Agent-based Argumentation for Ontology Alignments

Loredana Laera and **Valentina Tamma** and **T.J.M. Bench-Capon**[1] and **Jérôme Euzenat** [2]

Abstract. When agents communicate they do not necessarily use the same vocabulary or ontology. For them to interact successfully they must find correspondences between the terms used in their ontologies. While many proposals for matching two agent ontologies have been presented in the literature, the resulting alignment may not be satisfactory to both agents and can become the object of further negotiation between them.

This paper describes our work constructing a formal framework for reaching agents' consensus on the terminology they use to communicate. In order to accomplish this, we adapt argument-based negotiation used in multi-agent systems to deal specifically with arguments that support or oppose candidate correspondences between ontologies. Each agent can decide according to its interests whether to accept or refuse the candidate correspondence. The proposed framework considers arguments and propositions that are specific to the matching task and related to the ontology semantics. This argumentation framework relies on a formal argument manipulation schema and on an encoding of the agents preferences between particular kinds of arguments. The former does not vary between agents, whereas the latter depends on the interests of each agent. Therefore, this work distinguishes clearly between the alignment rationales valid for all agents and those specific to a particular agent.

1 Introduction

When agents transfer information, they need a conceptualisation of the domain of interest and a shared vocabulary to communicate facts with respect to this domain. The conceptualisation can be expressed in a so-called *ontology*. An ontology abstracts the essence of the domain of interest and helps to catalogue and distinguish various types of objects in the domain, their properties and relationships (see, e.g. [14]). An agent can use such a vocabulary to express its beliefs and actions, and so communicate about them. Ontologies thus contribute to semantic interoperability when agents are embedded in open, dynamic environments, such as the Web, and its proposed extension the Semantic Web [7]. It has long been argued that in this type of environment there cannot be a single universal shared ontology, that is agreed upon by all the parties involved, as it would result in imposing a standard communication vocabulary. Interoperability therefore relies on the ability to reconcile different existing ontologies that may be heterogeneous in format and partially overlapping [22]. This reconciliation usually exists in the form of correspondences (or mapping) between agent ontologies and to use them in order to interpret or translate messages exchanged by agents. The underlying problem is usually termed an *ontology alignment* problem [13].

There are many matching algorithms able to produce such alignments [17]. In general, alignments can be be generated by trustable alignment services that can be invoked in order to obtain an alignment between two or more ontologies, and use it for translating messages [12]. Alternatively, they can be retrieved from libraries of alignments. However, the alignments provided by such services may not suit the needs of all agents. Indeed agents should be able to accept or refuse a proposed correspondence according to their own interests. In order to address this problem, we develop a formal framework for reaching agents consensus on the terminology they need to use in order to communicate. The framework allows agents to express their preferred choices over candidate correspondence. This is achieved adapting argument-based negotiation used in multi-agent systems 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 kind of mapping that can be supported by any one argument is clearly specified.

In order to compute preferred 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 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 an 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 defines the problem of reaching agreement over ontology alignments among agents. In section 3 we present in detail the argumentation framework and how it can be used. Section 4 defines the notion of agreeable alignments for two agents, and proposes a procedure to find these agreeable alignments. Next, in section 5, an example is provided to illustrate the idea. Section 6 points out some related work. Finally, section 7 draws some concluding remarks and identifies directions for further exploration.

2 Reaching agreement over ontology alignments

Before describing the framework, we first need to delimit the problem of reaching agreement over ontology alignments and state the assumptions upon which we build the theoretical framework.

In this paper, we concentrate on agents situated in a system, that need to display *social ability* and communicate in order to carry out some task. Each agent has a name, a role and a knowledge base. In some agent models, the basic knowledge base of an agent may be

[1] University of Liverpool,email: lori,valli,tbc@csc.liv.ac.uk
[2] INRIA Rhône-Alpes, email:Jerome.Euzenat@inrialpes.fr

consist of a set of beliefs, a set of desires and a set of intentions. However, for the purpose of this paper, we do not need to distinguish between beliefs, desire and intentions, and we will simply assume that an agent has a knowledge base where it stores facts about the domain it knows (which correspond to an ontology). Moreover, we do not make explicit use of the agent role.

