process to enable agents to reach an agreement on a mapping when they differ in their ordering of argument types. Another interesting topic for future work would be to investigate how to argue about the whole alignments, and not only the individual candidate mapping. These arguments could occur when a global similarity measure between the whole ontologies is applied<sup>16</sup>

## References

1. K. Aberer, P. Cudr-Mauroux, A. M. Ouksel, T. Catarci, M.-S. Hacid, A. Illarramendi, V. Kashyap, M. Mecella, E. Mena, E. J. Neuhold, and et al. Emergent semantics principles

<sup>16</sup> **Acknowledgements** The research has been partially supported by Knowledge Web (FP6-IST 2004-507482) and PIPS (FP6-IST 2004-507019). Special thanks to Floriana Grasso and Ian Blacoe for their comments.

- and issues. In *Proceedings of Database Systems for Advances Applications, 9th International Conference, DASFAA 2004*, 2004.
2. S. C. Bailin and W. Truszkowski. Ontology negotiation: How agents can really get to know each other. In *Proceedings of the First International Workshop on Radical Agent Concepts (WRAC 2002)*, 2002.
  3. T. Bench-Capon. Representation of case law as an argumentation framework. In *Proceedings of Legal Knowledge and Information Systems (Jurix 2002)*, 2002.
  4. T. Bench-Capon. Value based argumentation frameworks. In *Proceedings of Non Monotonic Reasoning*, pages 444–453, 2002.
  5. T. Bench-Capon. Persuasion in practical argument using value-based argumentation frameworks. In *Journal of Logic and Computation*, volume 13, pages 429–448, 2003.
  6. T. Berners-Lee, J. Hendler, and O. Lassila. The semantic web. *Scientific American*, 284(5):34–43, 2001.
  7. R. Dhamankar, Y. Lee, A. Doan, A. Halevy, and P. Domingos. imap. In *Proceedings of the International Conference on Management of Data (SIGMOD)*, pages 383–394.
  8. S. Doutre, T. Bench-Capon, and P. E. Dunne. Determining preferences through argumentation. In *Proceedings of AI\*IA'05*, pages 98–109, 2005.
  9. P. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. In *Artificial Intelligence*, volume 77, pages 321–358, 1995.
  10. M. Ehrig and S. Staab. Qom - quick ontology mapping. In *Proceedings of the International Semantic Web Conference*, 2004.
  11. J. Euzenat. Alignment infrastructure for ontology mediation and other applications. In Hepp, editor, *Proceedings of the First International workshop on Mediation in semantic web services*, 2005.
  12. J. Euzenat and P. Valtchev. Similarity-based ontology alignment in owl-lite. In *Proceedings of European Conference on Artificial Intelligence (ECAI04)*, 2004.
  13. N. Green. Argumentation in medical genetics. In *Proceedings of Workshop on Computational Models of Natural Argument*, pages 39–44, 2003.
  14. T. R. Gruber. A translation approach to portable ontology specifications. *Knowledge Acquisition*, 5(2):199–220, 1993.
  15. L. Laera, V. Tamma, J. Euzenat, T. Bench-Capon, and P. T. Reaching agreements over ontology alignments. In *Proceedings of the Fifth International Semantic Web Conference (ISWC'06)*, 2006.
  16. I. Rahwan, S. D. Ramchurn, N. R. Jennings, P. McBurney, S. Parsons, and L. Sonenberg. Argumentation-based negotiation. In *The Knowledge Engineering Review*, volume 18, pages 343–375, 2003.
  17. P. Shvaiko and J. Euzenat. A survey of schema-based matching approaches. *Journal on data semantics*, 4:146–171, 2005.
  18. P. Shvaiko, F. Giunchiglia, P. Pinheiro da Silva, and D. McGuinness. Web explanations for semantic heterogeneity discovery. In *Proceedings of ESWC*, pages 303–317, 2005.
  19. C. Sierra, N. R. Jennings, P. Noriega, and S. Parsons. A framework for argumentation-based negotiation. In *Proceedings of the 4th International Workshop on Intelligent Agents IV, Agent Theories, Architectures, and Languages*, 1997.
  20. N. Silva, P. Maio, and J. Rocha. An approach to ontology mapping negotiation. In *Proceedings of the Workshop on Integrating Ontologies*, 2005.
  21. J. van Diggelen, R. Beun, F. Dignum, R. van Eijk, and J.-J. Meyer. A decentralized approach for establishing a shared communication vocabulary. In *Proceedings of the AMKN*, 2005.
  22. P. Visser, D. Jones, T. Bench-Capon, and M. Shave. Assessing heterogeneity by classifying ontology mismatches. In N. Guarino, editor, *Proceedings of the FOIS'98*, 1998.