

Reaching Agreement over Ontology Alignments

Loredana Laera¹, Valentina Tamma¹, Jérôme Euzenat²,
Trevor Bench-Capon¹, and Terry Payne³

¹ Department of Computer Science, University of Liverpool, UK
{lori, valli, tbc}@csc.liv.ac.uk

² INRIA Rhône-Alpes, Montbonnot, France
Jerome.Euzenat@inrialpes.fr

³ Department of Electronics and Computer Science, University of Southampton, UK
trp@ecs.soton.ac.uk

Abstract. When agents communicate, they do not necessarily use the same vocabulary or ontology. For them to interact successfully, they must find correspondences (mappings) between the terms used in their respective 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 thus may necessitate additional negotiation to identify a mutually agreeable set of correspondences.

We propose an approach for supporting the creation and exchange of different arguments, that support or reject possible correspondences. Each agent can decide, according to its preferences, whether to accept or refuse a candidate correspondence. The proposed framework considers arguments and propositions that are specific to the matching task and are based on 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. Whilst the former does not vary between agents, the latter depends on the interests of each agent. Thus, this approach distinguishes clearly between alignment rationales which are valid for all agents and those specific to a particular agent.

1 Introduction

Ontologies play an important role in inter-agent communication, by providing the definitions of the vocabularies used by agents to describe the world [11]. An agent can use such a vocabulary to express its beliefs and actions, and so communicate about them. Ontologies contribute to semantic interoperability when agents are embedded in open, dynamic environments, such as the Web and its proposed extension, the Semantic Web [4]. However, in this type of environment there cannot be a single universally shared ontology that is agreed upon by all the parties involved, as this would result in imposing a standard communication vocabulary. Instead, every agent will typically use its own private ontology, which may not be understandable by other agents. Interoperability therefore relies on the ability to reconcile different existing ontologies that may be heterogeneous in format and with partially overlapping domains [19]. This reconciliation usually relies on the existence of correspondences (or mappings) between agent ontologies, and using them in order to interpret or translate messages exchanged by agents. The underlying problem is usually termed an *ontology matching* problem [10].

There are many matching algorithms that are able to produce such alignments [14]. In general, alignments can be generated by independent, trustable alignment services that can be invoked in order to obtain an alignment between two ontologies, and use it for translating messages [9]. Alternatively, they can be retrieved from libraries of alignments. However, in an open environment where autonomous agents try to pursue their own objectives, the acceptability of a partial alignment provided by such services cannot be taken for granted. For a given context, agents might have different and inconsistent perspectives; *i.e.* *interests* and *preferences*, on the acceptability or not of a candidate mapping, each of which may be rationally acceptable. This may be due, for instance, 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 exist within an open environment, and thus have no prior knowledge of either the existence or constraints of other agents. These constraints are also relevant in Semantic Web Service applications, where services performing the same tasks may advertise their capabilities differently, or where service requests, and service offers may be expressed by using different ontologies, and thus need to be reconciled dynamically at run time.

In order to address this problem, we present a framework that allows agents to reach a consensus on the terminology they use in order to communicate. The framework we present was primarily motivated by open-agent environments, and although in the remainder of the paper we refer to agents, the framework can equally be applied to semantic web services. The framework allows agents to express their preferred choices over candidate correspondences. This is achieved by 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 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 [3] 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 defines the assumptions underlining the framework. In Section 3 we present in detail the argumentation framework and how it can be used. Section 4 defines the various categories of arguments that can support or attack mappings. Section 5 defines the notion of agreed and agreeable alignments for agents, and a procedure to find them is proposed in Section 6. Next, in section 7, an example is provided to illustrate the argumentation process. Section 8

presents some related work, and finally, Section 9 draws some concluding remarks and identifies directions for further exploration.

2 Assumptions Underlining the Framework

In this paper, we focus on autonomous agents situated in an open system. Each agent has a name and a knowledge base, expressed using an ontology. Moreover, we assume that the *mental attitudes* of an agent towards correspondences are represented in terms of *interests* and *preferences*. These represent the motivations of the agent, that determine whatever a mapping is accepted or rejected. Indeed, each agent has a (partial or total) pre-ordering of preferences over different types of ontology mismatches (*Pref*).

