

CANARD: An approach for generating expressive correspondences based on competency questions for alignment

Elodie Thiéblin, Guilherme Sousa, Olivier Haemmerlé and Cássia Trojahn *

Institut de Recherche en Informatique de Toulouse, Toulouse, France

E-mails: elodie@thieblin.fr, guilherme.santos-sousa@irit.fr, ollivier.haemmerle@irit.fr, cassia.trojahn@irit.fr

Editor: Jérôme Euzenat, University Grenoble Alpes & INRIA, France

Solicited review: Three anonymous reviewers

Abstract. Ontology matching aims at making ontologies interoperable. While the field has fully developed in the last years, most approaches are still limited to the generation of simple correspondences. More expressiveness is, however, required to better address the different kinds of ontology heterogeneities. This paper presents CANARD (Complex Alignment Need and A-box based Relation Discovery), an approach for generating expressive correspondences that rely on the notion of competency questions for alignment (CQA). A CQA expresses the user knowledge needs in terms of alignment and aims at reducing the alignment space. The approach takes as input a set of CQAs as SPARQL queries over the source ontology. The generation of correspondences is performed by matching the subgraph from the source CQA to the similar surroundings of the instances from the target ontology. Evaluation is carried out on both synthetic and real-world datasets. The impact of several approach parameters is discussed. Experiments have showed that CANARD performs, overall, better on CQA coverage than precision and that using existing same : As links, between the instances of the source and target ontologies, gives better results than exact label matches of their labels. The use of CQA improved also both CQA coverage and precision with respect to using automatically generated queries. The reassessment of the counter-example increased significantly the precision, to the detriment of runtime. Finally, experiments on large datasets showed that CANARD is one of the few systems that can perform on large knowledge bases, but depends on regularly populated knowledge bases and the quality of instance links.

Keywords: Ontology matching, complex alignment, competency question for alignment, user needs

1. Introduction

Ontology matching is the task of generating a set of correspondences between the entities of different ontologies. This is the basis for a range of other tasks, such as data integration, ontology evolution, or query rewriting. While the field has fully developed in the last decades, most works are still dedicated to the generation of simple correspondences (i.e., those linking one single entity of a source ontology to one single entity of a target ontology). However, simple correspondences are insufficient for covering the different kinds of ontology heterogeneities. More expressiveness is achieved by complex correspondences, which can better express the relationships between entities

* Corresponding author. E-mail: cassia.trojahn@irit.fr.

of different ontologies. For example, the piece of knowledge that a conference paper has been accepted can be represented as (i) a class IRI in a source ontology, or (ii) as a class expression representing the papers having a decision of type acceptance in a target ontology. The correspondence $\langle \text{ekaw:Accepted_Paper}, \exists \text{cmt:hasDecision.cmt:Acceptance}, \equiv, \rangle$ ¹ expresses an equivalence between the two representations of “accepted paper”.

Earlier works in the field have introduced the need for expressive alignments [15,34], and different approaches for generating them have been proposed in the literature afterwards. These approaches rely on diverse methods, such as correspondence patterns [9,22,23], knowledge-rules [13] and association rules [40], statistical methods [18,35], genetic programming [19] or still path-finding algorithms [20]. The reader can refer to [27] for a survey on complex matching approaches. All these proposals, however, intend to cover the full common scope of the ontologies and often need a large number of common instances.

While the matching space for generating complex correspondences is not $\mathcal{O}(mn)$ as for the simple alignment generation (m and n being respectively the number of entities of the source and target ontologies), but higher than $\mathcal{O}(2^{mn})$, a space reduction strategy can be based upon on two assumptions. First, it may be the case that the user does not need the alignment to cover the full scope of the ontologies. Focusing on the user’s needs can reduce the search space. Reducing the search space definitely impacts the matching performance, in particular when dealing with large knowledge bases. The second assumption is that for each knowledge need, the ontologies share at least one instance.

This paper presents CANARD (Complex Alignment Need and A-box based Relation Discovery), a system that discovers expressive correspondences between populated ontologies based on the notion of Competency Questions for Alignment (CQAs). Correspondences involving logical constructions. Correspondences involving transformation functions are out of the scope of this paper. CQAs represent the user knowledge needs and define the scope of the alignment. They are competency questions that need to be satisfied over two ontologies. CANARD takes as input a set of CQAs translated into SPARQL queries over the source ontology. The answer to each query is a set of instances retrieved from a knowledge base described by the source ontology. These instances are matched with those of a knowledge base described by the target ontology. The generation of the correspondence is performed by matching the subgraph from the source CQA to the lexically similar surroundings of the target instances.

The main contributions of the paper can be summarised as follows: (i) detailing a scalable CQA-based matching approach able to generate complex correspondences involving logical constructors; (ii) discussing the impact of the different design and implementation choices (CQAs vs. queries, reassessment with counter-examples, etc.); (iii) evaluating the approach on both synthetic and real-world benchmarks; and (iv) comparing the proposed approach to state-of-the-art systems. The paper extends the work in [29] in several directions: (i) providing a deeper description of the steps of the approach; (ii) discussing the impact of the different design and implementation choices; (iii) extending the comparison of the approach to systems participating in the most recent OAEI campaigns; and (iv) presenting an optimized version of the system that improves runtime. The source code is available² under the GNU Lesser General Public License v2.1.

The rest of this paper is organized as follows. The next section introduces ontology matching and CQA (Section 2), followed by an overview of the proposed approach (Section 3). The details of the approach are then presented (Section 4). Next, the experiments are presented (Section 5), followed by a discussion on the main related work (Section 7). Finally, the conclusions and future work end the paper (Section 8).

2. Foundations

2.1. Complex ontology alignment

Ontology matching (as in [7]) is defined as the process of generating an alignment A between two ontologies: a source ontology o and a target ontology o' . A is a set of correspondences $\langle e_1, e_2, r, n \rangle$. Each correspondence expresses a relation r (e.g., equivalence (\equiv), subsumption (\sqsubseteq , \sqsupseteq)) between two members e_1 and e_2 , and n expresses

¹Example from the ontologies in the OAEI Conference Dataset: <https://oaei.ontologymatching.org/2022/conference/index.html>.

²<https://gitlab.irit.fr/melodi/ontology-matching/complex/canard>

its level of confidence $[0..1]$. A member can be a single ontology entity (class, object property, data property, individual) of respectively o and o' or a complex construction which is composed of entities using constructors. Two kinds of correspondences are considered depending on the type of their members:

- a correspondence is **simple** if both e_1 and e_2 are single entities (IRIs): $\langle o:Paper, o':Paper, \equiv, I \rangle$;
- a correspondence is **complex** if at least one of e_1 or e_2 involves a constructor: $\langle o:Accepted_Paper, \exists o':hasDecision.o':Acceptance, \equiv, I \rangle$.

A simple correspondence is noted (s:s), and a complex correspondence (s:c) if its source member is a single entity, (c:s) if its target member is a single entity, or (c:c) if both members are complex entities. An approach that generates complex correspondences is referred to as a “complex approach” or “complex matcher” below.

2.2. Competency questions for alignment (CQAs)

In ontology authoring, to formalize the knowledge needs of an ontology, competency questions (CQs) have been introduced as *ontology's requirements in the form of questions the ontology must be able to answer* [11]. A competency question for alignment (CQA) is a competency question which should (in the best case) be covered by two ontologies, *i.e.*, it expresses the knowledge that an alignment should cover if both ontologies' scopes can answer the CQA. The first difference between a CQA and a CQ is that the scope of the CQA is limited by the intersection of its source and target ontologies' scopes. The second difference is that this maximal and ideal alignment's scope is not known *a priori*. As CQs [21], a CQA can be expressed in natural language or as SPARQL queries. Inspired from the predicate arity in [21], the notion of **question arity**, which represents the arity of the expected answers to a CQA is adapted, as introduced in [28]:

- a *unary* question expects a set of instances or values, *e.g.*, *Which are the accepted papers?* (*paper1*), (*paper2*);
- a *binary* question expects a set of instances or value pairs, *e.g.*, *What is the decision of which paper?* (*paper1*, *accept*), (*paper2*, *reject*); and
- an *n-ary* question expects a tuple of size 3 or more, *e.g.*, *What is the rate associated with which review of which paper?* (*paper1*, *review1*, *weak accept*), (*paper1*, *review2*, *reject*).

In CANARD, CQAs are limited to *unary* and *binary*, of *select* type, and no modifier. This is a limitation in the sense that it does not deal with specific kinds of SPARQL queries, such as the ones involving CONSTRUCT and ASK. The approach does not deal with transformation functions or filters inside the SPARQL queries and only accepts queries with one or two variables. However, as classes and properties are unary and binary predicates, these limitations still allow the approach to cover ontology expressiveness. Questions with a binary or counting type have a corresponding selection question. For example, the question *Is this paper accepted?* has a binary type: its answers can only be *True* or *False*. The question *How many papers are accepted?* is a counting question. These two questions have the same selection question: *What are the accepted papers?*. We also restrain the question polarity to *positive* because a negative question implies that positive information is being negated. For example, the question *Which people are not reviewers?* is a negation of the question *Who are the reviewers?*. The CQA we consider has no modifier. The question arity of the CQA is limited to *unary* and *binary* because ontologies are mostly constructed using unary predicates (classes or class expressions) and binary predicates (object or data properties).

3. Overview of CANARD

CANARD takes as input a set of CQA in the form of SPARQL SELECT queries over the source ontology. It requires that the source and target ontologies have an *Abox* with at least one common instance for each CQA. The answer to each input query is a set of instances, which are matched with those of a knowledge base described by the target ontology. The matching is performed by finding the surroundings of the target instances which are lexically similar to the CQA. The idea behind the approach is to rely on a few examples (answers) to find a generic rule that

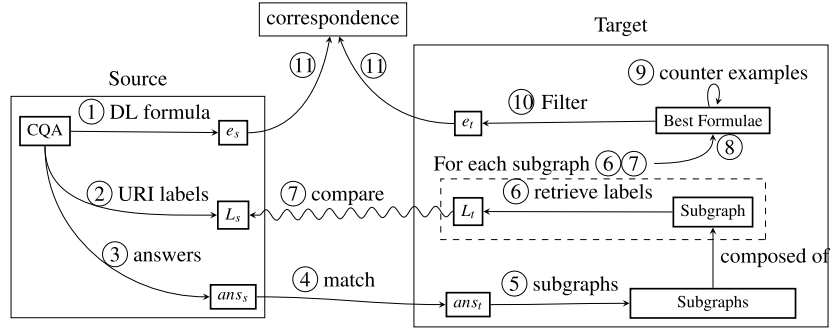


Fig. 1. Schema of the general approach.

