Matching Terminological Heterogeneous Ontologies by Exploiting Partial Alignments

Frederik Christiaan Schadd*, Nico Roos†
Department of Knowledge Engineering, Maastricht University
Maastricht, The Netherlands
Email: *frederik.schadd@maastrichtuniversity.nl, †roos@maastrichtuniversity.nl

***Abstract*—Matching ontologies which utilize significantly heterogeneous terminologies is a challenging task for existing matching techniques. These techniques typically exploit lexical resources in order to enrich the ontologies with additional terminology such that more terminological matches can be found. However, they are limited by the availability of an appropriate lexical resource for each matching task. For this scenario, we propose a new technique exploiting partial alignments. We evaluate our technique on a dataset which is characterized by matching problems with significant terminological heterogeneities. Further, we compare our technique with the performance of matching systems utilizing lexical resources to establish whether a partial-alignment-based matcher can perform similarly to a lexical-based matcher. Lastly, we provide a performance indication of a system utilizing both partial alignments and lexical resources.**

Keywords–Ontologies; Semantic Interoperability; Ontology Alignment; Partial Alignments.

I. Introduction

Semantically structured data facilitates many services which are used in a modern society, ranging from agent communication [1] to semantic querying systems [2]. An important criterion for the functionality of such systems is their ability to access multiple sources of semantically structured data. The structure of this data is determined by an ontology, which can be defined using expressive languages, such as the Resource Description Framework Schema (RDFS) or the Web-Ontology-Language (OWL). Information exchange between two sources is possible if both are based on the same ontology.

A common issue is that two sources of semantic data are based on two different ontologies modelling the same domain. The ontologies can differ with regard to their terminology, structure, scope or granularity [3]. In order to transfer data between different ontologies, the data has to be transformed such that it complies with the ontology of the receiving knowledge system. For this to happen, a mapping between the two ontologies is needed. This mapping specifies for every concept in the first ontology whether there is an concept in the second ontology modelling the same information. The process of creating such a mapping is known as ontology mapping.

Ontologies can differ by their applied terminologies. If two ontologies are created by different domain experts then it may happen that the experts prefer different terms in order to refer to the same concepts. This problem can be exacerbated if the two ontologies conform to different design principles, thus varying with regard to their naming formats. Alternatively, the ontologies may simply differ with regard to the used natural language, which can occur in international data-exchange situations. In these scenarios, the two given ontologies have very little overlap with regard to their terminologies, a problem which we refer to as a *terminological gap*. Name-based approaches for ontology mapping are thus unlikely to produce satisfying mapping results in such scenarios.

Terminological gaps between ontologies are typically overcome by exploiting additional resources. Existing techniques exploit lexical resources in order to identify additional names for ontology concepts [4], thus increasing the chance that corresponding concepts are associated with similar names. These techniques however require the presence of an appropriate lexical resource which is modelled with in such detail that alternative labels can be extracted for all concepts of both ontologies. Thus, it is not always the case that a suitable lexical resource is available. However, it might be that there is a different type of resource available for this scenario, being a partial alignment [5]. A partial alignment is an incomplete mapping between the given ontologies stemming from previous matching efforts. An example of such an effort is a domain expert being unable to complete the mapping due to time constraints. The main problem is how a partial alignment can be exploited to aid the matching between ontologies between which exists a terminological gap.

In this paper, we tackle this problem by proposing a profile-based similarity which exploits the correspondences of the given partial alignment. A typical profile similarity creates a virtual document of each concept by gathering the encoded terminology of related concepts and itself. The core intuition is that two concepts are considered similar if their documents are similar. A key component here is that information of related concepts is exploited as well. Our approach is based on enriching the given ontologies by extracting the encoded semantic relations of each correspondence of a partial alignment, also known as *anchors*. We define an extension of a given profile similarity which utilizes the added relations in order to identify additional terminology for each concept. We evaluate our approach on a dataset consisting of matching problems with distinct terminological gaps. Further, we compare our approach to the performance of existing systems utilizing lexical resources. Lastly, we provide a performance indication for systems utilizing both lexical resources and partial alignments.

The remainder of this paper is structured as follows. We discuss relevant work in Section II. We introduce profile similarities and their ability to deal with terminological heterogeneous ontologies in Section III. The proposed approach is detailed in Section IV. The experimental results are presented in Section V. We present the conclusions and future research topics in Section VI.