In order for agents to communicate, they need to establish alignments between their ontologies¹. We assume that potential alignments are generated by a dedicated agent, called an *Ontology Alignment Service (OAS)* [10]. The alignment provided will consist 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 [14]. There are a number of approaches that an Ontology Alignment Service can use for deriving such correspondences. A correspondence which has been provided by an OAS, but for which no agreement has been made by the agents will be called a *candidate mapping*. Moreover, we assume that for each correspondence m , an OAS is able to provide a set of justifications G , that explain why it has generated a candidate mapping. Agents will use such information to exchange arguments supplying the reasons for their mapping choices. In addition, every agent has a private threshold value ϵ which will be compared to the degree of confidence that an OAS associates with each mapping. Although few approaches for ontology alignment provide such justification [15] [5], tools such as [8] combine different similarity metrics, and these measures can be used to extend the system and provide the required justifications.

An agent will apply its pre-ordering of preferences and threshold, ϵ , when generating the arguments for and against a candidate mapping. Furthermore, we note that the process of reaching agreement should be as automatic as possible and should not require any involvement from human users.

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 Value-based Argument Frameworks (*VAFs*) [3]. This work is an experimental research and a prototype of the framework is under development. We start with the presentation of Dung work [7], upon which the *VAFs* rely.

¹ Although the agents' ontologies may differ, we assume that ontologies are encoded in the same language, the standard OWL (<http://www.w3.org/OWL/>), thus eliminating the problem of integrating different ontology languages.

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 Dung’s work, arguments are atomic and cannot be analysed further. In this paper, however, 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 does not hold; σ 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. Given an argument framework we can use definitions from [7] to define acceptability of an argument.

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 4, 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 $\mathcal{R} \subseteq \mathcal{V} \times \mathcal{V}$ 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 agents (represented by different audiences) can have different perspectives on the same candidate mapping. Acceptability of an argument is defined in the following way:²

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

- 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: $(x, y) \in A$ and it is not the case that $\eta(y) \succ_{\mathcal{R}} \eta(x)$.
- 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 $\mathcal{R}$.
- 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 $(x, y) \notin A$ or $\eta(y) \succ_{\mathcal{R}} \eta(x)$.
- 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 $\mathcal{R}$.
- A subset S is a preferred extension for the audience $\mathcal{R}$ if it is a maximal admissible set with respect to $\mathcal{R}$.
- 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, [3] introduces the notion of objective and subjective acceptance as follows:

² Note that all these notions are now relative to some audience.

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.*

4 Categories of Arguments for Correspondences

As we mentioned in Section 2, 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. 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 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.

These categories correspond to the type of categorizations underlying ontology matching algorithms [19].

In our framework, we will use the types of arguments described above as types for the VAF; hence $\mathcal{V} = \{M, IS, ES, T, E\}$. 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, compared to that in a lightweight ontology where there is very little reliable semantic information on which to base such arguments.

Table 1 summarises a number of reasons capable of justifying candidate OWL ontology 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 conflicting values of σ , 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.

Therefore, 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, as

mentioned in section 2, 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 ε , an agent will generate arguments $x = \langle G, m, + \rangle$. If not, the agent will generate arguments against: $x = \langle G, m, - \rangle$. The generation is achieved by instantiating the argumentation schema.