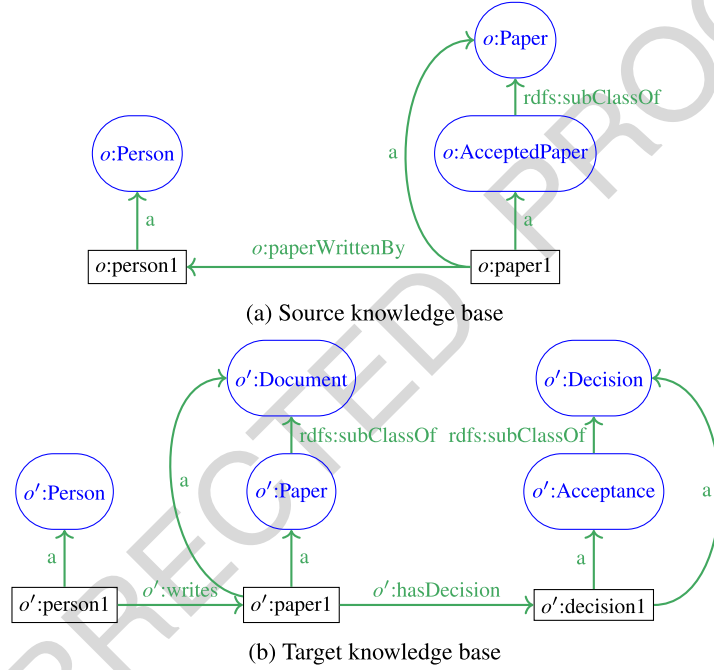


Fig. 2. Source and target knowledge bases.

describes more instances. The assumption that the user knows the source ontology and is able to write each CQA into a SPARQL query on the source ontology is made.

The overall approach is articulated in 11 steps, as depicted in Figure 1. The approach is based on **subgraphs** which are a **set of triples** for a unary CQA and a **property path** for a binary CQA. A lexical layer comparison is used to measure the similarity of the subgraphs with the CQA.

In the remainder of the paper, the examples consider the knowledge bases in Figure 2. They share common instances: $o:person1$ and $o':person1$, $o:paper1$ and $o':paper1$. Ontology o represents the concept of *accepted paper* as a class while o' models the same knowledge with a *has decision* property. The property *paper written by* is represented by a single property in o while in o' , the property *writes* links a *person* to a *document*. A criticism of this example could be that two knowledge bases may not represent the same conference, therefore they may not share common paper instances. However, these bases may have a different but overlapping scope. For example o could focus on the event organization part of a conference and o' on reviewer management. Before detailing the main steps of the approach, we instantiate the overall approach to deal with unary and binary queries.

```

SELECT ?x WHERE {
  ?x a ol:AcceptedPaper.
}

```

(a) Source unary CQA.

```

SELECT ?x ?y WHERE {
  ?x ol:paperWrittenBy ?y.
}

```

(b) Source binary CQA.

Fig. 3. Source CQAs.

3.1. Approach over a unary CQA

In the following, Figure 1 is instantiated for a unary CQA. The SPARQL CQA is that of Figure 3(a):

- ① Represent the SPARQL CQA as a DL formula e_s (e.g., $o:AcceptedPaper$) (Section 4.1).
- ② Extract lexical information from the CQA, L_s set labels of entities from the CQA (e.g., “accepted paper”).
- ③ Retrieve source answers ans_s of the CQA (e.g., $o:paper1$).
- ④ Find equivalent or similar target answers ans_t to the source instances ans_s (e.g. $o:paper1 \sim o':paper1$) (Section 4.2).
- ⑤ Extract the subgraphs of target answers (Section 4.3): for a unary query, this is the set of triples in which the target instances appear as well as the types (classes) of the subject or object of the triple (e.g. in DL, the description of $o':paper1$ would contain $\langle o':paper1, o':hasDecision, o':decision1 \rangle$, $\langle o':decision1, rdf:type, o':Decision \rangle$ and $\langle o':decision1, rdf:type, o':Acceptance \rangle$).
- ⑥ For each subgraph, retrieve L_t the labels of its entities (e.g., $o':hasDecision \rightarrow$ “decision”, $o':decision1 \rightarrow$ “decision for paper1”, $o':Decision \rightarrow$ “decision”).
- ⑦ Compare L_s and L_t (Section 4.4).
- ⑧ Select the subgraphs parts with the best similarity score, transform them into DL formulae (Section 4.3), and aggregate them (Section 4.5). In this example, the part of the subgraph that is the most similar to the CQA (in terms of label similarity) is $o':Acceptance$. The DL formula is therefore $\exists o':hasDecision.o':Acceptance$.
- ⑨ Reassess the similarity of each DL formula based on their counter-examples (Section 4.6 and Section 4.7). The counter-examples are common instances of the two knowledge bases which are described by the target DL formula but not by the original CQA.
- ⑩ Filter the DL formulae based on their similarity score (if their similarity score is higher than a threshold) (Section 4.8).
- ⑪ Put the DL formulae e_s and e_t together to form a correspondence (e.g., $\langle o:AcceptedPaper, \exists o':hasDecision.o':Acceptance, \equiv \rangle$) and express this correspondence in a reusable format (e.g., EDOAL). The confidence assigned to a correspondence is the similarity score of the DL formula computed.

3.2. Approach over a binary CQA

- ① Extract source DL formula e_s (e.g., $o:paperWrittenBy$) from SPARQL CQA (Section 4.1): `SELECT ?x ?y WHERE ?x ol:paperWrittenBy ?y.`
- ② Extract lexical information from the CQA, L_s set labels from the DL formula (e.g., “paper written by”).
- ③ Extract source answers ans_s of the CQA (e.g., a pair of instances ($o:paper1, o:person1$)).
- ④ Find equivalent or similar target answers ans_t to the source instances ans_s (e.g. $o:paper1 \sim o':paper1$ and $o:person1 \sim o':person1$) (Section 4.2).
- ⑤ Retrieve the subgraphs of target answers (Section 4.3): for a binary query, it is the set of paths between two answer instances as well as the types of the instances appearing in the path (e.g., a path of length 1 is found between $o':paper1$ and $o':person1$). The path is composed of only one property and there are no other instances than $o':paper1$ and $o':person1$ in this path. Their respective types are retrieved: ($o':Paper, o':Document$) for $o':paper1$ and ($o':Person$) for $o':person1$.
- ⑥ For each subgraph, retrieve L_t the labels of its entities (e.g., $o':writes \rightarrow$ “writes”, $o':Person \rightarrow$ “person”, $o':Paper \rightarrow$ “paper”, etc.).
- ⑦ Compare L_s and L_t (Section 4.4).

- ⑧ Select the subgraph parts with the best score and transform them into DL formulae (Section 4.3). Keep the best path variable types if their similarity is higher than a threshold. (e.g., the best type for the instance $o':paper1$ is $o':Paper$ because its similarity with the CQA labels is higher than the similarity of $o':Document$).
- ⑨ Reassess the similarity of each DL formula based on their counter-examples (Section 4.6 and Section 4.7).
- ⑩ Filter the DL formulae based on their similarity score (Section 4.8).
- ⑪ Put the DL formulae e_s and e_t together to form a correspondence (e.g., $\langle o:paperWrittenBy, dom(o':Paper) \sqcap o':writes^-, \equiv \rangle$ and express this correspondence in a reusable format (e.g., EDOAL). The confidence assigned to a correspondence is the similarity score of the DL formula computed.

The main difference with the case of unary CQAs is in Step ④ because the two instances of the pair answer are matched instead of one, Step ⑤ and Step ⑧ which deal with the subgraph extraction and pruning.

4. Main steps of the approach

This section details the steps ①, ④, ⑤, ⑦, ⑧, ⑨ and ⑩ of Figure 1 and illustrate them with examples.

4.1. Translating SPARQL CQAs into DL formulae

In Step ①, to translate a SPARQL query into a DL formula, the first step is to translate it into a FOL formula and then transform it into a DL formula. A SPARQL SELECT query (in the scope of our approach) is composed of a SELECT clause containing variable names and a basic graph pattern, i.e., a set of triples with variables sometimes with constructors (such as UNION or MINUS). First, the variables in the SELECT clause become the quantified variables of the formula. In unary CQA, the SELECT clause contains one variable. In binary CQA, the SELECT clause contains two variables. The SPARQL query of Figure 4, $?x$ becomes the quantified variable of our formula: $\forall x$. Then, the basic graph pattern is parsed to find what predicates apply to the quantified variables and add them to the formula. Each triple of the basic graph pattern is either a unary or a binary predicate. If new variables are added, an existential quantifier is used for them. In the example, we find the triple $\langle ?x, o':hasDecision, ?y \rangle$. The FOL formula becomes $\forall x, \exists y, o':hasDecision(x, y)$. We then recursively keep on exploring the basic graph pattern for each new variable introduced. After exploring the basic graph pattern for the variable $?y$, the FOL formula becomes $\forall x, \exists y, o':hasDecision(x, y) \wedge o':Acceptance(y)$. At the end of the process, we transform the basic graph pattern into a DL formula, which can also be translated into an EDOAL formula as shown below. $\forall x, \exists y, o':hasDecision(x, y) \wedge o':Acceptance(y)$ becomes in DL: $\exists o':hasDecision.o':Acceptance$. The FOL to DL equivalence is done as in [4].

4.2. Instance matching

In Step ④, the answers of the CQA over the source knowledge base that have been retrieved are matched with the instances of the target knowledge base. This instance matching phase relies on existing links (*owl:sameAs*, *skos:exactMatch*, *skos:closeMatch*, etc.) if they exist. If no such link exists, an exact label match is performed. When dealing with binary CQA whose results are an instance-literal value pair, the instance is matched as before (existing links or exact labels), and the literal value will be matched with an exactly identical value (the datatype is not considered) in the pathfinding step, detailed in Section 4.3.

4.3. Retrieving and pruning subgraphs

The approach relies on subgraphs, which are sets of triples from a knowledge base. These subgraphs are found (Step ⑤), pruned, and transformed into DL formulae (Step ⑧). The type of subgraphs for unary or binary CQAs is inspired by [39], which proposes an approach to find equivalent subgraphs within the same knowledge base.

```
SELECT ?x WHERE {
  ?x o2:hasDecision ?y.
  ?y a o2:Acceptance.
}
```

Fig. 4. SPARQL SELECT query with one variable in SELECT clause.