II. Related Work

Profile similarities have seen a rise in use since their inception. Initially developed for the Falcon-AO system [6], this type of similarity has seen use in ontology mapping systems, such as AML [7] and RiMoM [8]. These systems typically apply the same scope when gathering information for a concept profile, being the parent concepts and children concepts. Some systems, such as YAM++ [9], limit the scope to the information of the concept annotations and labels.

There exist some works that aim to extend the scope of exploited profile information in order to improve the effectiveness of the similarity. The deployed profile similarity in the mapping system PRIOR [10] extends the scope of exploited information to the grand-parent concepts and grand-children concepts, providing a larger scope of exploitable context.

The combination of profile similarities with a set of provided anchors has been tackled in [11]. However, [11] has some fundamental differences compared to this paper. In [11], ontology concepts are compared to the given anchors using a selection of similarity metrics, e.g., a string, instance, and lexical metric. Instead of extracting terminology, the concept profiles are created using the results of these similarity calculations. Concepts are matched if they exhibit comparable degrees of similarities towards the provides anchors. Therefore, this technique can only match terminological ontologies if appropriate similarity metrics are applied and both ontologies contain the exploited meta-data for these similarities.

A. Semantic Enrichment

One way in which additional information can be exploited is through *semantic enrichment*. *Semantic enrichment* describes any process which takes any ontology O as input and produces as output the enhanced ontology $E(O)$, such that $E(O)$ expresses more semantic information than O. Typically, a semantic enrichment process exploits resources such as stop word lists or lexical resources for this purpose. Semantic enrichment has been applied in ontology mapping systems in a non-profile context. Examples are the addition of synonyms to the concept descriptions by exploiting lexical resources. LogMap [4] is capable of adding information from WordNet or UMLS to the ontologies prior to mapping. YAM++ [9] uses a machine translator to generate English translations of labels prior to mapping. Multilingual ontology mapping has been specifically addressed in [12]. Ontologies are enriched with multilingual labels using a machine translator. A feature vector for each match candidate is constructed using a combination of similarities and aggregation techniques. Match candidates are then classified using a support vector machine.

A noteworthy application of semantic enrichment for a profile similarity is the work by Su et al. [13]. Here, the semantic enrichment process exploits a set of documents. Using a linguistic classifier and optional user input the documents are assigned to the ontology concepts, such that each assignment asserts that the ontology concept is discussed in its associated document. The concept profiles are then created by gathering terminological information from the assigned documents.

III. Profile Similarities and Terminological Gaps

Profile similarities are a robust and effective type of similarity metric and deployed in a range of state-of-the-art ontology matching systems [6][7][8]. They rely on techniques pioneered in the field of information retrieval [14], where the core problem is the retrieval of relevant documents when given an example document or query. Thus, the stored documents need to be compared to the example document or query in order to determine which stored document is the most relevant to the user. A profile similarity adapts the document comparison techniques by constructing a virtual document for each ontology concept, also referred to as the *profile* of that concept, and determines the similarity between two concepts x and y by comparing their respective profiles. The core intuition

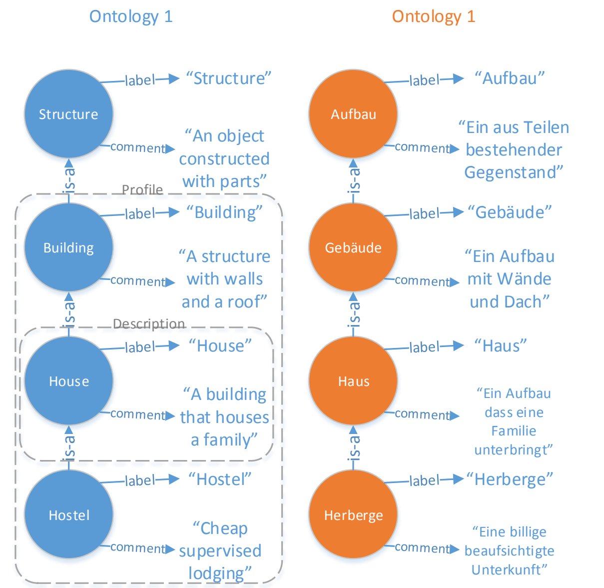

Figure 1. Illustration of a terminological gap between two ontologies modelling identical concepts.

of this approach is that x and y can be considered similar if their corresponding profiles can also be considered similar.