Table 1. Argument scheme for OWL ontological alignments

Mapping	σ	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'
$\langle e, e', n, \equiv \rangle$	-	$\nexists m_i = \langle ES(e), ES(e'), n', \equiv \rangle$	No neighbours of e and e' are mapped
$\langle e, e', n, \sqsubseteq \rangle$	-	$\nexists m_i = \langle ES(e), ES(e'), n', \equiv \rangle$	No neighbours of e are mapped to those of e'
$\langle e, e', n, \sqsubseteq \rangle$	-	$\exists m_i = \langle ES(e'), ES(e), n', \equiv \rangle$	(some or all) Neighbours of e' are mapped to 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'
$\langle c, c', n, \sqsubseteq \rangle$	-	$\nexists m_i = \langle IS(c), IS(c'), n', \equiv \rangle$	No properties of c are mapped to those of c'
$\langle c, c', n, \sqsubseteq \rangle$	-	$\exists m_i = \langle IS(c'), IS(c), n', \equiv \rangle$	(some or all) Properties of c' are mapped to those of c
$\langle c, c', n, \equiv \rangle$	+	$\exists m_i = \langle IS(c'), IS(c), n', \equiv \rangle$	The concepts c and c' have mapped properties
$\langle c, c', n, \equiv \rangle$	-	$\nexists m_i = \langle IS(c'), IS(c), n', \equiv \rangle$	No properties in c and c' are mapped
$\langle p, p', n, \equiv \rangle$	+	$\exists m_i = \langle IS(p), IS(p), n', \equiv \rangle$	The range and/or the domain of the property p is mapped with those of p'
$\langle p, p', n, \equiv \rangle$	-	$\nexists m_i = \langle IS(p), IS(p), n', \equiv \rangle$	The range and/or the domain of the properties p and p' are not mapped
$\langle p, p', n, \sqsubseteq \rangle$	-	$\nexists m_i = \langle IS(p), IS(p), n', \equiv \rangle$	The range and/or the domain of the properties p and p' are not mapped
$\langle i, i', n, \equiv \rangle$	+	$\exists m_i = \langle IS(i, i''), IS(i', i''), n', \equiv \rangle$	Each individual i and i' refers to a third instance i'' via two properties that are mapped
$\langle p, p', n, \equiv \rangle$	-	$\nexists m_i = \langle IS(i, i''), IS(i', i''), n', \equiv \rangle$	The properties that link each individual i and i' to a third instance i'' are not mapped
$\langle e, e', n, \equiv \rangle$	+	$\exists m_i = \langle E(e), E(e'), n', \equiv \rangle$	(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'
$\langle e, e', n, \equiv \rangle$	-	$\nexists m_i = \langle E(e), E(e'), n', \equiv \rangle$	No instances of e and e' are mapped
$\langle e, e', n, \sqsubseteq \rangle$	-	$\nexists m_i = \langle E(e), E(e'), n', \equiv \rangle$	No instances of e are mapped to those of e'
$\langle e, e', n, \sqsubseteq \rangle$	-	$\exists m_i = \langle E(e'), E(e), n', \equiv \rangle$	(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)
$\langle e, e', n, \equiv \rangle$	+	$URI(e) \approx_T URI(e')$	Entities' URIs share lexical features
$\langle e, e', n, \sqsubseteq \rangle$	-	$URI(e) \not\approx_T URI(e')$	Entities' URIs do not share lexical features

5 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 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. Doutre et al. [6] describe a means by which agents may negotiate a joint preferred extension on the basis of their partitioned arguments so as to maximise the number of desired arguments included, whilst identifying which optional arguments need to be included to support them.

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 supported³ 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.

6 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.

³ Note that a correspondence m is *supported* by an argument x if x is $\langle G, m, + \rangle$.

2. Given multiple agents, we simply merge their individual frameworks by forming the union of their individual argument sets and attack relations, and then extend the attack relations by computing the attacks between the arguments present in the framework of each agent with the arguments of all the other agents.
3. Then, for each VAF , we determine which of the arguments are undefeated by attacks from other arguments. We employ the algorithm in [2] 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.
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 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. An algorithm to find such agreed and agreeable alignments is available in Laera et al. [12].

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 [6] 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. However, such issues are the subject of ongoing research.

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.

7 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 ontology⁴ from the University of Toronto, based on bibTeX; whereas the second agent, Ag_2 , uses the General University Ontology⁵ from the French company Mondeca⁶. For space reasons, we will 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.

⁴ <http://www.cs.toronto.edu/semanticweb/maponto/ontologies/BibTex.owl>

⁵ <http://www.mondeca.com/owl/mones/univ.owl>

⁶ Note that ontology O_2 has been slightly modified for the purposes of this example.

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 will reason about the following candidate mappings, provided by the OAS:

$$\begin{aligned}
m_1 &= \langle O_1: Press, O_2: Periodical, n, = \rangle;^7 \\
m_2 &= \langle O_1: publication, O_2: Publication, n, = \rangle; \\
m_3 &= \langle O_1: hasPublisher, O_2: publishedBy, n, = \rangle; \\
m_4 &= \langle O_1: Magazine, O_2: Magazine, n, = \rangle; \\
m_5 &= \langle O_1: Newspaper, O_2: Newspaper, n, = \rangle; \\
m_6 &= \langle O_1: Organization, O_2: Organization, n, = \rangle.
\end{aligned}$$

As mentioned in Section 2, the generation of the arguments and counter-arguments is based on the agent's preferences and threshold. However, here we assume that all above candidate mappings have 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.

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, given in table 1, with respect to the interests and preferences $Pref$ of the agents and taking into consideration the justifications G , provided by the OAS. 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 1, 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 $O_2: Periodical$.