A **unary CQA** expects a set of single instances as an answer. The subgraph of a single instance is composed of a triple in which the instance is either the subject or the object, and the types (classes) of the object or subject of this triple. For example, $o':paper1$ is the subject of the triple $o':paper1 \ o':hasDecision \ o':decision1$ and $o':decision1$ has types (classes) $o':Acceptance$ and $o':Decision$. A subgraph of $o':paper1$ is therefore composed of the three following triples: (i) $\langle o':paper1, o':hasDecision, o':decision1 \rangle$, (ii) $\langle o':decision1, rdf:type, o':Acceptance \rangle$, (iii) $\langle o':decision1, rdf:type, o':Decision \rangle$.

When comparing the subgraph with the CQA labels, if the most similar object (resp. subject) type is more similar than the object (resp. subject) itself, the type is kept. Let us consider the *accepted paper* CQA. The most similar type of the triple of the object is $o':Acceptance$. Therefore, triple 3 is pruned. The object of triple 1 is $o':decision1$ and the most similar object type to the CQA is $o':Acceptance$. $o':Acceptance$ is more similar to the CQA than $o':decision1$. Therefore, $o':decision1$ becomes a variable and triple 2 stays in the subgraph. In order to translate a subgraph into a DL formula, we first translate this subgraph into a SPARQL query:

- The answer is transformed into a variable and put in the SELECT clause. In this example, $o':paper1$ becomes a variable $?x$ in the SELECT clause: `SELECT ?x WHERE`.
- The instances of the subgraphs that are not kept are transformed into variables. In this example, $o':decision1$ becomes a variable $?y$.
- These transformations are applied to the selected triples of the subgraph which become the basic graph pattern of the SPARQL query. In this example, the SPARQL query is the one in Figure 4.

Finally, the SPARQL query is transformed into a DL formula by using the same process as that described in Section 4.1: $\exists o':hasDecision.o':Acceptance$.

A **binary CQA** expects a set of pairs of instances (or pairs of instance-literal value) as an answer. Finding a subgraph for a pair of instances consists in finding a path between the two instances. The shortest paths are considered more accurate. Because finding the shortest path between two entities is a complex problem, paths of length below a threshold are sought. First, paths of length 1 are sought, then if no path of length 1 is found, paths of length 2 are sought, etc. If more than one path of the same length is found, all of them go through the following process. When a path is found, the types of instances forming the path are retrieved. If the similarity of the most similar type to the CQA is above a threshold, this type is kept in the final subgraph. For example, for a “*paper written by*” CQA with the answer $(o':paper1, o':person1)$ in the target knowledge, a subgraph containing the following triples is found: $\langle o':person1, o':writes, o':paper1 \rangle$, $\langle o':paper1, rdf:type, o':Paper \rangle$, $\langle o':paper1, rdf:type, o':Document \rangle$, $\langle o':person1, rdf:type, o':Person \rangle$. The most similar type of $o':person1$ is $o':Person$, which is below the similarity threshold. Triple 4 is then removed from the subgraph. The most similar type of $o':paper1$ is $o':Paper$. Triple 3 is therefore removed from the subgraph. $o':Paper$'s similarity is above the similarity threshold: triple 2 stays in the subgraph. The translation of a subgraph into a SPARQL query is the same for binary and unary CQAs. Therefore, the subgraph will be transformed into a SPARQL query and saved as a DL formula: $dom(o':Paper) \sqcap o':writes^-$.

4.4. Label similarity

In step ⑦, a label similarity metric is needed to compare two sets of labels L_s and L_t . A Levenshtein [14] distance-based similarity metric was chosen. It measures the minimum number of single-character edits (insertions, deletions, or substitutions) between two strings. The similarity between two sets of labels is the Cartesian product of the string similarities between the labels of L_s and L_t (1). *strSim* is the string similarity of two labels l_s and l_t (2).

$$sim(L_s, L_t) = \sum_{l_s \in L_s} \sum_{l_t \in L_t} strSim(l_s, l_t) \quad (1)$$

$$strSim(l_s, l_t) = \begin{cases} \sigma & \text{if } \sigma > \tau, \\ 0 & \text{otherwise} \end{cases} \quad \text{where } \sigma = 1 - \frac{levenshteinDist(l_s, l_t)}{\max(|l_s|, |l_t|)} \quad (2)$$

Table 1

Initial, extension, intension and final (in bold) formulae. The CQA considered is “accepted papers”

Initial formulae	Extension	Intension
$\exists o':hasDecision.\{o':accept\}$	$\exists o':hasDecision.\{o':accept,$	$\exists o':hasDecision.\top$
$\exists o':hasDecision.\{o':strongAccept\}$	$o':strongAccept, o':weakAccept\}$	
$\exists o':hasDecision.\{o':weakAccept\}$		
$\exists o':acceptedBy.\{o':person1\}$	$\exists o':acceptedBy.\{o':person1\}$	$\exists o':acceptedBy.\top$

4.5. DL formula aggregation

In Step ⑧ of the approach, when dealing with unary CQA, the DL formulae can be aggregated. It consists in transforming one or more formulae with a common predicate into a more generic formula. This aggregation only applies to formulae that contain an instance or a literal value and which were kept in the subgraph selection step. For example, this step would apply to a formula such as $\exists o':hasDecision.\{o':accept\}$. There are three steps to the aggregation. First, we create a first aggregated formula which we call the **extension** formula. It consists in merging the instances or literal values of the formulae with the same predicate into one set of values. Let us consider that through various answers to a CQA (e.g., $o':paper1$, $o':paper2$, etc.), we encountered the following formulae: $\exists o':hasDecision.\{o':accept\}$, $\exists o':hasDecision.\{o':strongAccept\}$, $\exists o':hasDecision.\{o':weakAccept\}$. The extension formula of these formulae is: $\exists o':hasDecision.\{o':accept, o':strongAccept, o':weakAccept\}$. The extension formula of a formula that does not share its predicate with any other is the formula itself. Then, an **intension** formula can be computed by replacing the set of values by the top class \top . The intension formula of the example formulae is: $\exists o':hasDecision.\top$. Finally, a choice is made between the extension or intension formulae based on the predicate similarity to the CQA. If the predicate is more similar than the values, the intension formula is kept. Otherwise, the extension formula is kept. In our example, the extension formula $\exists o':hasDecision.\{o':accept, o':strongAccept, o':weakAccept\}$ is kept.

We present two examples of initial formulae, with their respective intension and extension formulae in Table 1. These were obtained with the competency question “accepted paper”. In Table 1, the final formulae are in bold. Applied to the examples of Table 1:

- $o':accept$, $o':strongAccept$ and $o':weakAccept$ are more similar to the CQA than $o':hasDecision$. The extension form is chosen.
- $o':acceptedBy$ is more similar (based on labels) to the CQA than $o':person1$. The intension form is chosen.

4.6. Calculating the percentage of counter-examples

In Step ⑨, the approach refines the DL formula similarity score by looking for counter-examples (details about the similarity score are given in Section 4.7). A **counter-example** is a common instance of the source and target ontologies which is described by the DL formula found by the approach in the target ontology but which is not described by the CQA in the source ontology. For example, let us assume that the target formula e_t is $o':Paper$ for the “accepted paper” CQA. From the target ontology, the answers $o':paper1$, $o':paper2$, $o':paper3$ and $o':paper4$ are retrieved from e_t and matched to the source instances respectively $o:paper1$, $o:paper2$, $o:paper3$ and $o:paper4$. However, only $o:paper1$ and $o:paper2$ are accepted papers (and are described by the CQA) in the source ontology. Therefore $o:paper3$ and $o:paper4$ are counter-examples. The percentage of counter-examples is computed as follows. The answers $ans_t^{e_t}$ described by the target subgraph (e_t) are retrieved from the target knowledge. These answers are matched to source instances: $ans_s^{e_t}$. The percentage of counter-examples is the proportion of common instances $ans_s^{e_t}$ which are not answers to the CQA ($\neg(ans_s^{cqa})$). The equation for the percentage of counter-examples ($percCounterExamples$) is therefore:

$$percCounterExamples = \frac{|ans_s^{e_t} \sqcap \neg(ans_s^{cqa})|}{|ans_s^{e_t}|} \quad (3)$$

In the example, the percentage of counter-example is $\frac{2}{4} = 50\%$.

4.7. DL formula similarity

In Step ⑩, the formulae are filtered based on their similarity score with the CQA. The similarity score is a combination of:

Label similarity *labelSim* is the sum of the label similarity of each entity of the formula with the CQA.

Structural similarity *structSim*. This similarity was introduced to enhance some structural aspects of a formula. In the implementation of the approach, this value is set to 0.5 for a path between the two instances of the answer, and 0 for a unary CQA subgraph. Indeed, if the label similarity of the path is 0, the structural similarity hints that the fact that a path was found is a clue in favor of the resulting DL formula.

Percentage of counter examples *percCounterExamples* which is computed in Step ⑨ and detailed Section 4.6.

The similarity score is calculated with the following equation:

$$\text{similarity} = (\text{labelSim} + \text{structuralSim}) \times (1 - \text{percCounterExamples}) \quad (4)$$

For instance, consider the similarity of $\exists o': \text{hasDecision}.o': \text{Acceptance}$ with the unary CQA “accepted paper”.

- *labelSim* = 0.8 + 0.0 as $\text{sim}(\text{labels}(\text{CQA}), \text{labels}(o': \text{hasDecision})) = 0.0$ and $\text{sim}(\text{labels}(\text{CQA}), \text{labels}(o': \text{Acceptance})) = 0.8$
- *structuralSim* = 0.0 because it is a unary CQA
- *percCounterExamples* = 0.0

The similarity of this DL formula is $\text{similarity} = (0.8 + 0.0) \times (1 - 0) = 0.8$

4.8. DL formula filtering

In Step ⑩, the formulae are filtered. Only the DL formulae with a similarity higher than a threshold are put in correspondence with the CQA DL formula. If for a given CQA, there is no DL formula with a similarity higher than the threshold, only the best DL formulae with a non-zero similarity are put in the correspondence. The best DL formulae are the formulae with the highest similarity score. When putting the DL formula in a correspondence, if its similarity score is greater than 1, the correspondence confidence value is set to 1.

The definition of similarity can be seen as unusual, as it ranges from 0 to 1.5, while values of similarity are usually in [0,1]. We chose to have a structural similarity strong enough to exceed the sorting threshold for properties. In order to consider a non-structural similarity, a lexical measure has been considered. The structural similarity is only applied in the case of properties. In all cases, similarity/confidence is only used to filter out correspondences.

5. Evaluation

The approach has been automatically evaluated on a synthetic dataset (Populated Conference dataset), to measure the impact of various parameters on the approach. It was also evaluated on LOD repositories (Taxon dataset) to study how the approach performs when faced with Linked Open Data challenges such as large ontologies and millions of triples. Some of the knowledge bases chosen for this experiment are irregularly populated. This means that the same piece of knowledge can be represented in various ways in the same ontology and that all instances are not described identically. After detailing the evaluation parameters in Section 5.1 and the evaluation settings in Section 5.2, the results over the two datasets are presented (Sections 5.3 and 5.5, respectively). The discussion is then presented in Section 6.