As their origin implies, profile similarities are language-based techniques [15]. Language-based techniques interpret their input as an occurrence of some natural language and use appropriate techniques to determine their overlap based on this interpretation. A language-based technique might for instance perform an analysis on the labels of the concept in order to determine their overlap. For instance, given the two concepts *Plane* and *Airplane* a language-based analysis of their labels would result in a high score since the label *Plane* is completely contained within the label *Airplane*. Thus, despite the labels being different, a high similarity score would still be achieved. However, the degree of surmountable label-difference has a limit for language-based techniques. The labels of the concepts *Car* and *Automobile* have very little in common with regard to shared characters, tokens or length. Thus, many language-based techniques are unlikely to result in a high value.

Profile similarities have the advantage that they draw from wide range of information per concept. Thus terminological differences between the labels of two concepts can still be overcome by comparing additional information. This additional information typically includes the comments and annotations of the given concept and the information of semantically related concepts [6][10].

In order for two profiles to be similar, they must contain some shared terminology. For example, the concepts *House* and *Home* can still be matched if their parents contain the word *Building* or if a concept related *Home* contains the word *"House"*. In order for profile similarities to be effective, it is still required that the two given ontologies O_1 and O_2 exhibit some overlap with regard to their terminologies. However, this is not always the case as two ontologies can model the same domain using a completely different terminology. This can be the result of one ontology using synonyms, different naming conventions or the usage of acronyms. Furthermore, two ontologies might even be modelled in a different natural language. For example, one might need to match two biomedical ontologies where one is modelled in English and one in Latin. Thus, it is a real possibility that even if there is some overlap, there can exist corresponding parts of two ontologies exhibit little to no terminological overlap. The terminological gap between two ontologies is illustrated in Figure 1.

Figure 1 displays an example *Ontology 1* next to a series of concepts from a different ontology, *Ontology 2*, modelling the same entities. The terminological gap is illustrated through the fact that all information in *Ontology 2* is modelled in German instead of English. As we can see, comparing the concept *House* with its equivalent concept *Haus* using a typical profile similarity is unlikely to produce a satisfying result since the neither they nor their related concepts contain any overlapping terminology. Therefore, additional measures are necessary in order to ensure the effectiveness of profile similarities when the given ontologies have little to no shared terminology.

IV. Anchor-Based Profile Enrichment

A typical profile similarity is inadequate for ontology matching problems with significant terminological gaps. One way of tackling this issue is through semantic enrichment by exploiting lexical resources such as WordNet [16] or UMLS [17]. Techniques which fall under this category work by looking up each concept in the given resource and adding synonyms, additional descriptions or translations to the concept definition. However, these techniques rely on several assumptions: (1) the availability of an appropriate resource for the given matching problem, (2) the ability to locate appropriate lexical entries given the naming formats of the ontologies, and (3) the ability to disambiguate concept meanings such that no incorrect labels or comments are added to the concept definition. We can see that the performance of such techniques is severely impacted if any of these assumptions fail. If (1) and (2) fail then it is not possible to add additional information to the concept definition, thus causing the ontology concepts to be compared using only their standard profiles. To ensure the ability of identifying correct lexical entries when dealing with ambiguous concepts, one needs to apply a disambiguation technique. State-of-the-art disambiguation systems can achieve an accuracy of roughly 86% [18], meaning that even if a state-of-the-art system is applied there is still a significant proportion of concepts which would be associated with unrepresentative information based on incorrectly designated lexical entries.

If an appropriate lexical resource is not available, other measures are necessary to overcome the terminological gap. These typically are the exploitation of other ontological features, for example the ontology structure. However, it may be the case that instead of a lexical resource a different kind of resource is available to be exploited. For a given mapping problem it is possible that an incomplete alignment, also refereed to as *partial alignment*, is available as additional input. A partial alignment can stem from efforts such as a domain expert attempting to create an alignment, but being unable to complete it due to given circumstances, or from a high-precision system generating such an alignment. The correspondences of the given partial alignment can then be exploited in order to determine the unidentified correspondences.

Our approach aims at adapting profile similarities to be appropriate for matching problems with significant terminological gaps through the exploitation of partial alignments. It is based on the insight that an ontology will consistently use its own terminology. For instance, if an ontology uses the term *Paper* to refer to scientific articles, it is unlikely to use the equivalent term *Article* in the descriptions of other concepts instead, especially if the ontology is designed using a design principle that enforces this property [19]. However, if a partial alignment contains the correspondence *Paper-Article*, then one can use this insight to ones advantage. For instance, given the concept *Accept_Paper* a profile similarity is more likely to match it to its appropriate counterpart *Approve_Article* if the profile of *Accept_Paper* contains the term *'Article'*.