Ontology can be defined as a tuple [11] $\langle C, H_C, R_C, H_R, I, R_I, A^O \rangle$, where the concepts C are arranged in a subsumption hierarchy H_C. Relations R_C is a set of relation between single concepts. Relations (or properties) can also be arranged in a hierarchy H_R. Instances I of a specific concept are interconnected by property instances R_I . Axioms A^O can be used to infer knowledge from that already existing. We further assume that ontologies are encoded in the same language, the standard OWL[3], removing us from the problem of integrating the ontology languages.

In order for agents to communicate, they need to establish alignments between their ontologies. We assume that such an alignment is generated by an alignment service agent and consists of a set of correspondences. A correspondence (or a mapping) can be described as a tuple: $\langle e, e', R \rangle$, where e and e' are the entities (concepts, relations or individuals) between which a relation is asserted by the correspondence; and R is the relation (e.g., equivalence, more general, etc.), holding between e and e', asserted by the correspondence [17]. For example a equivalence correspondence will stand between the concept 'car' in an ontology O and the concept 'automobile' in an ontology O' . A correspondence delivered by such an algorithm and not yet agreed by the agents will be called a *candidate mapping*. Note that we assume that an alignment service agent is able to generate an alignment using an independently defined decision-making process. We make no assumptions about how the agents achieve such decisions, as this is an internal agent process separate from the argumentation framework we present here.

Therefore, let two autonomous agents be committed to two ontologies O and O'. The *reaching agreement* problem is defined as follows:

Definition *"Find an agreement on the correspondences between the vocabularies they use, expressed as an ontology alignment."*.

Figure 1 illustrates the situation. Note that the definition consider two agents that want to communicate, but it can easily be extended to multi-agent systems. It is noteworthy that the process of reaching agreement should be as automatic as possible and should not require any feedback from human users. Indeed, essential to our approach, is that ontological discrepancies are treated at the level of agents themselves, without the aid of an external observer. The framework accounts for the detection and handling of ontological discrepancies by the agents themselves, on the basis of their own subjective view on the world. Agents should work towards agreement on the basis of their interest and preference states. We believe that this approach is both theoretically and practically important for agent systems.

In the next section, we show how this can be achieved using argumentation. Note that the framework requires that agents are able to justify why they have selected a particular mapping when challenged, since they will exchange arguments supplying the reasons for such a choice.

[3] http://www.w3.org/OWL/

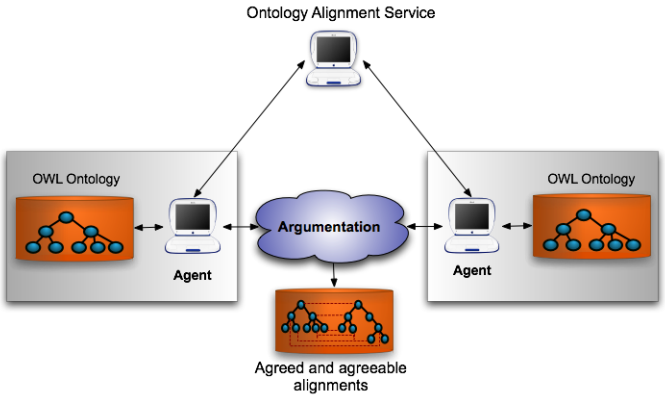

Figure 1. Reaching agreement over ontology alignments

3 Argumentation Framework

In order for the agents to consider potential mappings and the reasons for and against accepting them we use an argumentation framework. Our framework is based on the Value-based Argument Frameworks (VAFs) [5], a development of the classical argument systems of Dung [9]. We start with the presentation of Dung's framework, upon which the Value-based Argument Frameworks (VAFs) rely.

3.1 Classical argumentation framework

Definition 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 edges correspond to the elements of R. In Dung's work arguments are atomic and cannot be analysed further. In this paper, however, we are concerned only with arguments advocating mappings. We can therefore define arguments as follows:

Definition An argument $x \in AF$ is a triple $x = \langle G, m, \sigma \rangle$ where:

- m is a correspondence $\langle e, e', R \rangle$
- G is the grounds justifying a prima facie belief that the mapping does, or does not hold;
- σ is one of $\{+, -\}$ depending on whether the argument is that m does or does not hold.

