Deciding Agent Orientation on Ontology Mappings

Paul Doran[1], Terry R. Payne[1], Valentina Tamma[1], and Ignazio Palmisano[2]

[1] Department of Computer Science, University of Liverpool,
Liverpool L69 3BX, United Kingdom
{P.Doran,T.R.Payne,V.Tamma}@liverpool.ac.uk
[2] School of Computer Science, University of Manchester M13 9PL, UK
ignazio.palmisano@cs.manchester.ac.uk

Abstract. Effective communication in open environments relies on the ability of agents to reach a mutual understanding of the exchanged message by reconciling the vocabulary (ontology) used. Various approaches have considered how mutually acceptable mappings between corresponding concepts in the agents' own ontologies may be determined dynamically through argumentation-based negotiation (such as Meaning-based Argumentation, *MbA*). In this paper we present a novel approach to the dynamic determination of mutually acceptable mappings, that allows agents to express a private acceptability threshold over the types of mappings they prefer. We empirically compare this approach with the Meaning-based Argumentation and demonstrate that the proposed approach produces larger agreed alignments thus better enabling agent communication. Furthermore, we compare and evaluate the fitness for purpose of the generated alignments, and we empirically demonstrate that the proposed approach has comparable performance to the *MbA* approach.

1 Introduction

The problem of dynamic reconciliation of *ontologies* (vocabularies) used by agents during interactions has received significant attention [8,10,12], due to the growing adoption of mobile and service computing. In these scenarios, agents situated in open environments encounter unknown agents offering new services as a user's context or location changes. As the heterogeneity that permeates these environments increases, fewer assumptions on the vocabulary and content of these ontologies can be made, thus hindering seamless interaction between the agents.

The reconciliation of heterogeneous vocabularies has been investigated at length by research efforts in *ontology alignment* [7], which tries to determine suitable mappings between two ontologies. However, there are few traditional alignment approaches suitable for use in purely dynamic interaction scenarios as most require human intervention, or they align the ontologies at design time [11]. Although recent systems [6] have emerged that can generate alignments at

P.F. Patel-Schneider et al. (Eds.): ISWC 2010, Part I, LNCS 6496, pp. 161–176, 2010.

run time, these are often machine-learning based, requiring pre-labelled training data to guide the learning process.

Whilst this has been demonstrated to be effective when such data is available, it is not always suitable for all dynamic problems. Two agents may encounter each other for the first time with the aim of interacting to achieve some goal (where each agent may have its own preferences or policies over the terms and axioms used within a specific interaction). Whilst alignments may exist between the agents' ontologies, these may have been determined under different contexts or assumptions, and thus may not necessarily satisfy the current agents' preferences or policies. In order to address this limitation, and to consider the context within which the alignment is to be used, Laera *et. al.* [8] proposed in their Meaning-based Argumentation (*MbA*) approach the use of argumentation to select a set of mappings (i.e. an *alignment*) that is mutually acceptable to the negotiating agents, from the union of disparate, precomputed alignments where different alignments may have previously been generated (e.g. for previous agent-agent interactions) and then published or retained for future use.

Therefore, the problem can be cast as a search for a mutually acceptable set of mappings between two ontologies O_1 and O_2 (in the union of mappings previously computed), given the agents' individual, private preferences over the mapping type (i.e. terminological, extensional, etc.). Approaches such as those proposed by Laera *et al.* [8] and dos Santos *et al.* [10] assume that mappings have an associated confidence value, and based on this, utilise both an acceptance threshold, ϵ, and their preferences to determine whether or not a candidate mapping is suitable for a task.

The search is conducted collaboratively, through the use of argumentation. By specifying arguments that *support* (or *refute*) different mappings, the negotiating agents identify a subset of mappings that are considered mutually acceptable, which can subsequently be used to support further communication between the agents. The arguments are determined from the individual agent's preferences over the mapping types (which can vary, depending on the agents task or the expressive power of ontology it commits to) and its acceptance threshold. The argumentation converges on a set of *agreed* mappings, i.e. mappings that are mutually acceptable to the negotiating agents.

As the generation of arguments is directed by a single preference and acceptance threshold specified by each agent, this approach is susceptible to rejecting those mappings which, whilst not optimal, may still be considered acceptable to all the agents involved. This results in smaller alignments, which may fail to sufficiently support the agent's subsequent communication. This approach may also fail to reflect the true preference of an agent, as the different grounds supporting the choice or type of mapping may actually generate similar mappings in some cases.

In this paper, we demonstrate the effect of this limitation on the resulting alignment empirically, and propose a novel approach for generating arguments for each of the candidate mappings, utilising a weaker notion of *suitability* than that originally proposed. The *flexible approach for determining agents' orientation on* *ontology mappings* (*FDO*) proposed here provides a flexible mechanism for agents to decide whether they support or refute an argument about a mapping, and hence it allows agents to *compromise* over the suitable mappings; i.e by arguing in favour of an assertion that may not be amongst the preferred ones, but that facilitates the negotiation process in converging on a mutually acceptable solution. In this way, the agents create a larger consensus base, by increasing the number of arguments over which the agents negotiate, and that better reflect the agents' preferences over the type of mappings deemed to facilitate the exchange of messages. Whilst this approach results in agents relaxing some of their preferences over suitable mappings, we demonstrate that it produces a larger consensus over possible mappings due to the generation of a greater number of arguments in favour of the candidate mappings (compared to Laera *et al.*'s *MbA* approach), and better reflects the agents preferences than when only a single threshold and preference value is used. We also demonstrate that allowing the negotiation to take place over a larger set of arguments does not degrade the quality of the alignment produced, measured in terms of precision and recall over query answering tasks. Therefore, the contribution of this paper is twofold: we provide a novel approach to the determination of whether an agent supports or refutes an argument, and we provide an evaluation of this novel approach against the *MbA* approach.