A partial alignment PA is a set of correspondences, with each correspondence asserting a semantic relation between two concepts of different ontologies. The types of relations modelled in a partial alignment, e.g., $\sqsupseteq$, $\perp$, $\sqcap$ and $\equiv$, are typically also modelled in an ontology and thus exploited in the construction of a profile. Thus, by semantically annotating the given ontologies O_1 and O_2 with the correspondences of PA it becomes possible to exploit these newly asserted relations for the creation of the concept profiles. This enables us to construct the profiles of O_1 using a subset of the terminology of O_2, increasing the probability of a terminological overlap between the profiles of two corresponding concepts. This idea is illustrated in Figure 2.

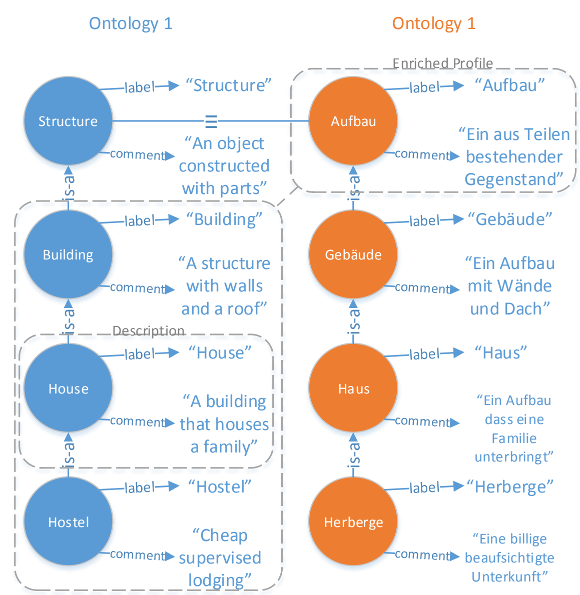

Figure 2. Two equivalent concepts being compared to a series of anchors.

Before we introduce our approach, we need to define a series terms and symbols that will be used in the following sections:

Correspondence A 5-tuple $< id, e_1, e_2, t, c >$ asserting the semantic relation t between entity $e_1 \in O_1$ and $e_2 \in O_2$ with a confidence of $c \in [0, 1]$.

Mapping/Alignment A set of correspondences, each asserting a relation between $e_1 \in O_1$ and $e_2 \in O_2$.

Partial Alignment A subset of an ideal alignment between ontologies O_1 and O_2.

Anchor A correspondence belonging to a partial alignment.

Collection of words: A list of unique words where each word has a corresponding weight in the form of a rational number.

+: Operator denoting the merging of two collections of words.

$\times$: Operator denoting element-wise multiplication of term frequencies with a weight.

depth(x): The taxonomy depth of concept x within its ontology.

D: The maximum taxonomical depth of a given ontology.

Next, it is necessary to provide a definition of a basic profile similarity upon which we can base our approach. For this, we provide a definition similar to the work by Mao et al. [10]. Neighbouring concepts are explored using a set of semantic relations, such as *isChildOf* or *isParentOf*. A base function of a profile similarity is the description of a concept, which gathers the literal information encoded for that concept. Let x be a concept of an ontology, the description $Des(x)$ of x is a collection of words defined as follows:

$$\begin{aligned} Des(x) \quad = \quad & \text{collection of words in the name of } x \\ & + \text{collection of words in the labels of } x \\ & + \text{collection of words in the comments of } x \\ & + \text{collection of words in the annotations of } x \end{aligned} \tag{1}$$

We define the profile of x as the merger of the description of x and the descriptions of semantically related concepts:

$$\begin{aligned} \textit{Profile}(x) = \quad & Des(x) + \textstyle\sum_{p \in P(x)} Des(p) + \\ & \textstyle\sum_{c \in C(x)} Des(c) + \sum_{r \in R(x)} Des(r) \end{aligned} \tag{2}$$

where

$$\begin{aligned} &P(x) = \{p | x \textit{ isChildOf } p\} \\ &C(x) = \{c | c \textit{ isChildOf } x\} \\ &R(x) = \{r | r \textit{ isRelatedTo } x \land r \notin P(x) \cup C(x)\} \end{aligned}$$

In order to compute the similarity between two profiles, they are parsed into a vector-space model and compared using the cosine similarity [20]. To bridge the terminological gap we aim to exploit the semantic relations provided by a given partial alignment PA, such that we can enhance the profile of a concept $x \in O_1$ using the terminology of O_2. We refer to this enlarged profile as the *anchor-enriched-profile*. For this, we explore the parents, children and properties of a concept x (or ranges and domains in case x itself is a property). If during this exploration a concept y is encountered which is mapped in a correspondence in PA to a concept $e \in O_2$, then *Profile*(x) is merged with $Des(e)$.

