Evaluating Ontology Alignment Systems in Query Answering Tasks

Alessandro Solimando[1], Ernesto Jiménez-Ruiz[2], and Christoph Pinkel[3]

[1] Dipartimento di Informatica, Bioingegneria, Robotica e Ingegneria dei Sistemi, Università di Genova, Italy
[2] Department of Computer Science, University of Oxford, UK
[3] fluid Operations AG, Walldorf, Germany

Abstract. Ontology matching receives increasing attention and gained importance in more recent applications such as ontology-based data access (OBDA). However, query answering over aligned ontologies has not been addressed by any evaluation initiative so far. A novel Ontology Alignment Evaluation Initiative (OAEI) track, Ontology Alignment for Query Answering (OA4QA), introduced in the 2014 evaluation campaign, aims at bridging this gap in the practical evaluation of matching systems w.r.t. this key usage.

1 Introduction

Ontologies play a key role in the development of the Semantic Web and are being used in many application domains such as biomedicine and energy industry. An application domain may have been modeled with different points of view and purposes. This situation usually leads to the development of different ontologies that intuitively overlap, but they use different naming and modeling conventions.

The problem of (semi-)automatically computing mappings between independently developed ontologies is usually referred to as the *ontology matching problem*. A number of sophisticated ontology matching systems have been developed in the last years [5]. Ontology matching systems, however, rely on lexical and structural heuristics and the integration of the input ontologies and the mappings may lead to many undesired logical consequences. In [1] three principles were proposed to minimize the number of potentially unintended consequences, namely: *(i) consistency principle*, the mappings should not lead to unsatisfiable classes in the integrated ontology; *(ii) locality principle*, the mappings should link entities that have similar *neighbourhoods*; *(iii) conservativity principle*, the mappings should not introduce alterations in the classification of the input ontologies. The occurrence of these violations is frequent, even in the reference mapping sets of the Ontology Alignment Evaluation Initiative[4] (*OAEI*) [6].

Violations to these principles may hinder the usefulness of ontology mappings. The practical effect of these violations, however, is clearly evident when ontology alignments are involved in complex tasks such as query answering [4].

[4] `http://oaei.ontologymatching.org/`

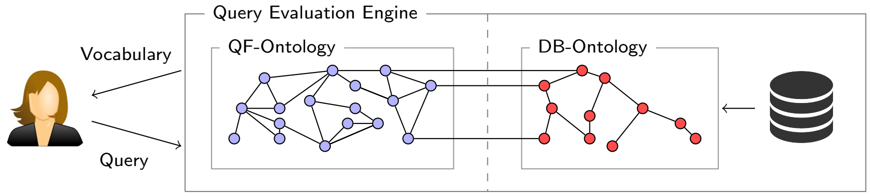

Fig. 1. Ontology Alignment in an OBDA Scenario

The traditional tracks of *OAEI* evaluate ontology matching systems w.r.t. scalability, multi-lingual support, instance matching, reuse of background knowledge, etc. Systems' effectiveness is, however, only assessed by means of classical information retrieval metrics (*i.e.*, precision, recall and f-measure) w.r.t. a manually-curated reference alignment, provided by the organisers. The new *OA4QA* track[5] evaluates those same metrics, but w.r.t. the ability of the generated alignments to enable the answer of a set of queries in an OBDA scenario, where several ontologies exist. Figure 1 shows an OBDA scenario where the first ontology provides the vocabulary to formulate the queries (QF-Ontology) and the second is linked to the data and it is not visible to the users (DB-Ontology). Such OBDA scenario is presented in real-world use cases (*e.g.*, Optique project[6] [2, 6]). The integration via ontology alignment is required since only the vocabulary of the DB-Ontology is connected to the data. The *OA4QA* will also be key for investigating the effects of logical violations affecting the computed alignments, and evaluating the effectiveness of the repair strategies employed by the matchers.

2 Ontology Alignment for Query Answering

This section describes the considered dataset and its extensions (Section 2.1), the query processing engine (Section 2.2), and the evaluation metrics (Section 2.3).

2.1 Dataset

The set of ontologies coincides with that of the *conference* track,[7] in order to facilitate the understanding of the queries and query results. The dataset is however extended with synthetic ABoxes, extracted from the *DBLP* dataset.[8]

Given a query q expressed using the vocabulary of ontology $\mathcal{O}_1$, another ontology $\mathcal{O}_2$ enriched with syntethic data is chosen. Finally, the query is executed over the aligned ontology $\mathcal{O}_1 \cup \mathcal{M} \cup \mathcal{O}_2$, where $\mathcal{M}$ is an alignment between $\mathcal{O}_1$ and $\mathcal{O}_2$. Referring to Figure 1, $\mathcal{O}_1$ plays the role of QF-Ontology, while $\mathcal{O}_2$ that of DB-Ontology.

[5] `http://www.cs.ox.ac.uk/isg/projects/Optique/oaei/oa4qa/`
[6] `http://www.optique-project.eu/`
[7] `http://oaei.ontologymatching.org/2014/conference/index.html`
[8] `http://dblp.uni-trier.de/xml/`

2.2 Query Evaluation Engine

The evaluation engine considered is an extension of the OWL 2 reasoner *HermiT*, known as *OWL-BGP*[9] [3]. OWL-BGP is able to process SPARQL queries in the SPARQL-OWL fragment, under the *OWL 2 Direct Semantics entailment regime*.[10] The queries employed in the *OA4QA* track are standard conjunctive queries, that are fully supported by the more expressive SPARQL-OWL fragment. SPARQL-OWL, for instance, also support queries where variables occur within complex class expressions or bind to class or property names.