The paper is organised as follows: the *MbA* approach is briefly summarised, followed by the description of our novel *FDO* approach for determining an agent's orientation on a mapping. This approach is then illustrated by means of an example, before being evaluated empirically. The results of the evaluation are then discussed, before concluding.

2 Arguing over Ontology Mappings

Meaning-based Argumentation (*MbA*), as proposed by Laera *et al.* [8], assumes that a number of precomputed alignments (i.e. sets of mappings) exist within some publicly available repository. A similar assumption is also made by dos Santos *et al.* [10], whereby such alignments are known (possibly computed *on-the-fly*) by different agents. Before presenting our *flexible approach for determining agent orientation*, we first give the formal definition of these alignments, and summarise the *MbA* approach[1].

A mapping between two agent ontologies O_1 and O_2 is described as a tuple: $m = \langle e, e', n, r\rangle$, where $e \in O_1$ and $e' \in O_2$ are the entities (concepts, properties or individuals) between which a relation, r, is asserted, such as equivalence, or subsumption, and n is a degree of confidence in this correspondence [7]. These mappings can either be computed offline and stored by a dedicated server, an *Ontology Alignment Service*, that provides the set of available *candidate mappings* the agents need to argue over [8], or they can be determined on the fly [10]. Whatever the approach used to generate the mappings, the argumentation process considers as input a set of pre-computed mappings, and a set

[1] We focus primarily on the *MbA* approach since the negotiation phase in dos Santos *et al.* is the same as the one used in *MbA*.

of justifications that motivate the existence of a mapping, that are provided by the mapping generation approach.

The Meaning-based Argumentation (*MbA*) process is based on the *Value-Based Argumentation Framework* (VAF) [3], which introduced the notions of *audience* and *preference values*. An audience represents a group of agents who share the same preferences over a set of values, with a single value being assigned to each argument. This framework extends the seminal work by Dung on the use of argumentation theory [5]. In Dung's framework, attacks always succeed; in essence they are all given equal value. For deductive arguments this suffices, but within the ontology alignment negotiation scenario [8] the persuasiveness of an argument could change depending on the audience, where an audience represents a certain set of preferences. Thus, the Value-Based Argumentation Framework (VAF) facilitates the assignment of 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. Hence, it is possible to systematically relate strengths of arguments to their motivations and to accommodate different audience interests.

Definition 1. *A Value-Based Argumentation Framework (VAF) is defined as* $\langle AR, A, \mathcal{V}, \eta \rangle$, *where:*

- $\langle AR, A \rangle$ *is an argumentation framework;*
- $\mathcal{V}$ *is a set of* k values *which represent the types of arguments;*
- $\eta : AR \rightarrow \mathcal{V}$ *is a mapping that associates a value* $\eta(x) \in \mathcal{V}$ *with each argument* $x \in AR$.

The types of arguments represented by $\mathcal{V}$ typically varies, depending upon the application. Within the *MbA* process, the values of $\mathcal{V}$ correspond to five different types of ontological mismatches that can occur between ontologies, as represented in Table 1.

In order to model the notion of different agents having different perspectives on the same candidate mappings, we define an *audience*, i.e. the representation of a preference ordering of $\mathcal{V}$. The notion of *audience* is central to the VAF. Audiences are individuated by their preferences over the values. Thus, potentially, there are as many audiences as there are orderings[2] of $\mathcal{V}$. The set of arguments is assessed by each audience in accordance to its preferences. An audience is defined as follows:

Definition 2. *An* audience *for a VAF is a binary relation* $\mathcal{R} \subseteq \mathcal{V} \times \mathcal{V}$ *whose irreflexive transitive closure,* $\mathcal{R}^*$, *is asymetric, 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* $(v_i, v_j) \in \mathcal{R}^*$.

As this notion allows different agents (represented by an audience) to have different perspectives on the same candidate mapping, we need to model what it means for an argument to be acceptable relative to some audience. This is defined within the VAF as follows:

[2] Number of audiences corresponds to the different combinations of the elements in $\mathcal{V}$; i.e. *Number of audiences* $= |\mathcal{V}|!$

Table 1. The classification of different types of ontological alignment approaches

Semantic	M	These methods utilise model-theoretic semantics to determine whether or not there is a correspondence between two entities, and hence are typically deductive. Such methods may include propositional satisfiability and modal satisfiability techniques, or logic based techniques.
Internal Structural	IS	Methods for determining the similarity of two entities based on the internal structure, which may use criteria such as the range of their properties (attributes and relations), their cardinality, and the transitivity and/or symmetry of their properties to calculate the similarity between them.
External Structural	ES	Methods for determining external structure similarity may evaluate the position of the two entities within the ontological hierarchy, as well as comparing parent, sibling or child concepts.
Terminological	T	These methods lexically compare the strings (tokens or n-grams) used in naming entities, or in the labels and comments concerning entities. Such methods may employ normalisation techniques (often found in Information Retrieval systems) such as stemming or eliminating stop-words, etc.
Extensional	E	Extension-based methods which compare the extension of classes, i.e., their set of instances. Such methods may include determining whether or not the two entities share common instances, or may use alternate similarity based extension comparison metrics.