When the set of such arguments and counter arguments have been produced, it is necessary to consider which of them should be accepted. Given an argument framework we can use definitions from [9] to define acceptability of an argument.

Definition Let $\langle AR, A \rangle$ be an argumentation framework. For R and S, subsets of AR, we say that:

- 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.

An argument x is *credulously accepted* if there is *some* preferred extension containing it; 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 reason giving arguments we need to be able to distinguish attacks from successful attacks, those which do defeat the attacked argument. One approach, taken in [1], is to rank arguments individually: an alternative, which we follow here, is to use a Value Based Argumentation framework (VAF) [5] which describes 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. VAFs are described in the next sub-section.

3.2 Value-based argumentation framework

We use the *Value-Based Argumentation Frameworks* (VAF) of Bench-Capon [5], to determine which mappings are acceptable, with respect to the different *audiences* represented by the different agents:

Definition 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 η: $AR \rightarrow \mathcal{V}$ is a mapping that associates a value $\eta(x) \in \mathcal{V}$ with each argument $x \in AR$

Definition 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, α is a *specific audience* (compatible with $\mathcal{R}$) if α 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 audiences (different agents) can have different perspectives on the same candidate mapping. [5] defines acceptability of an argument in the following way. Note that all these notions are now relative to some audience.

Definition 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 resoluble, and if it is, to determine the preferred extension with respect to a value ordering promoted by distinct audiences, [5] introduce the notion of objective and subjective acceptance as follows.

Definition *Subjective Acceptance.* Given an 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.

Definition *Objective Acceptance.* Given an VAF, $\langle AR, A, \mathcal{V}, \eta\rangle$, 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*. These definitions are particularly of interest in the case of the universal audience: subjective acceptability indicating that there is *at least one* specific audience (total ordering of values) under which x is accepted; objective acceptability that x must be accepted irrespective of the value ordering described by a specific audience; and, in contrast, x being indefensible indicating that no specific audience can ever accept x.

4 Arguing about correspondences

Our goal is to take advantage of value based argumentation so that agents can find the most mutually acceptable alignment. Section 4.1 defines the various categories of arguments that can support or attack mappings . Section 4.2 defines the notion of agreed and agreeable alignments for agents. Finally, in section 4.3 we demonstrate how the argumentation frameworks are constructed, in order to find such agreed and agreeable alignments.

4.1 Categories of arguments for correspondences

As we mentioned in Section 1, potential arguments are clearly identified and grounded on the underlying ontology languages, and the language of choice is the *de-facto* standard, OWL. Therefore, the grounds justifying correspondences can be extracted from the knowledge in ontologies. This knowledge includes both the extensional and intensional OWL ontology definitions. Our classification of the grounds justifying correspondences is the following:

semantic (M): the sets of models of some expressions do or do not compare;
internal structural (IS): the 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 of two entities with other entities do or do not compare;

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

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

These categories correspond to the type of categorizations underlying matching algorithms [22].

In our framework, we will use the types of arguments mentioned above as types for the value-based argumentation; hence $\mathcal{V} = \{M, IS, ES, T, E\}$. Therefore, 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 rather than in a lightweight ontology where there is very little reliable semantic information on which to base such arguments.

The reader may find it interesting to refer to the table 2, which summarises a number of reasons capable of justifying candidate OWL ontological alignments. Therefore, the table represents an (extensible) 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 different signs, 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 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. A list of arguments for or against accepting the mapping m, may be:

- The labels of the concept c and c' are synonymous. $\langle label(c) \approx label(c'), m, + \rangle$ (Terminological)
- Some of their instances are similar. $\langle E(c) \cap E(c') \neq \emptyset, m, + \rangle$ (Extensional)
- Some of their properties are similar. $\langle properties(c) \cap properties(c') \neq \emptyset, m, + \rangle$ (Internal Structural)
- Some of the super-classes of c and c' are dissimilar $\langle S(c) \cap S(c') = \emptyset, m, - \rangle$. (External Structural)

Similar arguments can be made for and against cases in which we consider properties or instances.

