

Arguing over ontology alignments

L. Laera¹, V. Tamma¹, J. Euzenat², T. Bench-Capon¹, and T. Payne³

¹ Department of Computer Science, University of Liverpool, UK

² INRIA Rhône-Alpes, Montbonnot, France

³ Computer Science, University of Southampton, UK

Abstract. In open and dynamic environments, agents will usually differ in the domain ontologies they commit to and their perception of the world. The availability of Alignment Services, that are able to provide correspondences between two ontologies, is only a partial solution to achieving interoperability between agents, because any given candidate set of alignments is only suitable in certain contexts. For a given context, different agents might have different and inconsistent perspectives that reflect their differing interests and preferences on the acceptability of candidate mappings, each of which may be rationally acceptable. In this paper we introduce an argumentation-based negotiation framework over the terminology they use in order to communicate. 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. Thus, this approach distinguishes clearly between the alignment rationales valid for all agents and those specific to a particular agent.

1 Introduction

Traditionally ontologies have been used to achieve semantic interoperability between software applications, as such applications provide the definitions of the vocabularies they use to describe the world [12], and they have proved especially effective when systems are embedded in open, dynamic environments, such as the Web and the Semantic Web [4]. Interoperability relies on the ability to reconcile the differences between heterogeneous ontologies [18]. This reconciliation usually relies on the existence of correspondences (or mappings) between different ontologies (*ontology alignment* [11]), and uses them in order to interpret or translate messages exchanged by applications. Such correspondences may be generated by a variety of different matching algorithms [16]⁴, and their production usually requires several steps. These can include the definition of an initial alignment, or the training over some examples, and these invariably involve some form of interpretation of preliminary results [10]. Therefore, approaches to ontology alignment can only be effective when used to support semantic interoperation at *design time* in closed or partially open environments, where the actors involved are often known, where ontology changes are controlled and thus the alignments can be established before the systems interact. However, these approaches are not sufficient to support semantic interoperation in open environments, where systems can dynamically join or leave and no prior assumption can be made on the ontologies to align. In

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

such environments, the different systems involved need to agree on the semantics of the terms used during the interoperation, and reaching this agreement can only come through some sort of negotiation process [1].

This paper extends the notion of reaching agreement through automated negotiation (*i.e.* without human intervention) by considering the type of systems that need to interoperate, which can affect how the negotiation should proceed. Specifically, *autonomous agents* (within an open environment) may perform different tasks depending on their state and the service providers they interact with. Thus, such agents will differ in the domain ontologies they commit to [12]; and their perception of the world (and hence the choice of vocabulary used to represent concepts). Imposing a single, universally shared ontology on agents is not only impractical because it would result in assuming a standard communication vocabulary (and thus violate the dynamics of open environments) but it also does not take into account the conceptual requirements of services that could appear in future. Instead, every agent assumes its own heterogeneous private ontology, which may not be understandable by other agents. The availability of *Alignment Services* that are able to provide correspondences between two ontologies is only the beginning of a solution to achieving interoperability between agents, as any given candidate set of alignments is only suitable in certain contexts. For a given context, agents might have different and inconsistent perspectives; *i.e.* *interests* and *preferences*, on the acceptability of a candidate mapping, each of which may be rationally acceptable. This may be due to the subjective nature of ontologies, to the context and the requirement of the alignments and so on. For example, an agent may be interested in accepting only those mappings that have linguistic similarities, since its ontology is too *structurally simple* to realise any other type of mismatch. In addition, any decision on the acceptability of these mappings has to be made dynamically (at run time), due to the fact that the agents have no prior knowledge of either the existence or constraints of other agents.

In order to address this problem, we present a framework to support agents to negotiate agreement on the terminology they use in order to communicate, by allowing them to express their preferred choices over candidate correspondences. This is achieved by adapting argument-based negotiation to deal specifically with arguments that support or oppose the proposed correspondences between ontologies. The set of potential arguments are clearly identified and grounded on the underlying ontology languages, and the kinds of mapping that can be supported by any such argument are clearly specified. Specifically, we use a value-based argumentation framework [2], allowing each agent to express its preferences between the categories of arguments that are clearly identified in the context of ontology alignment. Our approach is able to give a formal motivation for the selection of any correspondence, and enables consideration of an agents' interests and preferences that may influence the selection of a given correspondence. Therefore, this work provides a concrete instantiation of the "meaning negotiation" process that we would like agents to achieve. Moreover, in contrast to current ontology matching procedures, the choice of alignment is based on two clearly identified elements: (i) the argumentation framework, which is common to all agents, and (ii) the preference relations which are private to each agent.