We will define the set that describes the extended collection of *parentally-anchored-descriptions* (PAD) with regard to concept x in three variations. These gather the descriptions of anchored concepts from the ancestors of x. To measure the improvement caused by the addition of these sets, we also define the omission of any such descriptions. They are defined as follows:

$$\begin{aligned} &\text{PAD}_0(x, PA) = \emptyset \\ &\text{PAD}_1(x, PA) = \textstyle\sum_{e \in E} Des(e); \ \textit{where} \\ &\quad E = \{e | \exists < id, y, e, t, c > \in PA; y \textit{ isAncestorOf } x\} \\ &\text{PAD}_2(x, PA) = \textstyle\sum_{e \in E} \omega \times Des(e); \ \textit{where} \\ &\quad E = \{e | \exists < id, y, e, t, c > \in PA; y \textit{ isAncestorOf } x\} \\ &\quad \land \omega = \frac{D - |depth(x) - depth(y)|}{D} \end{aligned} \tag{3}$$

An interesting point to note is that PAD_2 utilizes the same set of concepts than PAD_1, but weighs their descriptions with respect to the concept's relative distance to x, such that the descriptions of closer concepts receive a higher weight.

Exploring the children of x, we define the merged collection of *child-anchored-descriptions* (CAD) in a similar way:

$$\begin{aligned} &\text{CAD}_0(x, PA) = \emptyset \\ &\text{CAD}_1(x, PA) = \textstyle\sum_{e \in E} Des(e); \ \textit{where} \\ &\quad E = \{e | \exists < id, y, e, t, c > \in PA; y \textit{ isDescendantOf } x\} \\ &\text{CAD}_2(x, PA) = \textstyle\sum_{e \in E} \omega \times Des(e); \ \textit{where} \\ &\quad E = \{e | \exists < id, y, e, t, c > \in PA; y \textit{ isDescendantOf } x\} \\ &\quad \land \omega = \frac{D - |depth(x) - depth(y)|}{D} \end{aligned} \tag{4}$$

Lastly, we can explore the relations defined by the properties of the ontology, being *isDomainOf* and *isRangeOf*. Defining O_c as the set of concepts defined in ontology O and O_p as the set of properties of O, we define the merged collection of *relation-anchored-descriptions* (RAD) in two variations as follows:

$$\begin{aligned} &\text{RAD}_0(x, PA) = \emptyset \\ &\text{RAD}_1(x, PA) = \\ &\begin{cases} \sum_{e \in E} Des(e); \ \textit{where} \\ \quad E = \{e | \exists < id, y, e, t, c > \in PA; x \textit{ isDomainOf } y\} \\ \quad \textit{if } x \in O_c \\ \sum_{e \in E} Des(e); \ \textit{where} \\ \quad E = \{e | \exists < id, y, e, t, c > \in PA; y \textit{ isDomainOf } x \lor \\ \quad y \textit{ isRangeOf } x\} \textit{ if } x \in O_p \end{cases} \\ &\text{RAD}_2(x, PA) = \\ &\begin{cases} \sum_{e \in E \cup F} Des(e); \ \textit{where} \\ \quad E = \{e | \exists < id, y, e, t, c > \in PA; x \textit{ isDomainOf } y\} \\ \quad \textit{and } F = \{f | \exists < id, y, f, t, c > \in PA \ \exists z \in O_p; \\ \quad x \textit{ isDomainOf } z \land y \textit{ isRangeOf } z\} \textit{ if } x \in O_c \\ \sum_{e \in E} Des(e); \ \textit{where} \\ \quad E = \{e | \exists < id, y, e, t, c > \in PA; y \textit{ isDomainOf } x \lor \\ \quad y \textit{ isRangeOf } x\} \textit{ if } x \in O_p \end{cases} \end{aligned} \tag{5}$$

The noteworthy difference between RAD_1 and RAD_2 is that if x is a concept and the domain of property z, then every range y of z will be explored as well. As an example, assume we are given the concepts *Car* and *Driver* being linked by the property *ownedBy*. Constructing the anchor-enriched-profile of *Car* using the set RAD_1 would mean that we only investigate if *ownedBy* is mapped in PA. Using RAD_2 means we also investigate *Driver*, which could provide additional context.