5.1. Matching approach set-up

Label similarity A threshold is applied to the similarity measure obtained: if the similarity between two labels is below a threshold τ , this similarity is considered noisy and is set to zero.

Table 2

Parameters of the evaluated variants of the approach: number of support answers (Nb. ans.), Levenshtein threshold in the similarity metric (Lev. thr.), type of instance matching strategy (Inst. match), computation of counter-examples (co.-ex.), CQAs input or query. In bold the parameter that has been changed with respect to the variant

Evaluated variant	Nb ans.	Lev. thr.	Inst. match	Co.-ex.	CQAs	query
baseline	10	0.4	links		✓	
Levenshtein	10	0.0–1.0	links		✓	
Support answers	1–100	0.4	links		✓	
exact label match	10	0.4	labels		✓	
query	10	0.4	links			✓
query+reassess	10	0.4	links	✓		✓
cqa+reassess	10	0.4	links	✓	✓	

Path length threshold The maximum path length sought is 3. Paths longer than that may bring noise in the correspondences, as the path-finding algorithm searches for all combinations of properties.

Structural similarity The structural similarity is 0 for a triple (in the case of a unary CQA) and 0.5 for a path found between two matched entities (in the case of a binary CQA). Finding a path between two instances (the matched answers of a binary CQA) is a hint that this subgraph can be correct. In opposition, the structure subgraphs for unary CQA are not that informative.

DL formula threshold The DL formulae with a similarity higher than 0.6 are kept. If a CQA has no DL formula with a similarity higher than 0.6, the best formulae are put in correspondence (the formulae having the best similarity, if this similarity is above 0.01). This threshold was chosen to be above the structural similarity threshold (0.5) for a path subgraph. Indeed, if two paths are found but only one has a label similarity above 0, then its associated DL formula will be the only one output. These thresholds were empirically chosen.

Approach variants The other parameters have been varied to create a set of *variants*. These variants are listed in Table 2. For each variant (lines in the table), the different parameters (number of support answers, Levenshtein threshold, type of instance matching strategy, and computation of counter-examples) have been varied. The values of the baseline approach were empirically chosen: a Levenshtein distance threshold of 0.4, 10 support answers, and no similarity value reassessment based on counter-examples. Note that the support answers correspond to the CQA answers with a match in the target knowledge base which are used to find subgraphs.

5.2. Evaluation settings

Evaluation datasets An automatic evaluation was performed on the populated version of the OAEI Conference benchmark [32]. This dataset is composed of 5 ontologies, with 100 manually generated CQAs. This evaluation measured the impact of various parameters on the approach. Second, a manual evaluation was carried out on the Taxon dataset about plant taxonomy, composed of 4 large populated ontologies: AgronomicTaxon [25], AgroVoc [5], DBpedia [3] and TaxRef-LD [16]. 6 CQAs from AgronomicTaxon have been manually generated. The CQA used in this evaluation are the one presented in [31] which were manually written from AgronomicTaxon CQs [25].

Evaluation metrics The evaluation metrics are based on the comparison of instance sets, as described in [30]. The generated alignment is used to rewrite a set of reference source CQAs whose results (set of instances) are compared to the ones returned by the corresponding target reference CQA. This metric shows the overall *coverage* of the alignment with respect to the knowledge needs and the best-rewritten query.³ A balancing strategy consists of calculating the intrinsic alignment *precision* based on common instances.

Given an alignment A_{eval} to be evaluated, a set of CQA reference pairs cqa_{pairs} (composed of source cqa_s and target cqa_t), kb_s the source knowledge base, kb_t a target knowledge base, and f an instance set (I) comparison

³The description of rewriting systems is out of the scope of this paper.

function:

$$\text{coverage}(A_{eval}, cqa_{pairs}, kb_s, kb_t, f) = \text{average}_{(cqa_s, cqa_t) \in cqa_{pairs}} f(I_{cqa_t}^{kb_s}, I_{bestq_t}^{kb_t}) \quad (5)$$

coverage is based on the *queryFmeasure* (also used for selecting the best-rewritten query). This is motivated by the fact that it better balances precision and recall. Given a reference instance set I_{ref} and an evaluated instance set I_{eval} :

$$QP = \frac{|I_{eval} \cap I_{ref}|}{|I_{eval}|} \quad QR = \frac{|I_{eval} \cap I_{ref}|}{|I_{ref}|} \quad (6)$$

$$\text{queryFmeasure}(I_{ref}, I_{eval}) = 2 \times \frac{QR \times QP}{QR + QP} \quad (7)$$

$$bestq_t = \underset{q_t \in \text{rewrite}(cqa_s, A_{eval}, kb_s)}{\text{argmax}} \quad \text{queryFmeasure}(I_{cqa_t}^{kb_s}, I_{q_t}^{kb_t}) \quad (8)$$

A best-match (query f-measure) aggregation over the reference CQA is performed. An average of the best-match scores gives the **CQA Coverage**.

Balancing *coverage*, *precision* is based on classical (i.e., scoring 1 for same instance sets or 0 otherwise) or non-disjoint functions f (as in the following):

$$\text{precision}(A_{eval}, kb_s, kb_t, f) = \text{average}_{(e_s, e_t) \in A_{eval}} f(I_{e_1}^{kb_s}, I_{e_2}^{kb_t}) \quad (9)$$

The CQA Coverage and Precision with the same scoring metric are finally aggregated in a **Harmonic Mean**.

For both *coverage* and *precision*, different functions f can be used for comparing instance sets (overlap, precision-oriented, recall-oriented etc.). These different functions are complementary. The **classical** (Equation (10)), **recall-oriented** (Equation (11)) and **precision-oriented** (Equation (12)) scoring functions are used in state-of-the-art works to emphasise whether the alignment favours precision or recall [6]. We have introduced the **overlap** metric to represent whether two queries have at least one common answer (Equation (13)). The **not disjoint** metric gives a 1 score to all the overlapping queries and the queries where I_{ev} and I_{ref} are empty sets.

$$\text{classical}(I_{ref}, I_{ev}) = \begin{cases} 1 & \text{if } I_{ev} \equiv I_{ref} \\ 0 & \text{otherwise} \end{cases} \quad (10)$$

$$\text{recall oriented}(I_{ref}, I_{ev}) = \begin{cases} 1 & \text{if } I_{ev} \supseteq I_{ref} \\ 0.5 & \text{if } I_{ev} \subseteq I_{ref} \\ 0 & \text{otherwise} \end{cases} \quad (11)$$

$$\text{precision oriented}(I_{ref}, I_{ev}) = \begin{cases} 1 & \text{if } I_{ev} \subseteq I_{ref} \\ 0.5 & \text{if } I_{ev} \supseteq I_{ref} \\ 0 & \text{otherwise} \end{cases} \quad (12)$$

$$\text{overlap}(I_{ref}, I_{ev}) = \begin{cases} 1 & \text{if } I_{ev} \not\sqcap I_{ref} \\ 0 & \text{otherwise} \end{cases} \quad (13)$$

$$\text{not disjoint}(I_{ref}, I_{ev}) = \begin{cases} 1 & \text{if } I_{ev} \not\sqcap I_{ref} \text{ or } I_{ev} \stackrel{\emptyset}{=} I_{ref} \\ 0 & \text{otherwise} \end{cases} \quad (14)$$

Such metrics have been used in the automatic evaluation of the controlled populated version of the Conference dataset. Given the uneven population of Taxon (i.e., a same piece of knowledge can be represented in various ways

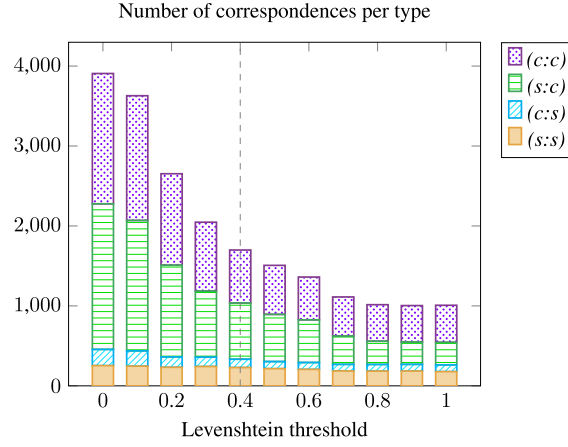


Fig. 5. Number of correspondences per type for each variant with a different Levenshtein threshold.

within the same ontology and that all instances are not described identically), a manual evaluation has been carried out instead to avoid entailing noise in the instance-based comparison.

Environment The approach and evaluation system has been executed on an Ubuntu 16.04 machine configured with 16 GB of RAM running under an i7-4790 K CPU 4.00 GHz \times 8 processors. The runtimes are given for a single run. The local SPARQL endpoints were run on the same machine with Fuseki 2.⁴

5.3. Results on populated conference

The approach has been run and evaluated on the populated conference 100% dataset on a local Fuseki 2 server. This choice is motivated by the fact that the search for a common instance is faster when the proportion of common instances in the source answers is higher. The implementation in Java of the evaluation system, as well as the Populated Conference dataset, is available.⁵

The *variants* of the approach have been compared to its baseline (Table 2). The parameters that are not described in this table such as path length threshold (3), DL formula filtering threshold (0.6), and structural similarity constants (0.5 for a path, 0 for a class expression) as presented in Section 5.1. This evaluation strategy allows for isolating the parameters and measuring their impact, as discussed in the following.

5.3.1. Impact of the threshold in the string similarity metric

An evaluation was performed with a Levenshtein threshold set between 0.0 and 1.0. Figure 5 shows the number of found correspondence per type, with the detailed results in Figure 6. The number of correspondences decreases when the Levenshtein threshold increases. Numerous correspondences obtained with a low Levenshtein threshold cover a lot of CQAs (high CQA Coverage) but contain a lot of errors (low Precision). The lower the threshold, the better the CQA Coverage and the lowest the Precision. The Harmonic Mean is the highest for a threshold of 0.4 in the similarity metric. The baseline approach Levenshtein threshold (0.4) was chosen based on this experiment.