Definition 3. *Let $\langle AR, A, \mathcal{V}, \eta \rangle$ be a VAF, with R and S as subsets of AR, and an audience $\mathcal{R}$:*

(a) *For $x, y \in AR$, x is a* successful attack *on y with respect to $\mathcal{R}$ if $(x, y) \in A$ and $\eta(y) \not\succ_{\mathcal{R}} \eta(x)$.*
(b) *$x \in AR$ is* acceptable *with respect to S with respect to $\mathcal{R}$ 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) *S is* conflict-free *with respect to $\mathcal{R}$ if for every $(x, y) \in S \times S$, either $(x, y) \notin A$ or $\eta(y) \succ_{\mathcal{R}} \eta(x)$*
(d) *A conflict-free set S is* admissible *with respect to $\mathcal{R}$ if every $x \in S$ is acceptable to S with respect to $\mathcal{R}$*
(e) *S is a* preferred extension *for the audience $\mathcal{R}$ if it is a maximal admissible set with respect to $\mathcal{R}$*
(f) *$x \in AR$ is* subjectively acceptable *if and only if x appears in the preferred extension for some specific audience.*
(g) *$x \in AR$ is* objectively acceptable *if and only if x appears in the preferred extension for every specific audience.*
(h) *$x \in AR$ is* indefensible *if it is neither subjectively nor objectively acceptable.*

Laera *et. al.* [8] subsequently adopted the VAF for the Meaning-based Argumentation (*MbA*) process, allowing agents to express preferences for different mapping types, and restricting the arguments to those concerning ontology mappings allowing agents to explicate their mapping choices. The definition of an agent and an argument are as follows:

Definition 4. *An agent, Ag_i, is characterised by the tuple $\langle O_i, VAF_i, Pref_i, \epsilon_i \rangle$ such that O_i is an ontology, VAF_i is a instance of a VAF, $Pref_i$ is an ordering over the possible values in $\mathcal{V}$ and ϵ_i is a private threshold between 0 and 1.*

Definition 5. *An argument $x \in AR$ is a triple $x = \langle G, m, \sigma \rangle$ where m is a mapping, G is the grounds justifying the* prima facie *belief that the mapping does or does not hold and σ is one of $\{+, -\}$ depending on whether the argument is that m does or does not hold.*

Thus, when arguing over ontology mappings using the VAF, an argument $x \in AR$ either supports or refutes a mapping m, depending on the value of σ. An agent determines this σ (i.e. decides whether to argue for or against a mapping) based on its preferences and threshold. Given the set of mappings $\mathcal{M} = \{m\}_{j=1,\ldots,p}$, such that p is the number of mappings, and the function[3] $\tau : M \rightarrow \mathcal{V} \mid \tau(m) = v \in \mathcal{V}$ then an agent can set the value of σ for an argument, x, about a mapping, m, as follows:

$$\sigma = \begin{cases} +, & \text{if } max(Pref_i) = \tau(m) \wedge n_m \geq \epsilon_i \\ -, & \text{otherwise} \end{cases} \tag{1}$$

The notion of an attack and counter-attack is also formally defined, whereby x is attacked by the assertion of its negation, $\neg x$.

Definition 6. *An argument $x \in AR$ attacks an argument $y \in AR$ if x and y are arguments for the same mapping, m, but with different σ. For example, if $x = \langle G_1, m, + \rangle$ and $y = \langle G_1, m, - \rangle$, x counter-argues y and vice-versa.*

The agents can now express, and exchange, their arguments about ontology mappings and decide from their perspective, audience, what arguments are in their preferred extension; but the agents still need to reach a mutually acceptable position with regards to what ontology alignment they actually agree upon. Laera *et. al.* define the notion of *agreed* and *agreeable* alignment as follows:

Definition 7. *An* agreed alignment *is the set of mappings supported by those arguments which are in every preferred extension of every agent.*

Definition 8. *An* agreeable alignment *extends the agreed alignments with those mappings supported by arguments in some preferred extensions of every agent.*

Thus, a mapping is *rejected* if it is in neither the agreed nor agreeable alignment. Given the context of agent communication it is rational for the agents to accept as many candidate mappings as possible [8], thus both sets of alignments are considered. The agents should only completely disagree when they want the opposite, indeed, the agents gain little by arguing and not reaching some kind of agreement.

The definition of audience is central to the notion of *acceptability* of an argument, since given a set of arguments, and their respective counter-arguments, the agents in an audience need to consider which of them they should accept. The acceptability of some arguments with respect to some audience, depends on the agents ability to determine a *preferred extension* that represents a consistent

[3] In some cases $\tau(m) = \eta(x_m)$, however in general this assumption does not hold.

position within an argumentation framework that can be defended against all attacks, and cannot be further extended without causing it to be inconsistent or open to attacks. The mappings supported in the preferred extensions form the mutually agreed set of mappings [8].

3 A Flexible Approach for Determining Agents' Orientation on Mappings

The meaning based negotiation approach by Laera *et al.* is the first attempt to tackle the problem of dynamic reconciliation of heterogeneous agent ontologies. Whilst the approach has the merit of having highlighted an important problem, the proposed solution presents a serious limitation, primarily due to the way σ is obtained.

In Laera's approach an agent argues only in favour of those arguments whose grounds have the *highest* ranking in the ordering of agent preferences, whilst all the other mappings are argued against. Hence, effectively the agents can only express one preference towards one type of mapping, and will argue against any other type of mapping, therefore greatly reducing the possibility to find a suitable agreement on a set of mappings. In other words, this approach fails to distinguish mappings that are *less preferred* from those mappings for which an agent is against.