2.3 Evaluation Metrics and Gold Standard

As already discussed in Section 1, the evaluation metrics used for the *OA4QA* track are the classic information retrieval ones (*i.e.*, precision, recall and f-measure), but on the result set of the query evaluation. In order to compute the gold standard for query results, the publicly available reference alignments *ra1* has been manually revised. The aforementioned metrics are then evaluated, for each alignment computed by the different matching tools, against the *ra1*, and manually repaired version of *ra1* from conservativity and consistency violations.

Three categories of queries will be considered in *OA4QA*: *(i)* basic, *(ii)* queries involving violations, *(iii)* advanced queries involving nontrivial mappings.

2.4 Impact of the Mappings in the Query Results

As an illustrative example, consider the aligned ontology $\mathcal{O}_{\mathcal{U}}$ computed using *confof* and *ekaw* as input ontologies ($\mathcal{O}_{confof}$ and $\mathcal{O}_{ekaw}$, respectively), and the *ra1* reference alignment between them. $\mathcal{O}_{\mathcal{U}}$ entails $ekaw{:}Student \sqsubseteq ekaw{:}Conf\_Participant$, while O_{ekaw} does not, and therefore this represents a conservativity principle violation. Clearly, the result set for the query $q(x) \leftarrow ekaw{:}Conf\_Participant(x)$ will erroneously contain any student not actually participating at the conference. The explanation for this entailment in $\mathcal{O}_{\mathcal{U}}$ is given below, where Axioms 1 and 3 are mappings from the reference alignment.

$$confof{:}Scholar \equiv ekaw{:}Student \tag{1}$$

$$confof{:}Scholar \sqsubseteq confof{:}Participant \tag{2}$$

$$confof{:}Participant \equiv ekaw{:}Conf\_Participant \tag{3}$$

The softening of Axiom 3 into $confof{:}Participant \sqsupseteq ekaw{:}Conf\_Participant$ represents a possible repair for the aforementioned violation.

3 Preliminary Evaluation

In Table 1 [11] a preliminary evaluation using the alignments of the *OAEI 2013* participants and the following queries is shown: *(i)* $q_1(x) \leftarrow ekaw{:}Author(x)$,

[9] https://code.google.com/p/owl-bgp/

[10] http://www.w3.org/TR/2010/WD-sparql11-entailment-20100126/#id45013

[11] $\#q(x)$ refers to the cardinality of the result set.

Category	Query	$\#\mathcal{M}$	Reference Alignment				Repaired Alignment			
			#q(x)	Prec.	Rec.	F-meas.	#q(x)	Prec.	Rec.	F-meas.
Basic	q_1	5	98	1	1	1	98	1	1	1
Violations	q_2	4	53	0.8	1	0.83	38	0.57	1	0.68
Advanced	q_3	7	-	-	-	-	182	1	0.5	0.67

Table 1. Preliminary query answering results for the OAEI 2013 alignments

over the ontology pair $\langle cmt, ekaw\rangle$; *(ii)* $q_2(x) \leftarrow ekaw{:}Conf\_Participant(x)$, over $\langle confof, ekaw\rangle$, involving the violation described in Section 2.4; *(iii)* and $q_3(x) \leftarrow confof{:}Reception(x) \cup confof{:}Banquet(x) \cup confof{:}Trip(x)$, over $\langle confof, edas\rangle$. The evaluation[12] shows the negative effect on precision of logical flaws affecting the computed alignments (q_2) and a lowering in recall due to missing mapping (q_3). For q_3 the results w.r.t. the reference alignment (*ra1*) are missing due to the unsatisfiability of the aligned ontology $\mathcal{O}_{confof} \cup \mathcal{O}_{edas} \cup ra1$.

4 Conclusions and Future Work

We have presented the novel OAEI track addressing query answering over pairs of ontologies aligned by a set of ontology-to-ontology mappings. From the preliminary evaluation the main limits of the traditional evaluation, for what concerns logical violations of the alignments, clearly emerged. As a future work we plan to cover increasingly complex queries and ontologies, including the ones in the Optique use case [6]. We also plan to consider more complex scenarios involving a single QF-Ontology aligned with several DB-Ontologies.

Acknowledgements. This work was supported by the EU FP7 IP project Optique (no. 318338), the MIUR project CINA (Compositionality, Interaction, Negotiation, Autonomicity for the future ICT society) and the EPSRC project Score!.

References

1. Jiménez-Ruiz, E., Cuenca Grau, B., Horrocks, I., Berlanga, R.: Logic-based Assessment of the Compatibility of UMLS Ontology Sources. J. Biomed. Semant. (2011)
2. Kharlamov, E., et al.: Optique 1.0: Semantic Access to Big Data: The Case of Norwegian Petroleum Directorate's FactPages. ISWC (Posters & Demos) (2013)
3. Kollia, I., Glimm, B., Horrocks, I.: SPARQL query answering over OWL ontologies. In: The Semantic Web: Research and Applications, pp. 382–396. Springer (2011)
4. Meilicke, C.: Alignments Incoherency in Ontology Matching. Ph.D. thesis, University of Mannheim (2011)
5. Shvaiko, P., Euzenat, J.: Ontology Matching: State of the Art and Future Challenges. IEEE Transactions on Knowl. and Data Eng. (TKDE) (2012)
6. Solimando, A., Jiménez-Ruiz, E., Guerrini, G.: Detecting and Correcting Conservativity Principle Violations in Ontology-to-Ontology Mappings. In: International Semantic Web Conference (2014)

[12] Out of the 26 alignments of OAEI 2013, only the ones shown in column $\#\mathcal{M}$ were able to produce a result (either for logical problems or for an empty result set due to missing mappings). Reported precision/recall values are averaged values.