⁷ 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	$\langle \exists m = \langle \text{superconcept}(\text{Press}), \text{superconcept}(\text{Periodical}), n, \equiv, \rangle, m_1, - \rangle$	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	$\langle \exists m = \langle \text{superconcept}(\text{publication}), \text{superconcept}(\text{Publication}), n, \equiv, \rangle, m_2, - \rangle$	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	$\langle \exists m = \langle \text{range}(\text{hasPublisher}), \text{range}(\text{publishedBy}), n, \equiv, \rangle, m_3, - \rangle$	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

- *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.

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. 1.** 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 . The argument H states that the mapping m_3 is correct, since O_1 : *hasPublisher* and O_2 : *publishedBy* are lexically similar. The argument 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*. The argument 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 two audiences, $\mathcal{R}_1$ and $\mathcal{R}_2$, the preferred extensions for the union of the argumentation frameworks (a) and (b) is shown in Table 4.

Table 4. Preferred Extensions

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 that are 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 be acceptable too, because the arguments supporting it are in some preferred extension for these audiences, as defined in section 5. The *agreed alignment* is then m_3, m_4, m_5 and m_6 , while the *agreeable alignment* adds m_2 . 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 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.

8 Related Work

There are few approaches in the literature which have addressed the use of argumentation or negotiation between agents w.r.t. ontology alignments. An ontology mapping negotiation [17] has been proposed to establish a consensus between different agents which use the MAFRA alignment framework. 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. van Diggelen et al. [18] 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. Bailin and Truszkowski [1] 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 converge on a single, shared ontology consisting of the 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. Contrastingly, significant research exists in the area of argumentation-based negotiation [16][13] in multi-agent systems. However, none has been apply in area of ontology alignments.

9 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 more sound and effective mutual understanding and communicative work in agents system.

In the current state of the implementation, the ontology alignments is provided manually. The next step is to extend the developed prototype to utilize an ontology alignment services in order to obtain the alignment automatically. An empirical evaluation is planned. Moreover, 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.

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.

References

1. S. C. Bailin and W. Truszkowski. Ontology Negotiation: How Agents Can Really Get to Know Each Other. In *Proceedings of the WRAC 2002*, 2002.
2. T. Bench-Capon. Value based argumentation frameworks. In *Proceedings of Non Monotonic Reasoning*, pages 444–453, 2002.
3. T. Bench-Capon. Persuasion in Practical Argument Using Value-Based Argumentation Frameworks. In *Journal of Logic and Computation*, volume 13, pages 429–448, 2003.
4. T. Berners-Lee, J. Hendler, and O. Lassila. The Semantic Web. *Scientific American*, 284(5):34–43, 2001.
5. R. Dhamankar, Y. Lee, A. Doan, A. Halevy, and P. Domingos. iMAP: Discovering complex semantic matches between database schemas. In *Proceedings of the International Conference on Management of Data (SIGMOD)*, pages 383–394, 2004.
6. S. Doutre, T. Bench-Capon, and P. E. Dunne. Determining Preferences through Argumentation. In *Proceedings of AI*IA'05*, pages 98–109, 2005.
7. 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.
8. M. Ehrig and S. Staab. QOM - Quick Ontology Mapping. In *Proceedings of the International Semantic Web Conference*, 2004.
9. J. Euzenat. Alignment infrastructure for ontology mediation and other applications. In M. Hepp, A. Polleres, F. van Harmelen, and M. Genesereth, editors, *Proceedings of the First International workshop on Mediation in semantic web services*, pages 81–95, 2005.
10. J. Euzenat and P. Valtchev. Similarity-based ontology alignment in OWL-Lite. In *Proceedings of the European Conference on Artificial Intelligence (ECAI 2006)*, 2004.
11. T. R. Gruber. A Translation Approach to Portable Ontology Specifications. *Knowledge Acquisition*, 5(2):199–220, 1993.
12. L. Laera, V. Tamma, T. Bench-Capon, and J. Euzenat. Agent-based Argumentation for Ontology Alignments. In *Proceedings of the Workshop on Computational Models of Natural Argument (CMNA 2006)*, 2006.
13. 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.
14. P. Shvaiko and J. Euzenat. A survey of schema-based matching approaches. *Journal on data semantics*, 4:146–171, 2005.
15. 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.
16. 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.
17. N. Silva, P. Maio, and J. Rocha. An Approach to Ontology Mapping Negotiation. In *Proceedings of the Workshop on Integrating Ontologies*, 2005.
18. 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.
19. 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.