In addition, this type of strict decision process could potentially increase the chance that inconsistent mappings are determined by the VAF. The walkthrough example presented in the next section illustrates an occurrence of this unlikely but possible event.

In this paper, we present an alternative approach that aims at recognising how agents can have different preferences over the types of mappings to use in interactions with other agents, and that these preferences can influence the decision making process behind the negotiation. An agent would ideally try to maximise the use of those types of mappings with the highest preferences, however, since it needs to interact with other agents (with their own preferences) then it might decide to *compromise*, i.e. to agree to use a less preferred mapping type if this facilitates communication.

This is the main motivation behind the novel approach to mapping selection that we present here. It builds on some of the notions presented in the previous section for the *MbA* approach, but gives agents more flexibility in deciding their orientation, i.e. whether to support or refute a mapping.

Given two agents ontologies O_1 and O_2, a mapping between $e \in O_1$ and $e' \in O_2$ is a tuple $m = \langle e, e', n, r \rangle$, as defined in the previous section. Analogously to *MbA* we define a VAF as a tuple $\langle AR, A, \mathcal{V}, \eta \rangle$ that is similar to the definition given in the previous section (likewise for the definition of mapping m). In the *flexible approach for determining agents' orientation on a mapping* (*FDO*) proposed here, we define an agent as a tuple $Ag_i = \langle O_i, VAF_i, Pref_i, \phi_i \rangle$, where O_i is an ontology, VAF_i is a instance of a VAF, $Pref_i$ is an ordering of the values in $\mathcal{V}$ and $\phi_i : \mathcal{V} \rightarrow [0, \cdots, 1]$ maps each v in $\mathcal{V}$ to a value $0 \leq z \leq 1$.

$\phi_i(v)$ represents the minimum confidence threshold for Ag_i to argue in favour of a mapping of type v.

Let us consider the function $\tau : \mathcal{M} \rightarrow \mathcal{V}$ that assigns a $v \in \mathcal{V}$ to every $m \in \mathcal{M}$, then the agent decides whether to be in favour or against the mapping as follows:

$$\sigma = \begin{cases} +, & \text{if } n_m \geq \phi_i(\tau(m)) \\ -, & \text{otherwise} \end{cases} \tag{2}$$

In our approach, an agent determines its orientation on a mapping solely on the basis of the minimum confidence threshold for arguing in favour of a mapping type, and no longer on the ordering of preferences. In this way, the agents express how much they prefer each of the possible mapping types, and how willing they are to argue in their favour. The ordering of preferences is now only used by the VAF when dealing with arguments and their attacks.

4 Illustrative Example

The following example illustrates how the proposed *FDO* approach differs from the original *MbA* approach, assuming the two ontologies illustrated in Figure 1, with the mappings given with their relevant mapping types. Mapping m_1 is a *Terminological* equivalence mapping between concepts A and C, with a confidence of 0.75, whereas mappings m_2 and m_3 are *External Structural* equivalence mappings: m_2 between concepts B and D (confidence 0.85); and m_3 between concepts B and E (confidence 0.65). Note that concepts D and E are disjoint, and thus an alignment containing both mappings m_2 and m_3 would be inconsistent.

Given two agents that wish to communicate: Ag_1 has the preference ordering ES$\succ$T; whereas Ag_2 has the preference ordering T$\succ$ES. Table 2 shows the sets of mappings that will be argued in favour of (+) or against (-). With the *MbA* approach, we assume that the acceptance threshold $\epsilon_1 = \epsilon_2 = 0.5$. Ag_1 will argue in favour of m_2 and m_3, and against m_1; whereas Ag_2 will argue against m_2 and m_3, but in favour of m_1. This is due to the fact that, in the case of Ag_1, only

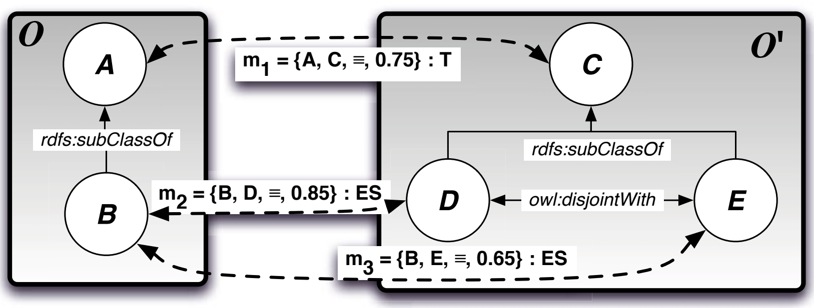

Fig. 1. An alignment between O and O'

Table 2. The arguments that support (+) or refute (-) different mappings, given thresholds and preferences

Approach	Mapping Type Preference	Acceptance Threshold	Arguments *in favor of* +	*against* -
MbA	ES $\succ$ T	0.5	$\{m_2, m_3\}$	$\{m_1\}$
	T $\succ$ ES	0.5	$\{m_1\}$	$\{m_2, m_3\}$
FDO	ES $\succ$ T	ES=0.5, T=0.7	$\{m_1, m_2, m_3\}$	$\{\}$
	T $\succ$ ES	T=0.5, ES=0.7	$\{m_1, m_2\}$	$\{m_3\}$

mappings of the first preference ordering were considered (subject to exceeding the acceptance threshold), and all other mappings were automatically refuted. The resulting attack graph is illustrated in Figure 2 (left), where each argument is assigned a label corresponding to its mapping, and the mapping type. These types are the values in the VAF, with each agent having a private preference ordering over them.