The remainder of this paper is structured as follows. Section 2 presents the argumentation framework and how it can be used. Section 3 defines the various categories of arguments that can support or attack mappings. Section 4 describes our agent model and discusses how agents should reach agreement. An example illustrating the argumentation process is given in Section 5, followed concluding remarks in Section 6⁵.

2 Argumentation Framework

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

For agents to communicate, they first need to establish a mutually acceptable set of alignments between their ontologies. Potential alignments are generated at design time (by a variety of different ontology-matching approaches [16]), and provided at run-time by a dedicated agent, called an *Ontology Alignment Service (OAS)* (Figure 1). An alignment consists of a set of 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 [16]. A *candidate mapping* is a correspondence (provided by an OAS) that could be used by the agents to align their ontologies. Each correspondence m is accompanied by a set of justifications G , which provide an explanation as to why the correspondence was generated⁷. This information is used by the agents when generating and exchanging arguments, for and against a candidate mapping. In addition, every agent has a private threshold value ε which will be compared to the degree of confidence, n , of a mapping, to decide whether it should be considered.

In order for the agents to consider potential mappings and the reasons for and against accepting them, we use an argumentation framework based on *Value-based Argument Frameworks (VAFs)* [2], that extends Dong’s classical argument system [7]⁸.

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 .

⁵ A survey of related work is given in an extended version of this paper [13].

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

⁷ Although few approaches for ontology alignment provide justifications [17, 5], tools such as [9] combine different similarity metrics which can be used to provide necessary justifications.

⁸ More details can be found in an extended version of this paper [13].

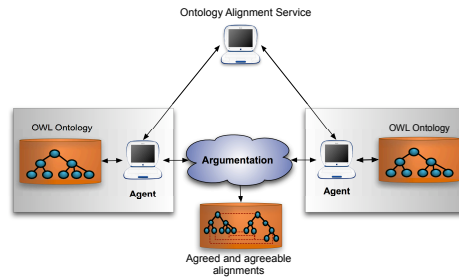


Fig. 1. Reaching agreement over ontology alignments

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

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

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 order to take into account that, for a given situation, agents might have different point of view, we are concerned by a set of audiences, which adhere to different argu-

ment with a different strengths. Therefore we use a Value-based Argumentation Framework, which prescribes different strengths to arguments on the basis of the values they promote and the ranking given to these values by the audience for the argument. This allows us to systematically relate strengths of arguments to their motivations, and to accommodate different audiences with different interests and preferences.

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

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

Definition 6. An audience for a VAF is a binary relation $\mathcal{R} \subset \mathcal{V} \times \mathcal{V}$ whose (irreflexive) transitive closure, \mathcal{R}^* , is asymmetric, i.e. at most one of (v, v') , (v', v) are members of \mathcal{R}^* for any distinct $v, v' \in \mathcal{V}$. We say that v_i is preferred to v_j in the audience \mathcal{R} , denoted $v_i \succ_{\mathcal{R}} v_j$, if $(v_i, v_j) \in \mathcal{R}^*$.

Let \mathcal{R} be an audience, 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:⁹

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

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

In order to determine whether the dispute is resolvable, and if it is, to determine the preferred extension with respect to a value ordering promoted by distinct audiences, [2] introduces the notion of objective and subjective acceptance as follows:

⁹ 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.

Next, we define the various types of arguments that can be distinguished for supporting or attacking correspondences.

3 Arguments for Correspondences

Potential arguments are clearly identified and grounded on the underlying ontology language OWL. Therefore, the grounds justifying correspondences can be extracted from the knowledge in ontologies¹⁰. 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 [18]. In our framework, we will use the types of arguments described above as types for the VAF ; hence $\mathcal{V} = \{M, IS, ES, T, E\}$. For example, an audience may specify that terminological arguments are preferred to semantic arguments, or vice versa. Note that this may vary according to the nature of the ontologies being aligned. Semantic arguments will be given more weight in a fully axiomatised ontology, compared to that in a lightweight ontology where there is very little reliable semantic information on which to base such arguments.

Table 1 presents a sample set of argument schemes, instantiations of which will comprise AR . Attacks between these arguments will arise when we have arguments for the same mapping but with conflicting values of 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.

In $VAFs$, arguments against or in favour of a candidate mapping are seen as grounded on their type. In this way, we are able to motivate the choice between preferred extensions by reference to the type ordering of the audience concerned. Moreover, the

¹⁰ This knowledge includes both the extensional and intensional OWL ontology definitions.

pre-ordering of preferences $Pref$ for each agent will be over \mathcal{V} , that corresponds to the determination of an audience.