5.3.2. Impact of the number of support answers

The approach has been evaluated with a number of support answers between 1 and 100. The runtime of the approach over the 20 oriented pairs of ontologies is displayed in Figure 7 and Figure 8 shows the number of correspondences per type. The evaluation results are shown in Figure 9. It could be observed that even with 1 answer as support, the CQA Coverage and Precision scores are high, which shows that the approach can make a generalization from a few examples. As expected, the bigger the number of support answers, the longer the process is to run. Some CQA has only 5 answers (only 5 conference instances in the population of the ontologies), which explains why the

⁴<https://jena.apache.org/documentation/fuseki2/>

⁵https://framagit.org/IRIT_UT2J/conference-dataset-population

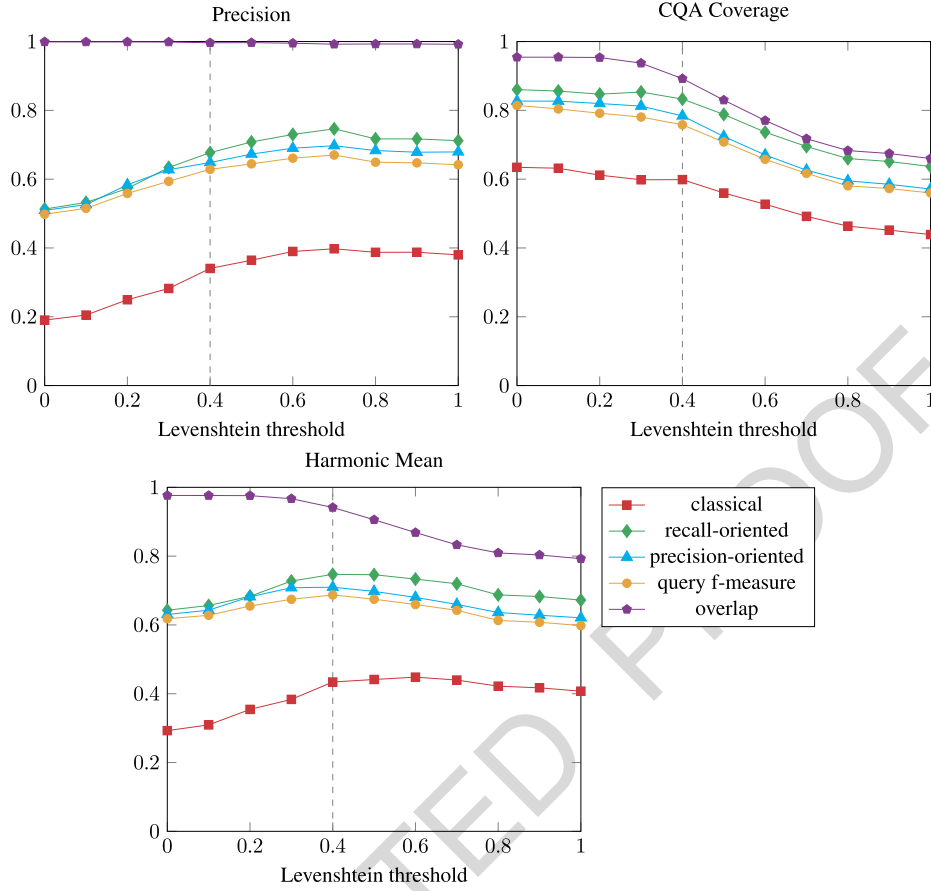


Fig. 6. Results of the evaluation with 10 support answers and variable Levenshtein threshold in the string similarity measure. The baseline results are highlighted by a vertical dashed line.

time rises linearly between 1 support answer and 5 support answers and has a lower linear coefficient for support instances over 5. The Precision scores get lower with more support answers. The main reason is that particular answer cases that are lexically similar to the CQA labels can be discovered when a lot of instances are considered. For example, the correspondence $\langle \text{cmt:Person}, \exists \text{ conference:has_the_last_name.}\{\text{"Benson"}\}, \equiv \rangle$ was discovered by the variant with 100 support answers. Indeed, "Benson" is lexically similar to "Person". The increase in the number of correspondences with the number of support answers shows that incorrect correspondences have been introduced.

The same problem occurs in the CQA Coverage: with 100 support answers, special cases having a higher similarity to the CQA than the expected formula can be found. As in the approach, the formulae are filtered, and when the similarity of the best formula is below a threshold (0.6), only the best one is kept. For example, with 10 support answers, the correspondence $\langle \text{conference:Rejected_contribution}, \exists \text{cmt:hasDecision.cmt:Rejection}, \equiv \rangle$ was found for the "rejected paper" CQA, the similarity of the target DL formula (0.43) was below the threshold (0.6) but it was the best formula so it was kept. For the 100 support answers, the correspondence $\langle \text{conference:Rejected_contribution}, \exists \text{cmt:hasSubjectArea.}\{\text{"entity consolidation"}, \text{"distribution"}, \text{"categorization"}\}, \equiv \rangle$ and had a DL formula similarity (0.44) higher than the expected formula, so only this correspondence was output. Therefore the *conference:Rejected_contribution* CQA could not be covered with this alignment. However, the overlap CQA Coverage gets slightly higher for a high number of support answers because accidental correspondences have been introduced. For example, the correspondence $\langle \text{conference:Topic}, \exists \text{rdfs:label.}\{\text{"compliance"}\}, \equiv \rangle$ was found with 100 support answers because "topic" and "compliance" have a 0.4 label similarity score. The *Topic* CQA over the *conference-*

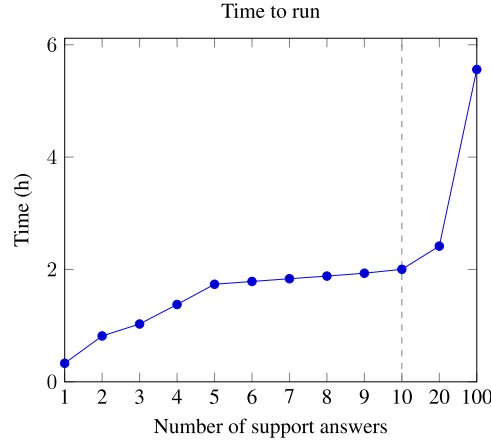


Fig. 7. Time taken by the approach to run for the 20 oriented pairs of ontologies with a different number of support answers.

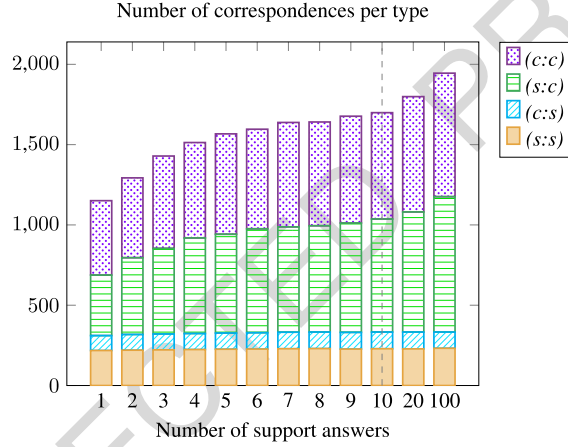


Fig. 8. Number of correspondence per type for each variant with a different number of support answers.

cmt pair was not covered by the variants of the approach with less than 100 support answers because no DL formula with a similarity above 0 was found.

5.3.3. Similar instances based on exact label match or existing links

A variant of the approach does not use existing links between instances, instead, it performs an exact label match between instances. Figure 10 shows the number of correspondences per type of the baseline and its variant. Figure 11 shows the results of the baseline and its *exact label match* variant. The use of an exact label match for the instance matching phase brings noise to the correspondences and lowers the Precision. The overlap Precision also decreases because the correspondences are not ensured to share a common instance. In the baseline approach, which uses *owl:sameAs* links, the support answers were by definition common instances, and outputting correspondences with no overlap was not possible (except when dealing CQA with literal values). For example, the paper submissions and their abstracts share the same title. Therefore, a rejected paper instance can be matched with its abstract in the target knowledge base. The following correspondence results from this wrong instance match: $\langle ekaw:RejectedPaper, \exists conference:is_the_first_part_of.conference:Rejected_contribution, \equiv \rangle$. This impacts the number of *(c:c)* correspondences which increases significantly when using the *exact label match*. Some ontologies use two data properties to link a person to their first and last name. The first and last names are then considered independent labels of the person instance. This induces confusion between two people sharing a first or a last name.

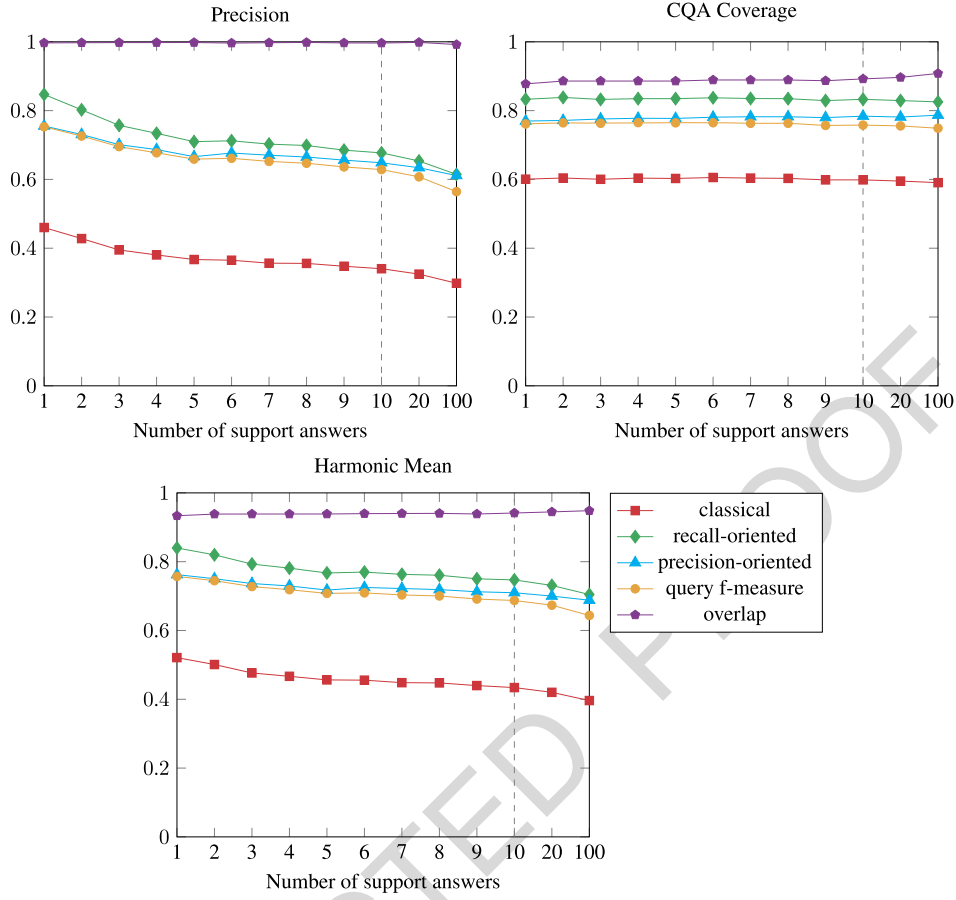


Fig. 9. Results of the evaluation with a 0.4 Levenshtein similarity threshold and a variable number of support answers. The baseline results are highlighted by a vertical dashed line.