Therefore, in VAF 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.

4.2 Agreed and agreeable alignments

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 more 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 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. [8] describes a means by which agents may negotiate a joint preferred extension on the basis of their partitioned arguments so to maximise the number of desired arguments included while identifying which optional arguments need to be included to support them.

Based on these above considerations, we thus define an *agreed alignment* as the set of correspondences supported [4] by those arguments which are in every preferred extension of every agent, and an *agreeable alignment* extends the agreed alignment with the correspondences supported by arguments which are in some preferred extension of every agent. The next section shows how the argumentation frameworks are constructed.

4.3 Constructing argumentation frameworks

Given a single agent, we could 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. 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.

If we have multiple agents, we can 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. We employ the algorithm in [4] for computing the preferred extensions of a value-based argumentation framework given a value ordering. The global view is considered by taking the union of these preferred extensions for each audience. Then, 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 the mappings that have only arguments 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. Algorithm 1 shows how to find such agreed and agreeable 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, as suggested in [8] which defines a dialogue process for evaluating the status of arguments in a VAF, and shows how this process can be used to identify mutually acceptable arguments. In the course of constructing a position, an ordering of values best able to satisfy the joint interests of the agents concerned is determined.

[4] Note that a correspondence m is *supported* by an argument x if x is $\langle G, m, + \rangle$

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}

AG:=$\emptyset$
AG_{ext}:=$\emptyset$
for all audience $\mathcal{R}_i$ **do**
  for all VAF **do**
    compute the preferred extensions for $\mathcal{R}_i$, $P_j(\langle AR, A, \mathcal{V}, \eta\rangle, \mathcal{R}_i)$, $j \geq 1$
  end for
  $P_k(\mathcal{R}_i)$:=$\bigcup_j P_j(\langle AR, A, \mathcal{V}, \eta\rangle, \mathcal{R}_i)$, $k \geq 1$
end for
$AGArg$:=$x \in \bigcap_{k,i} P_k(\mathcal{R}_i), \forall k \geq 1, \forall i \geq 0$
for all $x \in AGArg$ **do**
  if x is $\langle G, m, +\rangle$ **then**
    $AG := AG \cup \{m\}$
  else
    reject mapping m
  end if
end for
if $\exists\ m \in M$ such that m is neither in AG and rejected **then**
  $AGArg_{ext}$:=$x \in \bigcap_i P_k(\mathcal{R}_i), \forall i \geq 0, k \geq 1$
  for all $x \in AGArg_{ext}$ **do**
    if x is $\langle G, m, +\rangle$ **then**
      $AG_{ext} := AG_{ext} \cup \{m\}$
    end if
  end for
end if

The above technique considers sets of mappings and complete argumentation frameworks. If instead the problem is to determine the acceptability of a single mapping it may be more efficient to proceed by means of a dialectical exchange, in which a mapping is proposed, challenged and defended. Argument protocols have been proposed in e.g. [15]. Particular dialogue games have been proposed based on Dung's Argumentation Frameworks (e.g. [10]), and on VAFs [6].

5 A walk through example

Having described the framework, we will go through an practical example.

Let us assume that some agents need to interact with each others using two independent but overlapping ontologies. One ontology is the bibliographic ontology[5] from the University of Canada, based on the bibTeX record. The other is the General University Ontology[6] from the French company Mondeca[7]. For space reasons, we will only consider a subset of these ontologies, shown in figure 2 and figure 3, where the first and second ontologies are represented by O_1 and O_2 respectively.

We will reason about the following candidate mappings:
m_1=$\langle O_1\colon Press, O_2\colon Periodical, \_, =\rangle$,
m_2=$\langle O_1\colon publication, O_2\colon Publication, \_, =\rangle$,
m_3=$\langle O_1\colon hasPublisher, O_2\colon publishedBy, \_, =\rangle$,

The following mappings are taken to be already accepted:

[5] http://www.cs.toronto.edu/semanticweb/maponto/ontologies/BibTex.owl
[6] http://www.mondeca.com/owl/moses/univ.owl
[7] Note that ontology O_2 has been slightly modified for the purposes of this example.