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'
$\langle c, c', n, \sqsubseteq \rangle$	+	$\exists m_i = \langle IS(c), IS(c'), n', \equiv \rangle$	(some or all) Properties of concept c are mapped to those of concept c'
$\langle c, c', n, \sqsupseteq \rangle$	-	$\exists m_i = \langle IS(c), IS(c'), n', \equiv \rangle$	No properties of c are mapped to those of c'
$\langle e, e', n, \equiv \rangle$	+	$\exists m_i = \langle E(e), E(e'), n', \equiv \rangle$	(some or all) Instances of e and e' are mapped
$\langle e, e', n, \sqsubseteq \rangle$	+	$\exists m_i = \langle E(e), E(e'), n', \equiv \rangle$	(some or all) Instances of e are mapped to those of e'
$\langle e, e', n, \equiv \rangle$	+	$label(e) \approx_T label(e')$	Entities' labels share lexical features (e.g., synonyms and lexical variants)
$\langle e, e', n, \sqsubseteq \rangle$	-	$label(e) \not\approx_T label(e')$	Entities' labels do not share lexical features (e.g., homonyms)

Although in $VAFs$ there is always a unique non-empty preferred extension with respect to a specific audience, provided the AF does not contain any cycles in a single argument type, an agent may have multiple preferred extensions either because no preference between two values in a cycle has been expressed, or because a cycle in a single value exists. The first may be eliminated by committing to a specific audience, but the second cannot be eliminated in this way. In our domain, where many attacks are symmetric, two cycles will be frequent and in general an audience may have multiple preferred extensions.

Thus, given a set of arguments justifying mappings organised into an argumentation framework, an agent will be able to determine which mappings are acceptable by computing the preferred extensions with respect to its preferences. If there are multiple preferred extensions, the agent must commit to the arguments present in all preferred extensions, but it has some freedom of choice with respect to those in some but not all of them.

Based on the above considerations, we thus define an *agreed correspondence* and an *agreeable correspondence* as follows. An *agreed correspondence* is the correspondence supported¹¹ by those arguments which are in every preferred extension of every agent. An *agreeable correspondence* is the correspondence supported by arguments which are in some preferred extension of every agent. Thus, the agents will reach a common consensus over a specific mapping m only if the mapping m is an *agreed correspondence*. However, if a mapping m is an *agreeable correspondence* for a given agent Ag , this mean that such mapping can only be considered valid and consensual for that agent.

In the next section, we present a model of agents which put forward arguments and take into account other arguments coming from their interlocutors.

4 Model of Persuasive Agents

In this paper, we are assuming a multi-agent setting containing persuasive agents that do not use the same ontology. Each agent considers the repertoire of argument schemes

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

available to it, and is able to generate a set of arguments and counter-arguments by instantiating these schemes with respect to its interests. Moreover, the agents can record their interlocutors arguments in a commitment store CS [14] and individually evaluate them. Therefore, our persuasive agent can be defined as follows:

Definition 9. An agent Ag_i is defined by a 5-tuple $\langle O_i, VAF_i, CS_j^i, Pref, \varepsilon \rangle$ where O_i is the private ontology; $VAF_i = \langle AR_i, A_i, \mathcal{V}, \eta \rangle$ is the Valued-based Argumentation Framework of the agent Ag_i ; CS_j^i is a commitment store, i.e. a set of arguments where $CS_j^i(t)$ contains propositional commitments taken before or at time t between the Ag_i and other interlocutors; $Pref$ is the private pre-ordering of preferences over \mathcal{V} and ε is the private threshold value.

The set of arguments are not necessarily disjoint. The set of arguments shared by all agents are called *common arguments*: $AR_c \subseteq \bigcap_{x \in AR_i} AR_i \in VAF_i$. Instead, the values $\mathcal{V} = \{M, IS, ES, T, E\}$ are common and shared by all audiences.

In order to take into account the arguments notified in the commitment stores, we extend the definition of valued-based argumentation framework with the following:

Definition 10. An extended Value-Based Argumentation Framework VAF^+ is defined as $\langle AR^+, A^+, \mathcal{V}, \eta^+ \rangle$, where $AR^+ = AR \cup \{\bigcup_{j \neq i} CS_j^i\}$. The definition of A^+ and η^+ are now related to AR^+

Now, we can define the notion of *conviction* as follows:

Definition 11. Let Ag_i be an agent associated with the extended valued-based argumentation framework, VAF^+ and x be an argument provided by another agent Ag_j . The agent Ag_i is convinced by the argument x iff x is acceptable with respect to all audience \mathcal{R} , with $Ag_i \in \mathcal{R}$.