The *FDO* approach, however, assigns a separate acceptance threshold for each mapping type. Ag_1 assumes a 0.5 threshold for ES, but a 0.7 threshold for T, whereas Ag_2 assumes a 0.7 threshold for ES, and a 0.5 threshold for T. In this case, arguments are generated by Ag_1 in favour of all three mappings, whereas Ag_2 generates mappings in favour of m_1 and m_2, but against m_3. Although Ag_1 expresses a preference ordering for ES $\succ$ T, the confidence value of all three mappings exceeds the acceptance threshold for the different mapping types. The resulting attack graph is illustrated in Figure 2 (right).

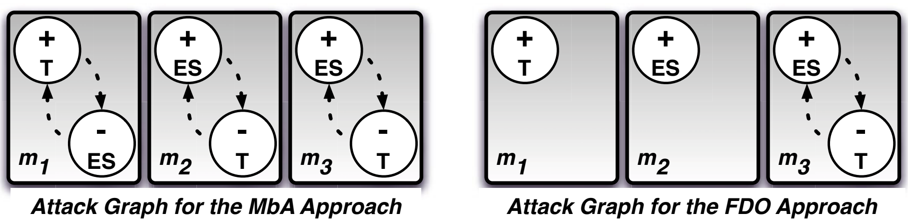

Fig. 2. Attack graphs for the *MbA* and *FDO* procedures

From the attack graphs shown in Figure 2 the preferred extensions for each audience can be computed for the *MbA* approach (see below). This does not produce an agreed alignment, but does produce an agreeable alignment, corresponding to $\{m_1, m_2, m_3\}$. However, as mentioned earlier, if this agreeable alignment were to be accepted by both agents, their ontologies would become inconsistent, thus making the ontologies and the resulting alignment unusable.

– T $\succ$ ES = $\{m_1+,\ m_2\text{-},\ m_3\text{-}\}$
– ES $\succ$ T = $\{m_1\text{-},\ m_2+,\ m_3+\}$

In contrast, the *FDO* approach produces an agreed alignment $\{m_1, m_2\}$, whereas mapping $\{m_3\}$ would only appear in an agreeable alignment. Thus, if the agreed alignment is accepted by both agents, they would be able to communicate with respect to concepts A, B, C, and D, but not with concept E.

5 Empirical Evaluation

The aim of the evaluation is to contrast the proposed *FDO* approach with the original *MbA* approach presented in [8]. Two hypotheses are explored: that the *FDO* approach generates a larger number of supporting arguments, resulting in more selected mappings that *MbA*; and that the increased number of mappings will better support communication tasks such as query answering (i.e. the resulting alignments are *fit for purpose*).

5.1 Evaluating the Generated Arguments

To explore the first hypothesis, the ratio of arguments in favour of mappings to those against was computed for both approaches, and the resulting mappings examined. This requires multiple candidate mappings based on different ontological grounds (and hence different mapping types) between ontologies of the same domain. Eleven ontologies were therefore taken from the *OAEI 2007 and 2008 Conference Track* repositories (with three exceptions[4]), as they represent different domain theories for the same, real-world domain (thus reflecting *real-world heterogeneity*) and can be used generate better pairwise alignments than ontologies from other tracks[5]. These ontologies (originally developed as part of the OntoFarm Project[6]) are listed in Table 3, complete with a brief characterisation in terms of the number of classes (named and anonymous) and properties (object and datatype), and their Description Logic expressivity[7].

For the evaluation, a total of 55 ontology pairs were identified[8]. The alignments between each ontology pair were generated using the *Alignment API* [7], which only produces mappings of type internal structural (IS), external structural (ES) and terminological (T); thus for our evaluation, we assume $\mathcal{V} = \{ES, IS, T\}$.

In order to investigate the differences depending on the threshold, 4 thresholds have been identified for each mapping type. The first, $\epsilon_1 = 0$ corresponds to the case where the agent will argue in favour of all arguments. The remaining

[4] These ontologies have memory requirements of >1.5GB.

[5] `http://oaei.ontologymatching.org/2007/conference/`

[6] `http://nb.vse.cz/~svatek/ontofarm.html`

[7] The expressivity of an ontology (and hence complexity of a reasoner) for a Description Logic is indicated by the concatenation of letters representing different DL operators [1].

[8] Note that the ordering of the ontologies in each pair is irrelevant; thus rendering an evaluation on the symmetric pairs unnecessary. Therefore, a total of $N(N-1)/2$ ontology pairs were used, where N correspond to the 11 ontologies listed in Table 3.

Table 3. Characteristics for the ontology test set

Ontology	Named Classes	Object Prop.	Datatype Prop.	Anon. Classes	Expressivity
cmt	29	49	10	11	$\mathcal{ALCHIF}(\mathcal{D})$
Conf	59	46	18	33	$\mathcal{ALCHIF}(\mathcal{D})$
confOf	38	13	23	42	$\mathcal{SHIF}(\mathcal{D})$
crs_dr	14	15	2	0	$\mathcal{ALCHIF}(\mathcal{D})$
edas	103	30	20	30	$\mathcal{ALCHIF}(\mathcal{D})$
ekaw	73	33	0	27	$\mathcal{SHIN}$
MICRO	31	17	9	33	$\mathcal{ALCHOIF}(\mathcal{D})$
OpenConf	62	24	21	63	$\mathcal{ALCHOI}(\mathcal{D})$
paperdyne	45	58	20	109	$\mathcal{ALCHOIF}(\mathcal{D})$
PCS	23	24	14	26	$\mathcal{ALCHIF}(\mathcal{D})$
sigkdd	49	17	11	15	$\mathcal{ALCHI}(\mathcal{D})$