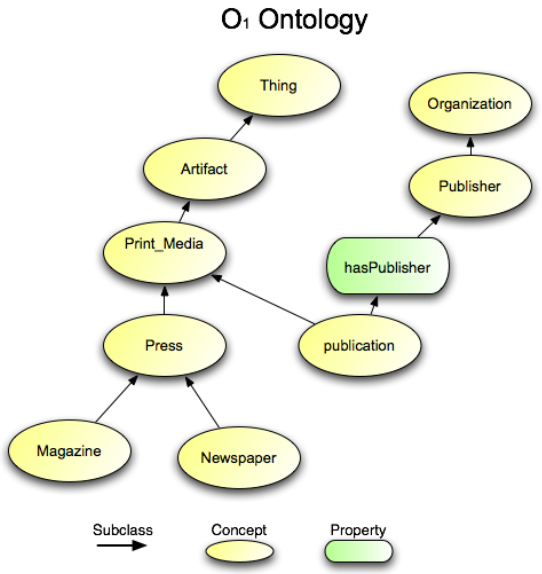

Figure 2. Ontology O_1

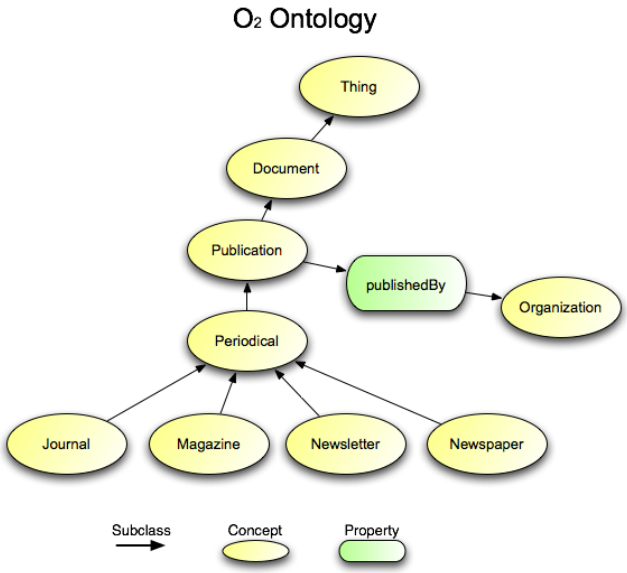

Figure 3. Ontology O_2

m_4=$\langle O_1\colon Magazine, O_2\colon Magazine, \_, =\rangle$,
m_5=$\langle O_1\colon Newspaper, O_2\colon Newspaper, \_, =\rangle$
m_6=$\langle O_1\colon Organization, O_2\colon Organization, \_, =\rangle$,

We begin by identifying a set of arguments and the attacks between them. This is achieved by instantiating the argumentation schemes, discussed previously, with respect to the interests of the agent. Table 1 shows each argument, labeled with an identifier, its type, and the attacks that can be made on it by opposing agents.

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 4, where nodes represent arguments, with the respective type value, and arcs represent the attacks. Now we can look in more detail at each argumentation framework.

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 similar 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 similar to two sub-concepts of O_2: $Periodical$, O_1: $Magazine$ and O_1: $Newspaper$, as established by m_4 and m_5. C pleads against m_1, because $Press$ and $Periodical$ do not have any lexical similarity.

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 similar (E). Argument F is based on the fact that O_1: $publication$ and O_2: $Publication$ have similar properties, O_1: $hasPublisher$ and O_1: $publishedBy$, as defined in m_3. F is then attacked by G, which states that the range of these properties, O_1: $Publisher$ and O_2: $Organization$, are not similar. This is in turn counter-attacked by the arguments H and I. The argument H states the mapping m_3 is correct, since O_1: $hasPublisher$ and O_1: $publishedBy$ are lexically similar. The argument I attacks the justification on 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: $Organization$.