The following correspondence was obtained by matching a person to another sharing the same first name: $\langle conference:has_the_first_name, edas:isReviewedBy^- \circ edas:isWrittenBy \circ edas:hasFirstName, \equiv \rangle$.

The baseline approach (existing *owl:sameAs* links) takes **2.0 hours** to run over the 20 pairs of ontologies whereas the exact label match approach takes **59.2 hours**. The long runtime for the exact label match approach can be explained by the necessary steps to find the exact label match answers. First, the labels of each source answer to the CQA must be retrieved. This query takes about 64 ms. Then, for each label of the source answer, a match is sought. The runtime of the query to retrieve all instances annotated by a given label is about 2 s. The reason is that this query contains a tautology. The choice of this query was made because some ontologies define their labeling properties instead of using *rdfs:label* or other widely used properties.

When using direct links, these steps are replaced by directly retrieving *owl:sameAs* links, which takes about 20 ms per source instance. If the number of common support answers between the source and target ontology is reached (in the baseline, when 10 support answers are found), the approach stops looking for new matches. However, when no common instance can be found, the approach looks for a match for every answer of the CQA. This fact coupled with the slow label queries results in a long time. When common instances exist but do not share the same exact labels, the approach also looks for matches for every source answer, without success. For example, *cmt* represents the full name of a person, and *conference* represents its first name and its last name in two different labels. For the CQA retrieving all the Person instances, the approach goes through the 4351 instances without finding any match.

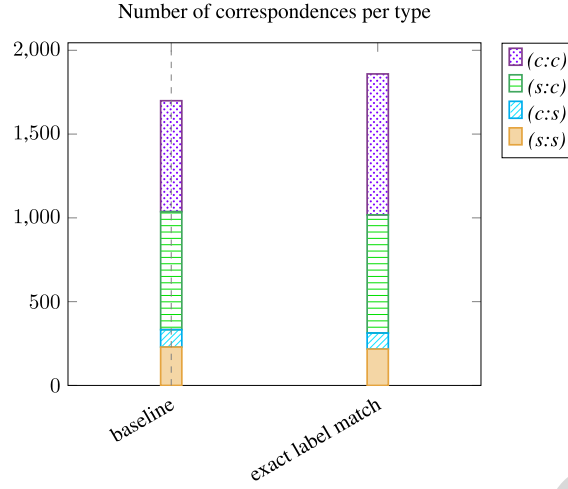


Fig. 10. Number of correspondence per type for the baseline and the variant based on exact label match.

5.3.4. CQAs or generated queries

In order to measure how the CQA impacts the results of the approach, the baseline approach is compared to a variant that does not rely on input CQA but automatically generates queries. Three types of SPARQL queries are generated for a given source ontology: *Classes*, *Properties*, and *Property-Value pairs*.

Classes For each *owl:Class* populated with at least one instance, a SPARQL query is created to retrieve all the instances of this class.

Properties For each *owl:ObjectProperty* or *owl:DatatypeProperty* with at least one instantiation, a SPARQL query is created to retrieve all the pairs of instances of this class.

Property-value pairs Inspired by the approaches of [17,18,35], SPARQL queries of the following form are created:

- SELECT DISTINCT ?x WHERE { ?x o1:property1 o1:Entity1. }
- SELECT DISTINCT ?x WHERE { o1:Entity1 o1:property1 ?x. }
- SELECT DISTINCT ?x WHERE { ?x o1:property1 "Value". }

Table 3 shows the number of generated queries per source ontology of the evaluation set.

The approach based on generated queries will not output a correspondence for each CQA in the evaluation. Therefore, the rewriting systems in the evaluation process will bring noise. The CQA Coverage scores are comparable as only the best result is kept. The Precision of the alignment output is computed by comparing the instances of the source and target members in their respective ontologies. These Precision scores give an indicator of the actual precision of these approaches.

The results of the evaluation of the baseline (based on CQAs) and the query variant are presented in Figure 12. Figure 13 shows the number of correspondences per type. The CQA Coverage scores when the approach is based on generated queries are between 10% and 20% lower than those obtained with CQA. Indeed, the (c:c) correspondences it retrieves are limited to the Class-by-Attribute-Value pattern on their source member. The Precision scores are not comparable because the ontologies were populated based on CQA and not on entities: a *Document* class may be populated with more or fewer instances given its subclasses. As the approach relies on common instances, the overlap Precision (percentage of correspondences whose member's instances overlap) is around 1.0. The classical Precision (percentage of correspondences whose members are strictly equivalent) is, however, rather low overall.

The baseline and the *query* variant both take **2.0 hours** to run on the 20 pairs of ontologies. Even if there are more queries to cover than CQA, the runtime of the *query* variant is compensated by the “difficulty” of the CQA: some CQAs contain unions or property paths and therefore take more time to be answered by the Fuseki server than the generated queries.

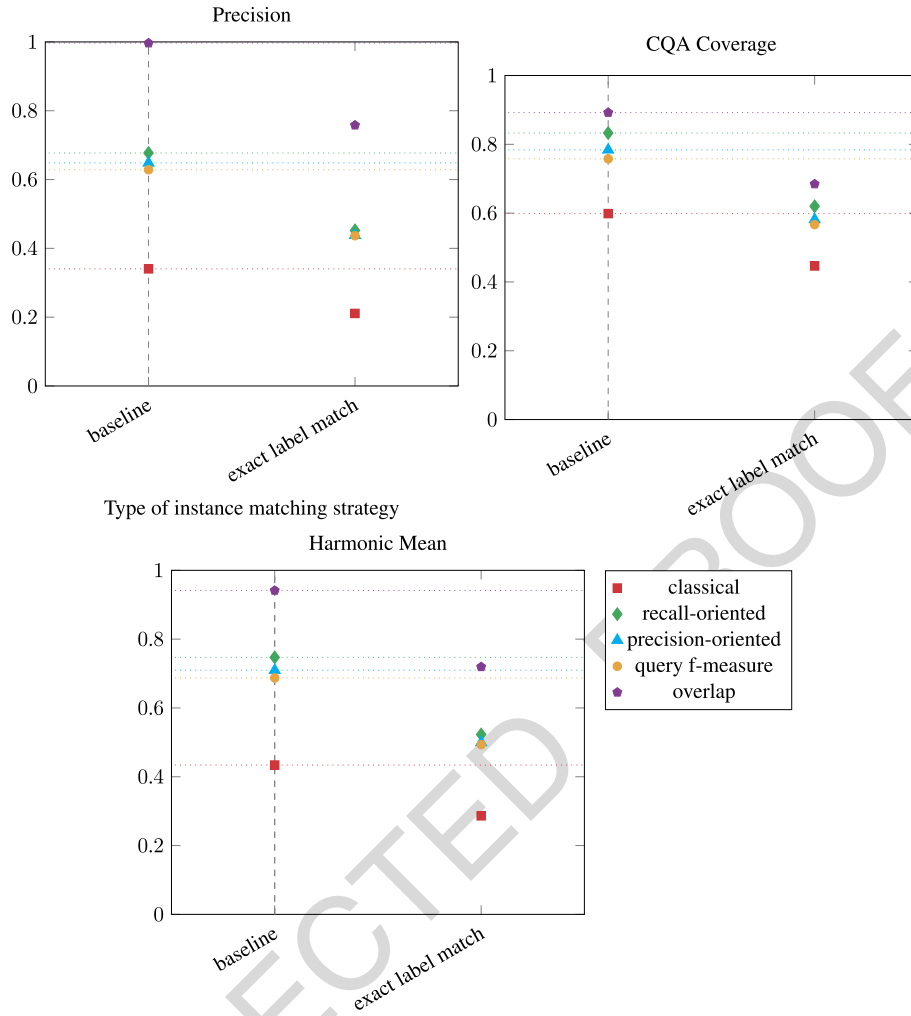


Fig. 11. Comparison of the approach results when relying on existing *owl:sameAs* links or on an exact label-based instance matching. The baseline results are highlighted by a vertical dashed line.

Table 3
Number of generated queries and CQAs per source ontology

Nb of queries	cmt	Conference	confOf	edas	ekaw
classes	26	51	29	43	57
properties	50	50	20	28	26
properties-value	30	20	0	5	15
TOTAL	106	121	49	76	98
CQAs	34	73	54	52	65

The number of (*s:s*) and (*s:c*) correspondences is much higher for the *query* variant. This approach generates 380 queries that express simple expressions (lines classes and properties of Table 3) and therefore, will give (*s:s*) or (*s:c*) correspondences if a match is found. In comparison, the baseline approach relies on 133 SPARQL CQAs which represent a simple expression, and 145 which represent a complex expression.

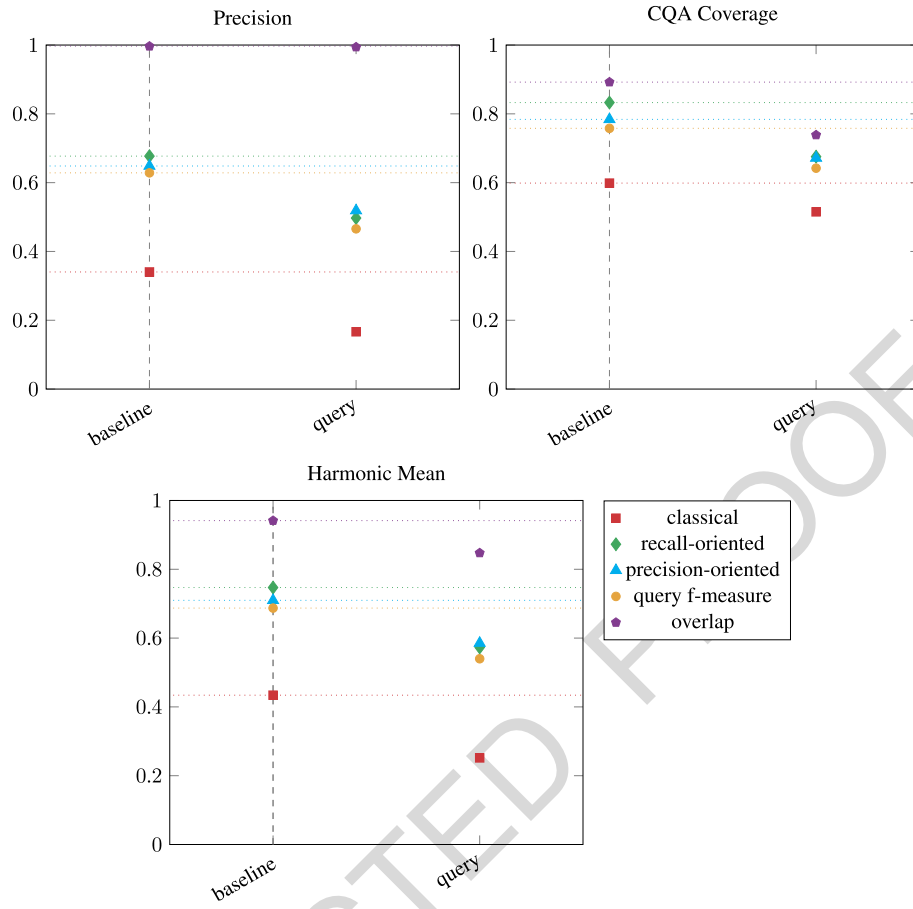


Fig. 12. Results for the baseline and the variant which generates queries (query).