Given a partial alignment PA between ontologies O_1 and O_2, and given a concept x, we define the *anchor-enriched-profile* of x as follows:

$$\begin{aligned} &\textit{Profile}^{AE}_{\kappa,\lambda,\mu}(x, PA) = \textit{Profile}(x) + \text{PAD}_\kappa(x, PA) + \\ &\text{CAD}_\lambda(x, PA) + \text{RAD}_\mu(x, PA) \end{aligned} \tag{6}$$

V. Experiments

In this section, we will detail the performed experiments to test the effectiveness of our approach and discuss the obtained results. A widely used way of evaluating a mapping A is by comparing it to a reference alignment R by calculating the standard measures of Precision, Recall and F-Measure [21]. When matching with a partial alignment PA the newly computed correspondences are typically merged with PA in order to create a complete mapping. However, this action creates a bias with respect to the measured alignment quality.

For instance, if PA contains half the correspondences of R, then the resulting Recall score cannot go below 0.5. It would hence be desirable to use a measure which only focuses on the correspondences that are contributed to PA to get a better indication of their quality. To achieve this, we will use adapted variants of Precision, Recall and F-Measure, which take the presence of a partial alignment into account. Given a computed alignment A, a reference alignment R and a partial alignment PA, the adapted measures of Precision and Recall are computed as follows:

$$P^*(A, R, PA) = \frac{\mid A \cap R \cap \overline{PA} \mid}{\mid A \cap \overline{PA} \mid} \tag{7}$$

$$R^*(A, R, PA) = \frac{\mid A \cap R \cap \overline{PA} \mid}{\mid R \cap \overline{PA} \mid} \tag{8}$$

The adapted F-Measure can then be computed as follows:

$$F^*(A, R, PA) = \frac{2 * P^*(A, R, PA) * R^*(A, R, PA)}{P^*(A, R, PA) + R^*(A, R, PA)} \tag{9}$$

A. *Multi-Farm*

In this section we will present the results of our evaluation on the *Multi-Farm-sameOnto* dataset. This data-set stems from the OAEI 2014 [21] competition. The terminologies of the ontologies in this dataset vary greatly since it is designed to be a cross-lingual dataset. The set consists of 8 ontologies that are modelled using 9 languages (including English). For each pair of ontologies a set of mapping tasks exists consisting of every possible combination of selecting different languages. We generate the partial alignments by randomly sampling the reference alignment with the condition that $R(PA, R) = 0.5$ and aggregate the results of 100 evaluations for each task. This evaluation is repeated for every possible combination of κ, λ and μ. The result of this evaluation is presented in Table I.

Table I. AGGREGATED ADAPTED PRECISION, RECALL AND F-MEASURE FOR ALL VARIATIONS ON THE MULTI-FARM DATASET.

κ	λ	μ	P^*	R^*	F^*
0	0	0	0.418	0.278	0.326
0	0	1	0.657	0.433	0.510
0	0	2	0.630	0.405	0.481
0	1	0	0.500	0.324	0.381
0	1	1	0.675	0.469	0.543
0	1	2	0.666	0,453	0.529
0	2	0	0.512	0.333	0.393
0	2	1	0.688	0.475	0.552
0	2	2	0.678	0.457	0.535
1	0	0	0.521	0.376	0.423
1	0	1	0.667	0.529	0.583
1	0	2	0.659	0.518	0.574
1	1	0	0.594	0.409	0.470
1	1	1	0.691	0.559	0.611
1	1	2	0.688	0.555	0.609
1	2	0	0.601	0.417	0.478
1	2	1	**0.699**	0.565	0.619
1	2	2	0.695	0.562	0.615
2	0	0	0.523	0.385	0.433
2	0	1	0.674	0.538	0.592
2	0	2	0.661	0.522	0.577
2	1	0	0.591	0.411	0.471
2	1	1	0.690	0.562	0.614
2	1	2	0.685	0.554	0.607
2	2	0	0.597	0.421	0.481
2	2	1	0.698	**0.570**	**0.622**
2	2	2	0.692	0.562	0.614

First, by comparing the performance of the baseline configuration $Profile^{AE}_{0,0,0}$ to any configuration of our approach we can easily see that our approach improves upon the performance of the baseline. Adding the sets PAD or CAD using either variation typically resulted in an F-Measure of 0.39-0.43, an improvement of 0.07 to 0.11 when compared to the baseline. Curiously, enriching the profiles using RAD alone typically resulted in a F^* score of approximately 0.5. This could indicate that for this dataset the concept annotations more often contain terms of related concepts than ancestors or descendants.

Looking at dual-combinations between PAD, CAD and RAD we can see a consistent increase in performance. Of these combinations, $Profile^{AE}_{1,1,0}$ resulted in the lowest F-Measure of 0.47, while $Profile^{AE}_{1,0,1}$ resulted in the highest F-Measure of 0.583. We can also observe that combinations which include a variation of the RAD-set in the enriched profiles typically performed better than combinations that didn't.