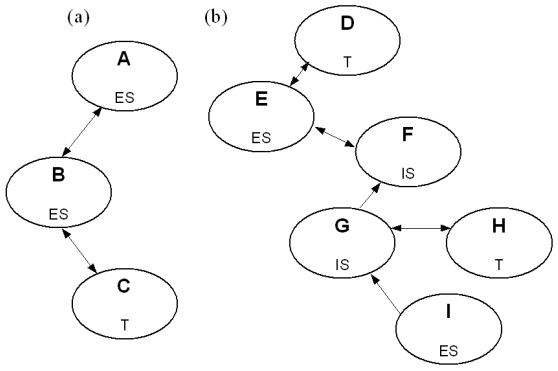

Figure 4. Value-Based Argumentation Frameworks

The above analysis gives different, but sometimes overlapping reasons to argue for and against several candidate mappings. 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 ($ES \succ_{\mathcal{R}_2} T$). For $\mathcal{R}_1$, we get two preferred extensions for the union of the argumentation frameworks $\{A, C, D, F, I, H\}$, and $\{A, C, D, E, I, H\}$, since E and F form a two cycle between types about which no preference has been expressed. For $\mathcal{R}_2$, however, the preferred extensions are $\{A, C, D, F, I, H\}$, $\{B, D, F, I, H\}$, $\{A, C, E, I, H\}$ and $\{B, E, I, H\}$, as there is a two cycle in ES which is no longer broken by C and no preference has been expressed between ES and IS. Therefore, the arguments that are accepted by both audiences are only $\{I, H\}$. 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 mapping m_2 will be accepted (it is accepted by $\mathcal{R}_1$ and acceptable to $\mathcal{R}_2$). m_3 will be accepted because H is agreed acceptable for these audiences. The *agreeable alignment* is then m_2 and m_3. Interestingly, in this scenario, should an agent wish to reject the mappings m_2 and m_3, it can achieve this by considering a new audience $\mathcal{R}_3$, in which internal structure is valued more then 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 $\{E, G, I\}$, 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 Related work

There are few approaches in the literature which have tackled the problem of agents negotiating about ontology alignments. An ontology mapping negotiation [19] has been proposed to establish a consensus between different agents which use the MAFRA alignment framework [20]. The approach is based on the utility and meta-utility functions used by the agents to establish if a mapping is accepted, rejected or negotiated. However, the approach is highly dependent on the use of the MAFRA framework and cannot be flexibly applied in other environments. [21] present an approach for agreeing on a common grounding ontology, in a decentralised way. Rather than being the goal of any one agent, the ontology mapping is a common goal for every agent in the system. [3] present an ontology negotiation protocol which enables agents to exchange parts of their ontology, by a process of successive interpretations, clarifications, and explanations. However, the end result of this process is that each agent will have the same ontology made of some sort of union of all the terms and their relations. In our context, agents keep their own ontologies, that they have been designed to reason with, while keeping track of the mappings with other agent's ontologies.

Unlike other approaches cited above, our work takes into consideration agents interests and preferences that may influence the selection of a given correspondence.

Contrastingly, significant research exists in the area of argumentation-based negotiation [18][16] in multi-agent systems. However, it has fundamentally remained at the level of a theoretical approach, and the few existing applications are concerned with legal cases and recently, in political decision-making [2].

7 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

Table 1. Arguments for and against the correspondences m_1, m_2 and m_3

SupC = super-classes, SubC = sub-classes, Pr = properties, Lb = label, Rg = Range, Sb = sibling-classes

Id	Argument	A	$\mathcal{V}$
A	$\langle SupC(Press) \cap SupC(Periodical) = \emptyset, m_1, - \rangle$	B	ES
B	$\langle SubC(Press) \cap SubC(Periodical) = \emptyset, m_1, + \rangle$	A,C	ES
C	$\langle Lb(Press) \not\approx Lb(Periodical), m_1, - \rangle$	B	T
D	$\langle Lb(publication) \approx Lb(Publication = \emptyset), m_2, + \rangle$	E	T
E	$\langle SupC(publication) \cap SupC(Publication), m_2, - \rangle$	D,F	ES
F	$\langle Pr(publication) \cap (Publication) \neq \emptyset, m_2, + \rangle$	E	IS
G	$\langle Rg(hasPublisher) \not\approx Rg(publishedBy), m_3, - \rangle$	F,H	IS
H	$\langle Lb(hasPublisher) \approx Lb(publishedBy), m_3, + \rangle$	G	T
I	$\langle SupC(Publisher) \cap (Organization \neq \emptyset), m_4, + \rangle$	G	ES

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 any correspondence, and enables consideration of an agent's interests and preferences that may influence the selection of a correspondence. An implementation of the framework is under development. Thus the effective results of an empirical evaluation are expected in the near future. Moreover, in future work we intend to investigate use of a negotiation process to enable agents to reach an agreement on a mapping when they differ in their ordering of argument types.