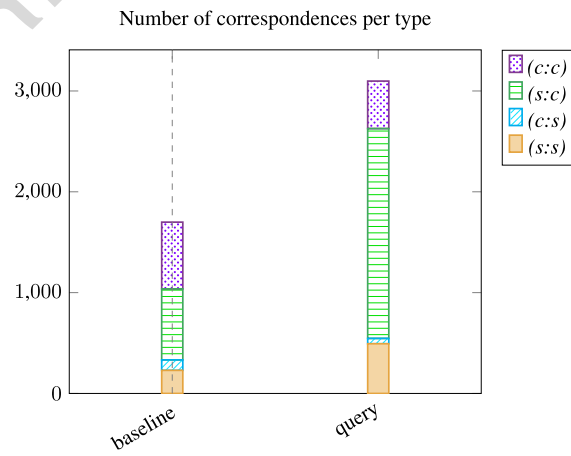


Fig. 13. Number of correspondence per type for the baseline and the variant which generates queries.

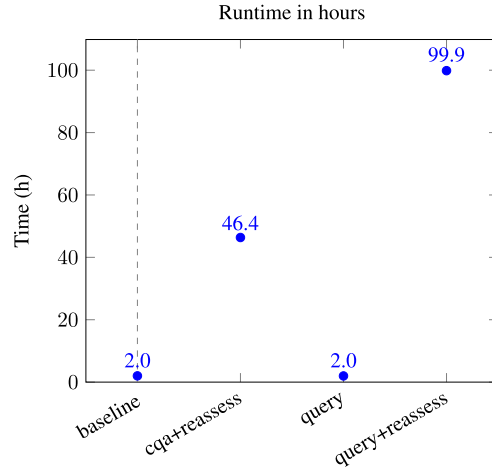


Fig. 14. Runtime of the baseline and its variants over the 20 oriented pairs of ontologies.

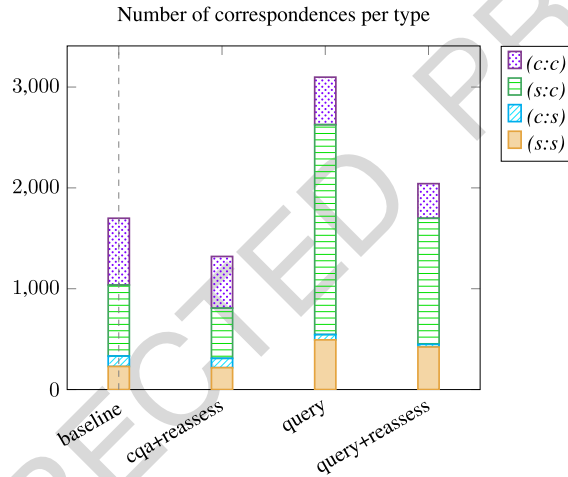


Fig. 15. Number of correspondence per type for the baseline, the variant which generates queries (*query*) and their equivalent variants with similarity reassessment based on counter-examples.

5.3.5. Similarity reassessment with counter-examples

The baseline, the *query* variant and their equivalent were run with a similarity reassessment phase. The runtime of the variants is presented in Figure 14. Figure 15 shows the number of correspondences per type output by the baseline and its variants. The results of this evaluation are presented in Figure 16.

The reassessment phase (finding counter-examples) increases the runtime by far, especially when running queries. It took **46.4 hours** to run the *cqa+reassess* approach and **99.9 hours** to run the *query+reassess* over the 20 pairs of ontologies when it only took **2.0 hours** for the baseline or *query* versions. The baseline approach and the generated query variants have approximately the same runtime over the 20 pairs of ontologies. However, for a similar runtime, the results of the approach with the CQA are better than those with the generated queries.

As expected, the reassessment phase decreases the number of correspondences as they are filtered. It entails an increase in Precision. The Precision of *cqa+reassess* is between 8% and 15% higher than that of the baseline. The Precision of *query+reassess* is between 6% and 17% higher than that of the *query* variant.

The CQA Coverage remains the same for the baseline and *cqa+reassess*. The CQA Coverage score of *query+reassess* is about 3% lower than that of *query*. As more specific correspondences are preferred over more general ones during the similarity reassessment phase, it leaves fewer possibilities during the rewriting phase.

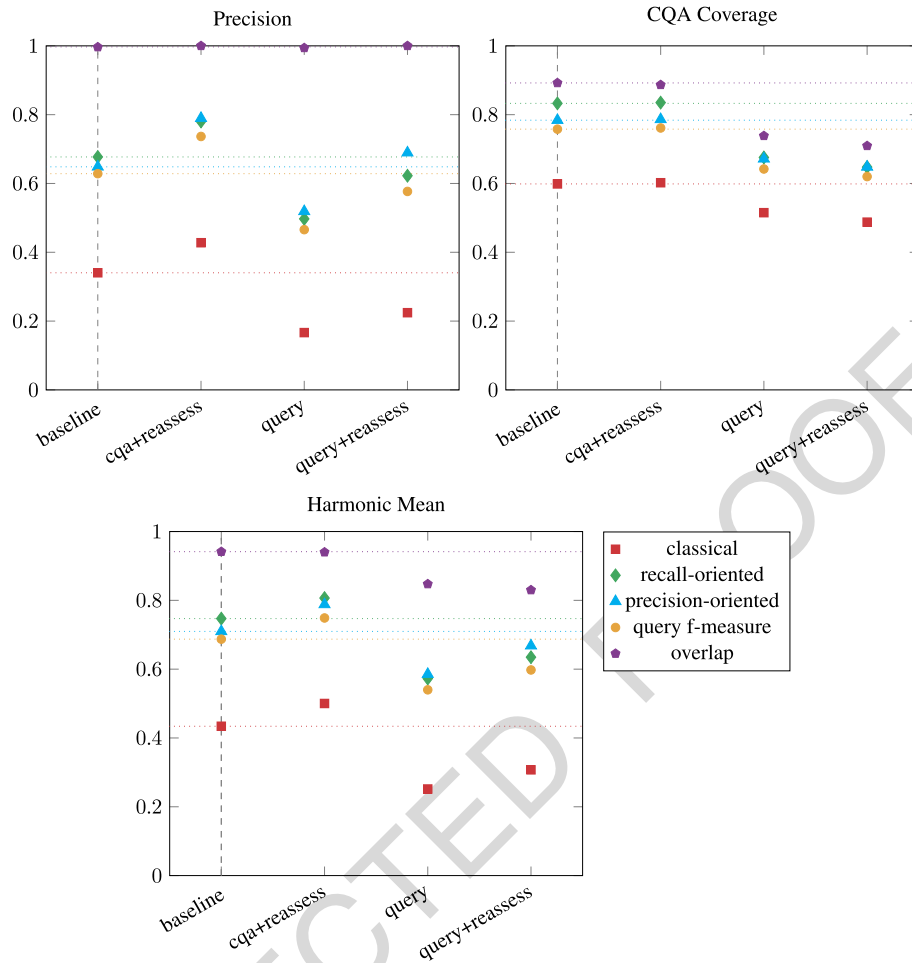


Fig. 16. Results for the baseline, the variant which generates queries (query) and their equivalent with a counter-example-based similarity reassessment.

5.4. Comparison with existing approaches

The generated alignments for the Conference were compared with three reference alignments (two task-oriented alignments which vary in the types of correspondences and expressiveness – query rewriting alignment set and ontology merging alignment set – and the simple reference alignment from the OAEI Conference dataset) and two complex alignments generated from existing approaches (Ritze 2010 and AMLC):

Query rewriting the query rewriting oriented alignment set⁶ from [26] – 10 pairs of ontologies

Ontology merging the ontology merging oriented alignment set⁶ from [26] – 10 pairs of ontologies

ra1 the reference simple alignment⁷ from the OAEI conference dataset [37] – 10 pairs of ontologies

Ritze 2010 the output alignment⁸ from [23] – complex correspondences found on 4 pairs of ontologies

AMLC the output alignment⁹ from [9] – output alignments between 10 pairs of ontologies

⁶<https://doi.org/10.6084/m9.figshare.4986368.v7>

⁷<http://oaei.ontologymatching.org/2018/conference/>

⁸<https://code.google.com/archive/p/generatingcomplexalignments/downloads/>

⁹<http://oaei.ontologymatching.org/2018/results/complex/conference/>

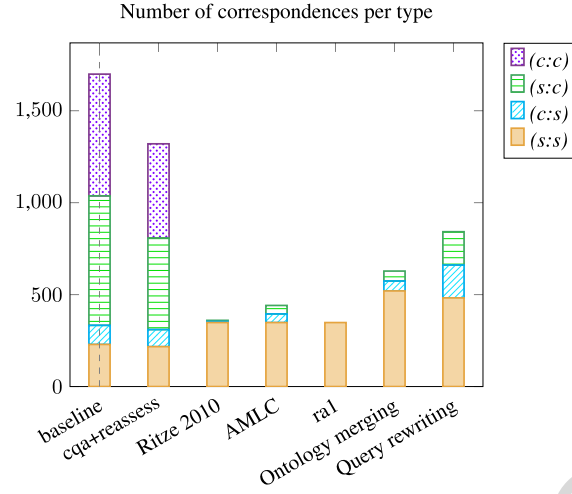


Fig. 17. Number of correspondence per type for the proposed approach, reference alignments and complex alignment generation approaches. The alignments of Ritze 2010 and AMLC include ra1.