Given this model, in order to determine the acceptability of a potential correspondence, it needs to proceed by means of a dialectical exchange, in which a mapping is proposed, challenged and defended. Many argument protocols have been proposed, e.g. [15]. Particular dialogue games have been proposed based on Dung's Argumentation Frameworks, e.g. [8], and on VAFs [3].

In this paper, we are not considering any specific protocol or persuasive dialogue. However, the idea of a dialogue is that agents reply to each other in order to reach the interaction goal, i.e. an agreement. Thus, given a set of social and autonomous agents, and a set of potential correspondences $\{m_1, \dots, m_i, \dots\}$, an agent initiates a persuasion dialogue when it wants present its viewpoint to the other agents. Specifically, for each mapping m_i , if the agent wants to accept that mapping, it will put forward arguments for m_i . In the negative case, it will put forward arguments against. If the other agents have no arguments against/for the mapping, it closes the dialogue. If the players have the same convictions, the the arguments is accepted and the dialogue closes. Otherwise, the goal of the dialogue is the resolution of the conflict by verbal means, and thus with an exchange of arguments and counter-arguments.

The dialogue between agents can thus consist simply of the exchange of individual arguments, from which they can compute acceptable mappings over the CS , by computing the preferred extensions. 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 [13] a detailed approach to argue over alignments and complete argumentation framework, with a common set of arguments, is proposed.

5 A Walk through Example

Let us assume that some agents or services need to interact with each other using two independent but overlapping ontologies. The first agent, Ag_1 uses the bibliographic ontology¹² from the University of Toronto, based on bibTeX; whereas the second agent, Ag_2 , uses the General University Ontology¹³ from Mondeca¹⁴. For space reasons, we only consider a subset of these ontologies, shown in Table 2, where the first and second ontologies are represented by O_1 and O_2 respectively.

We will assume that the set of candidate mappings, provided by the Ontology Alignment Service (OAS), is the following::

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

The generation of the arguments and counter-arguments of the Ag_1 and Ag_2 are achieved by instantiating the argumentation schemes, discussed previously, with respect to the agent's preferences and threshold. However, here we assume a degree of confidence n that is above the threshold of both 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 agent's audience. The agents Ag_1 and Ag_2 take on the part, respectively, of the audience \mathcal{R}_1 and \mathcal{R}_2 . For space reasons, we will only evaluate the mapping m_1 ¹⁶.

The argumentation starts, with the agent Ag_1 that wants to reject the mapping m_1 and will thus argue against it, forwarding an argument A . A states that none of the super-concepts of the concept $O_1: \textit{Press}$ are mapped to any super-concept of $O_2: \textit{Periodical}$. The agent Ag_2 , instead, does not agree and counter-argues with an argument B . B argues for m_1 , because two sub-concepts of $O_1: \textit{Press}$, $O_1: \textit{Magazine}$ and $O_1: \textit{Newspaper}$, are mapped to two sub-concepts of $O_2: \textit{Periodical}$, $O_2: \textit{Magazine}$ and $O_2: \textit{Newspaper}$, as established by m_4 and m_5 . The agent Ag_1 attacks B with the argument C , because $O_1: \textit{Press}$ and $O_2: \textit{Periodical}$ do not have any lexical similarity. The agent Ag_2

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

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

¹⁴ Note that ontology O_2 has been slightly modified for the purposes of this example.

¹⁵ m_1 states an equivalence correspondence with confidence n between the concept \textit{Press} in the ontology O_1 and the concept $\textit{Periodical}$ in the ontology O_2

¹⁶ An extended version of this example is provided in [13].

Table 2. Excerpts of O_1 and O_2 ontologies

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

does not have any other argument to reply to C but it supports the correspondences m_4 , m_5 and m_6 by six arguments. 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 D and I .

This position is illustrated in Figure 2, 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.

Table 3 shows these arguments, labelled with an identifier Id , its type \mathcal{V} , and the attacks A that can be made on it by opposing arguments.