thresholds are generated by determining the mean $\bar{x}$ and standard deviations of the confidence values for all the mappings for each of the types in $\mathcal{V}$, generated for the evaluation. Thus, $\epsilon_2 = \bar{x} - stdev(x)$, $\epsilon_3 = \bar{x}$, and $\epsilon_4 = \bar{x} + stdev(x)$. Whilst the upper limit ($\epsilon = 1$) was considered, this would have resulted in the agents arguing against all the mappings, resulting in empty alignments. The four levels have been varied independently, producing four actual preferences for each ordering; this produces 144 preferences for each pair of ontologies (again, discarding duplicates). The total number of argumentation situations is, therefore, 7920.

Each experimental argumentation scenario (AS) is defined by the following tuple:

$$AS = (O_1, O_2, P_1, P_2, A^{\sigma+}, A^{\sigma-}, M_{acc})$$

where the set of mappings over which to argue is determined univocally by the ontologies O_1 and O_2, together with the alignment technique used, with P_1 and P_2 representing the actual sets used depending on the approach. For *MbA*, $P_1 = (Pref_1, \epsilon_1)$, $P_2 = (Pref_2, \epsilon_2)$, i.e. for each agent we use the pair composed of the preference ordering and the threshold. For *FDO*, $P_x = (Pref_x, \phi_x)$, but in this case the $Pref_x$ is used only by the VAF (*not* in determining the agent orientation). $A^{\sigma+}$ and $A^{\sigma-}$ represent the set of arguments in favour and against any of the mappings in the argumentation respectively, while M_{acc} represents the set of accepted mappings, i.e., the mappings belonging to at least one preferred extension of one agent. These latter three parameters are recorded for each evaluation.

To compare the results between different ontologies, an index relating $A^{\sigma-}$ to the total number of arguments used has been defined; $NegArgs(AS) : AS \rightarrow [0, 1]$, where:

$$NegArg = \frac{|A^{\sigma-}|}{(|A^{\sigma-}| + |A^{\sigma+}|)}$$

The results have been grouped into nine scenarios based on the first mapping type of each agent preference $Pref_x$. Thus, each row entry in Table 4 is labeled by an Argument Scenario (AS) pair, such that the two values correspond to the first preferred mapping type of Ag_1 and Ag_2 respectively. The results present the averages[9] over each of the subsequent preference values; i.e. the pair (ES,IS) averages values for Ag_1 preferences ES $\succ$ (IS $\succ$ T | T $\succ$ IS), whereas for Ag_2, IS $\succ$ (ES $\succ$ T | T $\succ$ ES), etc. To compare scenarios based on these pairs, a comparison was made between *FDO* and *MbA* by pairing same ordering and same thresholds, since the structure of the preferences is the same for both approaches.

Table 4. Average number of arguments for each scenario

Argument Scenario	**FDO Approach**				**MbA Approach**			
	$A^{\sigma+}$	$A^{\sigma-}$	M_{acc}	NegArgs	$A^{\sigma+}$	$A^{\sigma-}$	M_{acc}	NegArgs
(ES, ES)	5230	2591	1364	0.34	1498	6533	739	0.8
(ES, IS)	5685	2896	1325	0.35	2560	6310	33	0.72
(ES, T)	5720	2698	1358	0.33	1209	7680	92	0.84
(IS, ES)	4626	2216	1136	0.33	1802	4640	20	0.73
(IS, IS)	5230	2591	1364	0.34	3032	4752	439	0.64
(IS, T)	6413	3132	1490	0.33	2177	6388	175	0.76
(T, ES)	4416	2479	1050	0.36	987	5828	73	0.85
(T, IS)	4237	2170	1036	0.35	1488	5135	111	0.77
(T, T)	5230	2591	1364	0.34	700	6880	418	0.89

When using *MbA*, the proportion of arguments against mappings averaged 78%, significantly greater than the 34% average of arguments that were generated against mappings with *FDO*. This can be clearly seen when examining the number of mappings that were generated when using *FDO* (for example, 1325 mappings on average for (ES, IS), compared to only 32.67 with *MbA*). This higher number of negative arguments generated by *MbA* suggests that it may result in a higher probability of generating empty alignments, thus resulting in unnecessary communication failure. Whilst these results support our hypothesis, it raises questions as to the suitability and hence fitness of the accepted mappings for a given task, which is addressed below.

5.2 Fitness Evaluation

The above evaluation demonstrated that the *FDO* approach produced a greater number of arguments in favour of mappings being generated than when using *MbA*, resulting in a larger number of mutually acceptable mappings. However, it is unclear whether the increase in mappings will result in a better alignment between two ontologies. To address this, new alignments were generated and evaluated (in terms of precision and recall) for a typical *query-answering* task. An alignment was selected to answer simple queries against one of the ontologies

[9] Note that these results include the arguments generated by both agents over all the mappings considered.

involved in the alignment, and the results compared to that achieved when a set of hand-crafted *reference* mappings (from the OAEI Alignment Challenge) were used. To investigate how the availability of different alignments affects the task, four alignment systems (*Asmov*, *Falcon*, *Lily* and *OntoDNA* [13] were used to generate the alignments, and the evaluations were conducted over different alignment combinations.