ACKNOWLEDGEMENTS

The research presented in this paper is partially supported under Knowledge Web European network of excellence (FP6-IST 2004-507482) and PIPS (FP6-IST 2004-507019). Special thanks to Floriana Grasso for her comments and suggestions.

REFERENCES

[1] L. Amgoud and C. Cayrol, 'On the Acceptability of Arguments in Preference-Based Argumentation', in *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*, eds., G. Cooper and S. Moral, (1998).

[2] K. Atkinson, T. J .M. Bench-Capon, and P. McBurney, 'Persuasive political argument', in *Proceedings of the Fifth International Workshop on Computational Models of Natural Argument (CMNA 2005)*, eds., C. Reed F. Grasso and R. Kibble, pp. 44–51, (2005).

[3] 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).

[4] T. Bench-Capon, 'Value based argumentation frameworks', in *Proceedings of Non Monotonic Reasoning*, pp. 444–453, (2002).

[5] T. Bench-Capon, 'Persuasion in Practical Argument Using Value-Based Argumentation Frameworks.', in *Journal of Logic and Computation*, volume 13, pp. 429–448, (2003).

[6] T. J .M. Bench-Capon, 'Agreeing to Differ: Modelling Persuasive Dialogue Between Parties Without a Consensus About Values', in *Informal Logic*, volume 22, pp. 231–245, (2002).

[7] T. Berners-Lee, J. Hendler, and O. Lassila, 'The Semantic Web', *Scientific American*, **284**(5), 34–43, (2001).

[8] S. Doutre, T.J.M. Bench-Capon, and P. E. Dunne, 'Determining Preferences through Argumentation', in *Proceedings of AI*IA'05*, pp. 98–109, (2005).

[9] P.H. Dung, 'On the Acceptability of Arguments and its Fundamental Role in Nonmonotonic Reasoning, Logic Programming and n-person Games', in *Artificial Intelligence*, volume 77, pp. 321–358, (1995).

[10] P. Dunne and T. J .M. Bench-Capon, 'Two Party Immediate Response Disputes: Properties and Efficiency', in *Artificial Intelligence*, volume 149, pp. 221–250, (2003).

[11] M. Ehrig and S. Staab, 'QOM - Quick Ontology Mapping', in *Proceedings of the International Semantic Web Conference*, (2004).

[12] J. Euzenat, 'Alignment infrastructure for ontology mediation and other applications', in *Proceedings of the First International workshop on Mediation in semantic web services*, eds., Martin Hepp, Axel Polleres, Frank van Harmelen, and Michael Genesereth, pp. 81–95, (2005).

[13] J. Euzenat and P. Valtchev, 'Similarity-based ontology alignment in OWL-Lite', in *Proceedings of European Conference on Artificial Intelligence (ECAI04)*, (2004).

[14] T. R. Gruber, 'A Translation Approach to Portable Ontology Specifications', *Knowledge Acquisition*, **5**(2), 199–220, (1993).

[15] P. McBurney and S. Parsons, 'Locutions for Argumentation in Agent Interaction Protocols', in *Proceedings of International Workshop on Agent Communication, New-York (NY US)*, pp. 209–225, (2004).

[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, pp. 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] 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).

[19] N. Silva, P. Maio, and J. Rocha, 'An Approach to Ontology Mapping Negotiation', in *Proceedings of the Workshop on Integrating Ontologies*, (2005).

[20] N. Silva and J. Rocha, 'MAFRA Semantic Web Ontology MApping FRAmework', in *Proceedings of the Seventh Multi-Conference on Systemics, Cybernetics and Informatics*, (2003).