Table 3. Arguments for and against the correspondences m_1 , m_4 and m_5

Id	Argument	A	\mathcal{V}
A	$\langle \bar{\exists}m = \langle superconcept(Press), superconcept(Periodical), n, \equiv, \rangle, m_1, - \rangle$	B,D,I	ES
B	$\langle \exists m = \langle subconcept(Press), subconcept(Periodical), n, \equiv, \rangle, m_1, + \rangle$	A,C	ES
C	$\langle Label(Press) \not\approx_T Label(Periodical), m_1, - \rangle$	B	T
D	$\langle Label(Magazine) \approx_T Label(Magazine), m_4, + \rangle$		T
E	$\langle \exists m = \langle siblingConcept(Magazine), siblingConcept(Magazine), n, \equiv, \rangle, m_4, + \rangle$		ES
F	$\langle \exists m = \langle superconcept(Magazine), superconcept(Magazine), n, \equiv, \rangle, m_4, + \rangle$		ES
G	$\langle Label(Newspaper) \approx_T Label(Newspaper), m_5, + \rangle$		T
H	$\langle \exists m = \langle siblingConcept(Newspaper), siblingConcept(Newspaper), m_5, + \rangle$		ES
I	$\langle \exists m = \langle superconcept(Newspaper), superconcept(Newspaper), n, \equiv, \rangle, m_5, + \rangle$		ES




Fig. 2. Value-Based Argumentation Frameworks

Finally, we can compute the acceptability of the arguments, computing the preferred extensions (see Table 5). Therefore, the arguments accepted by both audiences

Table 4. Preferred Extensions

Preferred Extensions for the framework (a)	Audience
$\{A, C, D, E, G, H\}$	\mathcal{R}_1
$\{A, C, D, E, F, G\}, \{B, I, D, E, F, G\}$	\mathcal{R}_2
$\{A, C, D, E, F, G\}, \{B, I, D, E, F, G\}$	

are $\{D, E, G, H\}$. Arguments A, C are, however, both 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 for the agent Ag_1 (since B is unacceptable to \mathcal{R}_1), while the mappings m_4 and m_5 will both be accepted (they are both accepted by \mathcal{R}_1 and both acceptable to \mathcal{R}_2). The *agreed correspondence* are then m_4 and m_5 .

6 Summary and Outlook

In this paper we have outlined a framework that provides a novel way for agents, who use different ontologies, to argue and reach agreement over ontology 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.

Our approach is able to give a formal motivation for the selection of a correspondence, and enables consideration of an agent's interests and preferences that may influence the selection of a correspondence. We believe that this approach will aim at reaching mutual understanding and communicative work in agents system more sound and effective. Future work will include experimental testing in order to demonstrate the practicality of our approach. An interesting topic for future work would be to investigate how to argue about the whole alignments, and not only the individual candidate mapping.

7 Acknowledgements

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

References

1. K. Aberer and et al. Emergent semantics principles and issues. In *Proceedings of Database Systems for Advances Applications, 9th International Conference, DASFAA 2004*, 2004.

2. T. Bench-Capon. Persuasion in practical argument using value-based argumentation frameworks. In *Journal of Logic and Computation*, volume 13, pages 429–448, 2003.
3. T. J. M. Bench-Capon. Agreeing to differ: Modelling persuasive dialogue between parties without a consensus about values. In *Informal Logic*, volume 22, pages 231–245, 2002.
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. In *Proceedings of the International Conference on Management of Data (SIGMOD)*, pages 383–394.
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. P. Dunne and T. J. M. Bench-Capon. Two party immediate response disputes: Properties and efficiency. In *Artificial Intelligence*, volume 149, pages 221–250, 2003.
9. M. Ehrig and S. Staab. Qom - quick ontology mapping. In *Proceedings of the International Semantic Web Conference*, 2004.
10. J. Euzenat. Alignment infrastructure for ontology mediation and other applications. In Hepp, editor, *Proceedings of the First International workshop on Mediation in semantic web services*, 2005.
11. J. Euzenat and P. Valtchev. Similarity-based ontology alignment in owl-lite. In *Proceedings of European Conference on Artificial Intelligence (ECAI 04)*, 2004.
12. T. R. Gruber. A translation approach to portable ontology specifications. *Knowledge Acquisition*, 5(2):199–220, 1993.
13. L. Laera, V. Tamma, J. Euzenat, T. Bench-Capon, and T. Payne. Reaching agreements over ontology alignments. In *Proceedings of the Fifth International Semantic Web Conference (ISWC'06)*, 2006.
14. N. Maudet and B. Chaib-draa. Commitment-based and dialogue-game based protocols-news trends in agent communication language. *Knowledge Engineering Review*, 17(2):157–179, 2003.
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)*, pages 209–225, 2004.
16. P. Shvaiko and J. Euzenat. A survey of schema-based matching approaches. *Journal on data semantics*, 4:146–171, 2005.
17. 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.
18. 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.