Table 5. Precision(P), Recall (R) and F-Measure (FM) values for a selection of combinations of alignments (where each alignment system is referenced by their initials)

		Base			**FDO**			**MbA**		
	O_1, O_2	R	P	FM	R	P	FM	R	P	FM
A/L/O	(cmt, ekaw)	0.60	1	0.75	0.60	1	0.75	0.58	1	0.74
	(cmt, sigkdd)	0.19	1	0.32	0.19	1	0.32	0.10	0.81	0.18
	(confOf, ekaw)	0.55	1	0.71	0.55	1	0.71	0.43	1	0.60
A/F/L	(cmt, confOf)	0.83	0.94	0.88	0.83	0.99	0.91	0.77	1	0.87
	(confOf, ekaw)	0.90	0.93	0.91	0.9	0.99	0.94	0.75	1	0.85
	(confOf, sigkdd)	1	0.96	0.98	1	0.99	1	0.59	0.61	0.60
A/O	(cmt, ekaw)	0.60	1	0.75	0.60	1	0.75	0.58	1	0.74
	(cmt, sigkdd)	0.19	1	0.32	0.19	1	0.32	0.10	0.81	0.18
	(confOf, ekaw)	0.55	1	0.71	0.55	1	0.71	0.43	1	0.60
A/F	(cmt, confOf)	0.83	0.94	0.88	0.83	0.99	0.91	0.77	1	0.87
	(confOf, ekaw)	0.90	0.93	0.91	0.90	0.99	0.94	0.75	1	0.85
	(confOf, sigkdd)	1	0.96	0.98	1	0.99	1	0.59	0.61	0.60
L/O	(cmt, ekaw)	0.60	1	0.75	0.60	1	0.75	0.58	1	0.74
	(cmt, sigkdd)	0.19	1	0.32	0.19	1	0.32	0.10	0.81	0.18
	(confOf, ekaw)	0.55	1	0.71	0.55	1	0.71	0.43	1	0.60
F	(cmt, confOf)	0.83	0.94	0.88	0.83	0.99	0.91	0.77	1	0.87
	(confOf, ekaw)	0.90	0.93	0.91	0.90	0.99	0.94	0.75	1	0.85
	(confOf, sigkdd)	1	0.96	0.98	1	0.99	1	0.59	0.61	0.60
O	(cmt, ekaw)	0.60	1	0.75	0.60	1	0.75	0.58	1	0.74
	(cmt, sigkdd)	0.19	1	0.32	0.19	1	0.32	0.10	0.81	0.18
	(confOf, ekaw)	0.55	1	0.71	0.55	1	0.71	0.43	1	0.60
F/L	(cmt, confOf)	0.83	0.94	0.88	0.83	0.99	0.91	0.77	1	0.87
	(confOf, ekaw)	0.90	0.93	0.91	0.90	0.99	0.94	0.75	1	0.85
	(confOf, sigkdd)	1	0.96	0.98	1	0.99	1	0.59	0.61	0.60

The query-answering tasks were evaluated by querying instances from various knowledge-bases (KBs) defined using the different ontologies. In each case, queries were constructed by considering each named concept in one ontology O_1, and querying the KB for O_2. To overcome the ontological heterogeneity, the query was resolved using $O_2 \cup M$, where M was the alignment used. As the resulting instance set depends on the generated alignment, a reference "*gold standard*" instance set was constructed by using the hand-crafted *reference* alignment. To evaluate scenarios where alternate alignments were available from the different alignment systems used, alignments were generated by all of the systems, resulting in 12 different alignments, where each one was partitioned between three or five ontology pairs. Query answering tasks were performed for three cases: when all the mappings in the alignments were aggregated and used without any use of the argumentation process (i.e. *Base*); when *MbA* was used; and when *FDO* was used. In each case, the answers generated for each query were analysed and compared with that obtained when using the Gold Standard set, and the Precision (P), Recall (R) and F-measure (FM) results (using these classical Information Retrieval measures) are reported in Table 5[10].

The results suggest that in most cases, there is a slight improvement in the success of a task when *FDO* is used (compared to *Base*) for the scenarios listed in Table 5, with an average F-measure of 0.83 (compared to 0.82 for *Base*). This contrasts sharply with *MbA*, which achieves only an average F-measure of 0.72. In general, the precision of *FDO* is higher or comparable with that exhibited by *MbA*. Interestingly, when *FDO* is compared with the base case in general, a marked increase in precision is observed. *Base* already represents a best-case scenario, in which the different alignment systems are tuned in order to provide the best accuracy when computing the mappings, and therefore typically generate only those mappings for which the system has the highest level of confidence. These results suggest that the further filtering of results due to the use of *FDO* pays off in terms of the increase in precision.

6 Related Work

A number of solutions have been proposed that attempt to resolve ontological mismatches within open environments [14,4,8,9]. An ontology mapping negotiation [14] was proposed to establish a consensus between different agents using the MAFRA alignment framework. It was based on the utility and meta-utility functions used by the agents to establish if a mapping is accepted, rejected or negotiated, making it highly dependent on the MAFRA framework and unsuitable for other environments.

Bailin and Truszkowski [2] present an ontology negotiation protocol that enabled agents to exchange parts of their ontology, by a process of successive interpretations, clarifications, and explanations. The result was that each agent would converge on a single, shared ontology. However, within an open environment, agents may not always want to modify their own ontologies, as this may affect subsequent communication with other agents.