[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.R.S. Visser, D.M. Jones, T.J.M. Bench-Capon, and M.J.R. Shave, 'Assessing Heterogeneity by Classifying Ontology Mismatches', in *Proceedings of the FOIS'98*, ed., N. Guarino, (1998).

Table 2. Argument scheme for OWL ontological alignments

Mapping	σ	**Grounds**	**Comment**
$\langle e, e', \sqsubseteq \rangle$	+	$S(e) \subseteq S(e')$	(some or all) neighbours (e.g., super-entities, sibling-entities, etc.) of e are similar in those of e'
$\langle e, e', \sqsubseteq \rangle$	-	$S(e') \subseteq S(e)$	no neighbours of e are similar in those of e'
$\langle e, e', \sqsubseteq \rangle$	-	$S(e') \subseteq S(e)$	(some or all)neighbours of e' are similar in those of e
$\langle e, e', \equiv \rangle$	+	$S(e) \cap S(e') \neq \emptyset$	Entities have similar neighbours (e.g., super-entities, sibling-entities, etc.)
$\langle e, e', \equiv \rangle$	-	$S(e) \cap S(e') = \emptyset$	Entities does not have similar neighbours
$\langle c, c', \sqsubseteq \rangle$	+	$properties(c) \subseteq properties(c')$	(some or all) properties of c are similar in those of c'
$\langle c, c', \sqsubseteq \rangle$	-	$properties(c') \not\subseteq properties(c)$	no properties of c are similar in those of c'
$\langle c, c', \sqsubseteq \rangle$	-	$properties(c') \subseteq properties(c)$	(some or all) properties of c' are included in those of c
$\langle c, c', \equiv \rangle$	+	$properties(c) \cap properties(c') \neq \emptyset$	the concepts c and c' have common properties
$\langle c, c', \equiv \rangle$	-	$properties(c) \cap properties(c') = \emptyset$	no properties in c and c' are similar
$\langle p, p', \equiv \rangle$ $\langle p, p', \sqsubseteq \rangle$	+	$I(p) \approx I(p')$	Properties have similar structure (e.g., range, domain or cardinality)
$\langle p, p', \equiv \rangle$ $\langle p, p', \sqsubseteq \rangle$	-	$I(p) \not\approx I(p')$	Properties do not have similar structure
$\langle i, i', \equiv \rangle$ $\langle i, i', \sqsubseteq \rangle$	+	$properties(i, i'') \approx properties(i', i'')$	Each individual i and i' referees to a third instance i" via similar properties
$\langle p, p', \equiv \rangle$ $\langle p, p', \sqsubseteq \rangle$	-	$properties(i, i'') \not\approx properties(i', i'')$	The properties that link each individual i and i' to a third instance i" are dissimilar
$\langle e, e', \sqsubseteq \rangle$	+	$E(e) \subseteq E(e')$	(some or all) instances of e are similar in those of e'
$\langle e, e', \sqsubseteq \rangle$	-	$E(e) \not\subseteq E(e')$	no instances of e are similar in those of e'
$\langle e, e', \sqsubseteq \rangle$	-	$E(e') \subseteq E(e)$	(some or all) instances of e' are similar in those of e
$\langle e, e', \equiv \rangle$	+	$E(e) \cap E(e') \neq \emptyset$	e instances are similar in those of e' and/or vice versa.
$\langle e, e', \equiv \rangle$	-	$E(e) \cap E(e') = \emptyset$	Entities e and e' does not have common instances
$\langle e, e', \equiv \rangle$ $\langle e, e', \sqsubseteq \rangle$	+	$label(e) \approx label(e')$	Entities's labels are similar (e.g., synonyms and lexical variants)
$\langle e, e', \equiv \rangle$ $\langle e, e', \sqsubseteq \rangle$	-	$label(e) \not\approx label(e')$	Entities' labels are dissimilar (e.g., homonyms)
$\langle e, e', \equiv \rangle$ $\langle e, e', \sqsubseteq \rangle$	+	$URI(e) \approx URI(e')$	Entities' URIs are similar
$\langle e, e', \equiv \rangle$ $\langle e, e', \sqsubseteq \rangle$	-	$URI(e) \not\approx URI(e')$	Entities' URIs are dissimilar