Lastly, we can observe using all three types of description sets resulted in the highest measured F^* score. We can see that every combination of PAD, CAD and RAD resulted in an F^* score higher than 0.6. The best performing combination was $Profile^{AE}_{2,2,1}$ with an F^* score of 0.622.

Comparing RAD_1 with RAD_2 reveals that combinations which utilized RAD_1 performed slightly better than combinations which used RAD_2 instead. This implies that concepts which are related through properties are less likely to share terms, leading to many impact-less terms being added to the concept profiles.

B. *Comparison with Lexical Enrichment Systems*

The main goal behind this work is to provide an approach that allows the enrichment of concept profile by exploiting the relations of a provided partial alignment. The reason behind this is that current enrichment methods exploit primarily lexical resources, which rely on the presence of an appropriate resource. In the previous sections, we have established the performance of our approach using different configurations, datasets, and partial alignment sizes. In this section, we will provide some interesting context for these results. Specifically, we aim to compare the results of our approach with the performances of matching systems tackling the same dataset while exploiting lexical resources. This allows us to establish whether an approach exploiting a partial alignment can produce alignments of similar quality as approaches exploiting lexical resources. To do this, we will compare the performance of our approach on the Multi-Farm dataset [21] to the performances of the OAEI participants which competed in the 2014 evaluation. Here we will make the distinction between approaches utilizing no external resources, lexical resources and partial alignments. This allows us to see the benefit of exploiting a given type of external resource.

Furthermore, to provide an upper boundary for the potential performance on this dataset, we will also evaluate a method utilizing both lexical resources and partial alignments. To achieve this, we will re-evaluate the best performing configuration from sub-section V-A. However, the profiles of this re-evaluation will be additionally enriched by translating the concept labels using the *Microsoft Bing* translator. This will provide an indication of how well a system may perform when utilizing both appropriate lexical resources and partial alignments. The comparison can be seen in Table II. Performances of approaches utilizing partial alignments are denoted in adapted precision, recall and F-Measure.

Table II. PERFORMANCE COMPARISON BETWEEN OUR APPROACH AND THE OAEI 2014 COMPETITORS (MULTI-FARM DATASET).

Lex.	P. Align.	Matcher	Precision	Recall	F-Measure
yes	yes	$Profile^{AE}_{2,2,1}$ + Bing	*0.849*	*0.838*	*0.843*
yes	no	AML	0.95	0.48	0.62
yes	no	LogMap	0.94	0.27	0.41
yes	no	XMap	0.76	0.40	0.50
no	yes	$Profile^{AE}_{2,2,1}$	*0.698*	*0.570*	*0.622*
no	no	AOT	0.11	0.12	0.12
no	no	AOTL	0.27	0.01	0.02
no	no	LogMap-C	0.31	0.01	0.02
no	no	LogMapLt	0.25	0.01	0.02
no	no	MaasMatch	0.52	0.06	0.10
no	no	RSDLWB	0.34	0.01	0.02

From Table II, we can make several observations. First, we can observe that every system utilizing either lexical resources or partial alignments performs significantly better than systems which do not. This is an expected result given the nature of this dataset. Of the system which do not exploit resources AOT has the highest performance with an F-Measure of 0.12.

Comparing the performance of $Profile^{AE}_{2,2,1}$ to the performance of system exploiting only lexical resources reveals an interesting observation. Specifically, we can see that the performance of these systems is comparable. While the performances of LogMap and XMap were lower than $Profile^{AE}_{2,2,1}$, with an F-Measure of 0.62 the performance of AML is very close to the performance of $Profile^{AE}_{2,2,1}$. However, AML distinguishes itself from our approach by having a notably higher precision and a somewhat lower recall. In fact, all systems utilizing only lexical resources are characterized with a high precision, which implies that enriching ontologies using these resources only rarely leads to false-positive matches in terminology.

Lastly, we can observe the performance of our approach when paired with a lexical resource, specifically *Bing Translator*. The produced alignments reached an F^* score of 0.843, which is significantly higher than the OAEI participants. This implies that the correct correspondences which lexical-based systems find differ significantly from the correct correspondences of a partial-alignment-based system. From this we can conclude that the two types of resources are complementary for matching problems with significant terminological gaps.