The work by van Diggelen *et al.* [4] dynamically generates a *minimal* shared ontology, where minimality is evaluated against the ability of the different components to communicate with no information loss. The agents can explain concepts to each other via the communication mechanism; either by defining the concept in terms already understood or by invoking an extensional learning mechanism.

[10] In eight cases, the recall and precision of the *Base* and *FDO* evaluations were of value 1 (i.e. they returned only those instances in the "*gold standard*" instance set), and thus have not been included in the Table. In these cases, the precision of *MbA* was also 1, but the recall varied between 0.9 and 0.99.

However, the ontological model used here is limited and non-standard, as its expressivity supports only simple taxonomic structures, with no properties and few restrictions other than disjointness and partial overlap, and does not correspond to any of the OWL flavours[11]. As a consequence, its applicability to the augmentation of existing real-world, published, OWL ontologies on the web is limited.

dos Santos *et al.* [9,10] address the problem of generating a canonical alignment using an extended version of the VAF, which considers both the strength and value of an argument. They do not consider the problem of dynamically aligning two agent ontologies to facilitate communication and fail to consider the preferences of the agents.

7 Conclusions

This paper presents a novel mechanism for determining whether agents are in favour or against ontology mappings during a process of dynamic selection of mutually acceptable alignements. The *flexible approach for determining agents' orientation on ontology mappings* (*FDO*) allows agents to express a minimum acceptability thresholds for each of the mapping types to include in the alignment used during communication. In this respect *FDO* provides a more flexible framework the Meaning-based argumentation (*MbA*) approach in order to decide whether agents support or refute a mapping.

A systematic evaluation has been presented, aiming at assessing the performance of this novel mechanism over the 11 ontologies used in the OAEI 2007 initiative. In particular, the evaluation investigated whether the *FDO* approach generates larger set of mutually acceptable mappings than the original *MbA* approach, thus improving the possibility of finding an alignment agents can use to interact. In addition, we investigated whether these mappings are *fit for purpose* for a query answering task.

The results obtained suggest that the *FDO* approach produces a considerably larger set of mutually acceptable mappings by reducing the number of mappings an agent argues against when compared with *MbA*. The fitness for purpose evaluation shows that the *FDO* approach has a comparable if not higher F-measure than the case when no argumentation is used, and definitely outperforms *MbA*.

References

1. Baader, F., Calvanese, D., McGuinness, D.L., Nardi, D., Patel-Schneider, P.F. (eds.): The Description Logic Handbook: Theory, Implementation, and Applications. Cambridge University Press, Cambridge (2003)
2. Bailin, S.C., Truszkowski, W.: Ontology negotiation: How agents can really get to know each other. In: Truszkowski, W., Hinchey, M., Rouff, C.A. (eds.) WRAC 2002. LNCS, vol. 2564, pp. 320–334. Springer, Heidelberg (2003)

[11] The authors mention a reformulation of their model using Description Logics, but provide no formal proof of its soundness [4].

3. Bench-Capon, T.: Value based argumentation frameworks. In: Proceedings of Non Monotonic Reasoning, pp. 444–453 (2002)
4. van Diggelen, J., Beun, R.J., Dignum, F., van Eijk, R., Meyer, J.J.: Ontology negotiation in heterogeneous multi-agent systems: The anemone system. Applied Ontology 2(3-4), 267–303 (2007)
5. Dung, P.: On the Acceptability of Arguments and its Fundamental Role in Non-monotonic Reasoning, Logic Programming and n-person Games. In: Artificial Intelligence, vol. 77, pp. 321–358 (1995)
6. Eckert, K., Meilicke, C., Stuckenschmidt, H.: Improving ontology matching using meta-level learning. In: Aroyo, L., Traverso, P., Ciravegna, F., Cimiano, P., Heath, T., Hyvönen, E., Mizoguchi, R., Oren, E., Sabou, M., Simperl, E. (eds.) ESWC 2009. LNCS, vol. 5554, pp. 158–172. Springer, Heidelberg (2009)
7. Euzenat, J., Shvaiko, P.: Ontology Matching. Springer, Heidelberg (2007)
8. Laera, L., et al.: Argumentation over ontology correspondences in mas. In: 6th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2007), Honolulu, Hawaii, USA, May 14-18, p. 228 (2007)
9. dos Santos, C.T., Quaresma, P., Vieira, R.: Conjunctive queries for ontology based agent communication in mas. In: 7th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2008), Estoril, Portugal, May 12-16, vol. 2, pp. 829–836 (2008)
10. dos Santos, C.T., et al.: A cooperative approach for composite ontology mapping. Journal of Data Semantics 10, 237–263 (2008)
11. dos Santos, C.T., Euzenat, J., Tamma, V., Payne, T.R.: Argumentation for reconciling agent ontologies. In: SASFA 2010. Springer, Heidelberg (2010) (in press)
12. Sensoy, M., Yolum, P.: A cooperation-based approach for evolution of service ontologies. In: 7th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2008), Estoril, Portugal, May 12-16, vol. 2, pp. 837–844 (2008)
13. Shvaiko, P., Euzenat, J., Giunchiglia, F., He, B. (eds.): Proceedings of the 2nd International Workshop on Ontology Matching (OM-2007) Collocated with the 6th International Semantic Web Conference (ISWC 2007) and the 2nd Asian Semantic Web Conference (ASWC 2007), Busan, Korea, CEUR Workshop Proceedings, November 11, vol. 304. CEUR-WS.org (2008)
14. Silva, N., Maio, P., Rocha, J.: An approach to ontology mapping negotiation. In: Proceedings of the Workshop on Integrating Ontologies (2005)