VI. Conclusion

In this paper, we presented a technique aimed at tackling ontology mapping problems with significant terminological heterogeneities between the given ontologies. This technique exploits an existing partial alignment by enriching the given ontologies with the relations asserted in the partial alignment. We establish the performance of the approach on a dataset characterized by terminological heterogeneous mapping problems. A comparison with other matching systems reveals that the approach performs similarly to systems utilizing lexical resources. Combining our approach with a lexical resource reveals that a significantly higher performance can be achieved if both partial alignments and lexical resources are utilized.

References

[1] F. Bellifemine, A. Poggi, and G. Rimassa, "Jade–a fipa-compliant agent framework," in Proceedings of PAAM, vol. 99. London, 1999, p. 33.

[2] Y. Lei, V. Uren, and E. Motta, "Semsearch: A search engine for the semantic web," in Managing Knowledge in a World of Networks. Springer, 2006, pp. 238–245.

[3] J. Euzenat, "Towards a principled approach to semantic interoperability," in Proceedings of the IJCAI-01 Workshop on Ontologies and Information Sharing, 2001, pp. 19–25.

[4] E. Jiménez-Ruiz and B. Cuenca Grau, "Logmap: logic-based and scalable ontology matching," in ISWC 2011, 2011, pp. 273–288.

[5] F. C. Schadd and N. Roos, "Anchor-profiles: Exploiting profiles of anchor similarities for ontology mapping," in Proceedings of the 26th Belgian-Dutch Conference on Artificial Intelligence (BNAIC 2014), 2014, pp. 177–178.

[6] Y. Qu, W. Hu, and G. Cheng, "Constructing virtual documents for ontology matching," in Proceedings of the 15th international conference on World Wide Web, ser. WWW '06. ACM, 2006, pp. 23–31.

[7] I. Cruz, F. Antonelli, and C. Stroe, "Agreementmaker: efficient matching for large real-world schemas and ontologies," Proc. VLDB Endow., vol. 2, no. 2, Aug. 2009, pp. 1586–1589.

[8] J. Li, J. Tang, Y. Li, and Q. Luo, "Rimom: A dynamic multistrategy ontology alignment framework," IEEE Transactions on Knowledge and Data Engineering, vol. 21, no. 8, 2009, pp. 1218–1232.

[9] D. Ngo, Z. Bellahsene, and R. Coletta, "Yam++-a combination of graph matching and machine learning approach to ontology alignment task," Journ. of Web Sem., 2012.

[10] M. Mao, Y. Peng, and M. Spring, "A profile propagation and information retrieval based ontology mapping approach," in SKG 2007. IEEE, 2007, pp. 164–169.

[11] F. C. Schadd and N. Roos, "Anchor-profiles for ontology mapping with partial alignments," in Proceedings of the 12th Scandinavian AI Conference (SCAI 2013), 2013, pp. 235–244.

[12] D. Spohr, L. Hollink, and P. Cimiano, "A machine learning approach to multilingual and cross-lingual ontology matching," in ISWC 2011. Springer, 2011, pp. 665–680.

[13] X. Su and J. A. Gulla, "Semantic enrichment for ontology mapping," in Proc. of the 9th International Conference on Applications of Natural Languages to Information Systems, NLDB 2004, vol. 3136, 2004, p. 217.

[14] C. D. Manning, P. Raghavan, and H. Schütze, Introduction to information retrieval. Cambridge university press Cambridge, 2008, vol. 1.

[15] J. Euzenat and P. Shvaiko, Ontology Matching. Springer Berlin, 2007, vol. 18.

[16] G. A. Miller, "Wordnet: a lexical database for english," Communications of the ACM, vol. 38, November 1995, pp. 39–41.

[17] O. Bodenreider, "The unified medical language system (umls): integrating biomedical terminology," Nucleic acids research, vol. 32, no. suppl 1, 2004, pp. D267–D270.

[18] R. Navigli, "Word sense disambiguation: A survey," ACM Comput. Surv., vol. 41, no. 2, Feb. 2009, pp. 10:1–10:69.

[19] Y. Sure, S. Staab, and R. Studer, "Methodology for development and employment of ontology based knowledge management applications," ACM SIGMOD Record, vol. 31, no. 4, 2002, pp. 18–23.

[20] P.-N. Tan, M. Steinbach, and V. Kumar, Introduction to Data Mining, 1st ed. Addison Wesley, May 2005.

[21] Z. Dragisic, K. Eckert, J. Euzenat, D. Faria, A. Ferrara, R. Granada, V. Ivanova, E. Jimenez-Ruiz, A. Kempf, P. Lambrix et al., "Results of theontology alignment evaluation initiative 2014," in OM 2014, 2014, pp. 61–104.