The two approaches have been chosen because their implementations are available online and they output alignments in EDOAL. Ritze 2010 [23] and AMLC [9] both require simple alignments as input. They were run with ra1 as input. ra1 has then been added to Ritze 2010 and AMLC for the CQA Coverage evaluation. The Precision evaluation was made only on their output (ra1 correspondences excluded). Ritze 2010 took **58 minutes** while AMLC took about **3 minutes** to run over the 20 pairs of ontologies. Even though these two approaches are similar, this difference of runtime can be explained by the fact that Ritze 2010 loads the ontologies and parses their labels for each pattern while AMLC only loads the ontologies once. Moreover, Ritze 2010 covers 5 patterns while AMLC only covers 2. Some refactoring was necessary so that the alignments could be automatically processed by the evaluation system. The ra1 dataset had to be transformed into EDOAL, instead of the basic alignment format. The Alignment API could not be used to perform this transformation as the type of entity (class, object property, data property) must be specified in EDOAL. The Ritze 2010 alignments used the wrong EDOAL syntax to describe some constructions (*AttributeTypeRestriction* was used instead of *AttributeDomainRestriction*). The AMLC alignments were not parsable because of RDF/XML syntax errors. The entities in the correspondences were referred to by their URI suffix instead of their full URI (e.g., *Accepted_Paper* instead of http://ekaw#Accepted_Paper). Some correspondences were written in the wrong way: the source member was made out of entities from the target ontology and the target member was made out of entities from the source ontology. As the evaluation of these alignments was manual so far in the OAEI complex track, these errors had not been detected. The alignments' syntax has been manually fixed so that they can be automatically evaluated.

Figure 17 shows the number of correspondences per type over the 20 pairs of ontologies. These alignments were not directional so their number of (s:c) and (c:s) correspondences are identical.

Figure 18 shows the results of the baseline approach (*baseline*), the baseline approach with counter-example-based similarity reassessment (*cqa+reassess*), and the compared alignments. The Precision results should be considered carefully. First of all, the relation of the correspondence is not considered in this score: all correspondences are compared as if they were equivalences. The Ontology merging and Query rewriting alignments contain a lot of correspondences with subsumption relations so their classical Precision score is lower than the percentage of correct correspondences it contains. Second, the precision of the alignments is considered to be between the classical Precision and the percentage of correspondences whose members are either overlapping or both empty (*not disjoint*) due to the way the ontologies were populated.

Another limitation of the Precision score is related to the correspondences whose members are not populated in the dataset. For instance, $\langle \text{cmt:Preference}, \text{conference:Review_preference}, \equiv \rangle$ is a correct correspondence that was not detected as such in the Precision evaluation. The review preference of a reviewer for a paper was not part of the CQA for the population process. There is therefore no instance for either member of the correspondence.

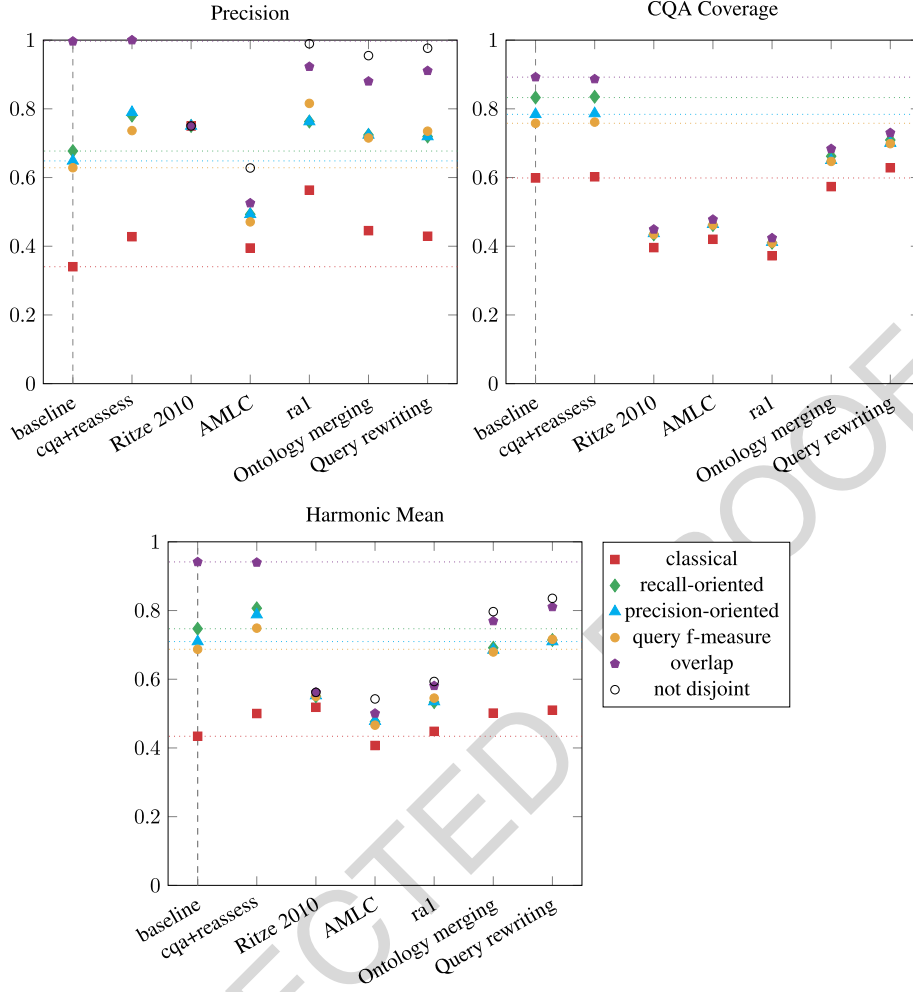


Fig. 18. Results of the proposed approach, reference alignments and complex alignment generation approaches.

To compensate for these errors, we use the *not disjoint* scoring metric in the Precision evaluation. The score for a correspondence is 1 when the members are overlapping or both empty and 0 otherwise. This metric gives the upper bound of the precision of an alignment. When calculating the Harmonic Mean of CQA Coverage and Precision, the *overlap* CQA Coverage was used with the *not disjoint* Precision score to give an upper bound. Indeed, in the CQA Coverage, the source query will never return empty results.

The CQA Coverage of Ritze 2010 and AMLC is higher than that of ra1 which they include. Overall, the CQA Coverage of the other alignments (Ontology Merging, Query Rewriting, ra1, Ritze 2010, and AMLC) is lower than the score of our approach. Indeed, ra1 only contains simple equivalence correspondences, Ritze 2010 and AMLC are mostly restrained to finding $(s:c)$ class expressions correspondences (and therefore do not cover binary CQA). The Ontology merging and query rewriting alignments are limited to $(s:c)$, $(c:s)$ correspondences.

Globally, the Query rewriting alignment outperforms the Ontology merging in terms of CQA Coverage except for the *edas-confOf* pair. In the Ontology merging alignments, unions of properties were separated into individual subsumptions which were usable by the rewriting system. In the Query rewriting alignment, the subsumptions are unions.

CANARD obtains the best CQA Coverage scores except for the classical CQA Coverage where the Query rewriting alignment is slightly better (0.62 vs. 0.60). It can generate $(c:c)$ correspondences which cover more CQA than the other alignments limited to $(s:s)$, $(s:c)$ and $(c:s)$.

Table 4
Number of taxa and plant taxa in each knowledge base of the track, in its original and reduced version

Version	AgronomicTaxon	AgroVoc	DBpedia	TaxRef-LD
Taxa (original)	32	8,077	306,833	570,531
Plant taxa (reduced)	32	4,563	58,257	47,058

The Precision of our approach is overall lower than the Precision of reference alignments (considering that their Precision score is between the classical and not disjoint score). Ritze 2010 only outputs equivalent or disjoint correspondences. Its Precision score is therefore the same (0.75) for all metrics. AMLC achieves a better classical Precision than our baseline approach but contains a high number of disjoint correspondences (37% of all the output correspondences had members whose instance sets were disjoint).

Overall, as expected, the Precision scores of the reference alignments are higher than those output by the matches. Our approach relies on CQAs and for this reason, it gets higher CQA Coverage scores than Ritze 2010 and AMLC. Moreover, these two matches both rely on correspondence patterns which limit the types of correspondences they can generate.

5.5. Evaluation on taxon

The Taxon dataset is composed of 4 ontologies that describe the classification of species: AgronomicTaxon [25], AgroVoc [5], DBpedia [3] and TaxRef-LD [16]. The CQA used in this evaluation are the one presented in [31] which were manually written from AgronomicTaxon CQs [25]. The ontologies are populated and their common scope is plant taxonomy. Their particularity, however, is that within the same dataset, the same information can be represented in various ways but irregularly across instances. For this reason, creating a set of references and exhaustive CQA is not easily feasible.

The knowledge bases described by these ontologies are large. The English version of DBpedia describes more than 6.6 million entities alone and over 18 million entities.¹⁰ The TaxRef-LD endpoint contains 2,117,434 instances¹¹ and the AgroVoc endpoint 754,874¹¹. AgronomicTaxon has only been populated with the wheat taxonomy and only describes 32 instances. The approach has been run on the distant SPARQL endpoints but server exceptions have been encountered, probably due to an unstable network connection or an overload of the servers. A reduced version of the datasets was then stored on a local machine to avoid these problems. The reduced datasets contain all the plant taxa and their information (surrounding triples, annotations, *etc.*) from the SPARQL endpoint of the knowledge bases. Table 4 shows the number of plant taxa in each knowledge base. Even though the number of instances was reduced, the knowledge bases are still large-scale.

The approach is run with the following settings: Levenshtein threshold: 0.4; Number of support answers: 1 and 10 (two runs); Instance matching: look for existing links (*owl:sameAs*, *skos:closeMatch*, *skos:exactMatch*) and if no target answer is found like that, perform an exact label match; No counter-example reassessment (computing the percentage of counter-examples would last too long on this dataset). The generated correspondences have been manually classified as equivalent, more general, more specific, or overlapping. The classical, recall-oriented, precision-oriented, and overlap scores have been calculated based on this classification.

5.5.1. Evaluation results

The number of correspondences per type is shown in Figure 19. The correspondences have been manually classified as equivalent, more general, more specific, or overlapping. The classical, recall-oriented, precision-oriented, and overlap scores have been calculated based on this classification. The results are shown in Figure 20.

Overall, the classical Precision and CQA Coverage scores are rather low. The Precision of the approach with 1 or 10 support answers is approximately the same. However, the CQA Coverage is higher with 10 instances. In comparison with the Conference dataset, this can be explained by the differences in population between the

¹⁰Statistics from the 2016-10 release <https://wiki.dbpedia.org/develop/datasets/dbpedia-version-2016-10>.

¹¹Tested on 2019/04/12.

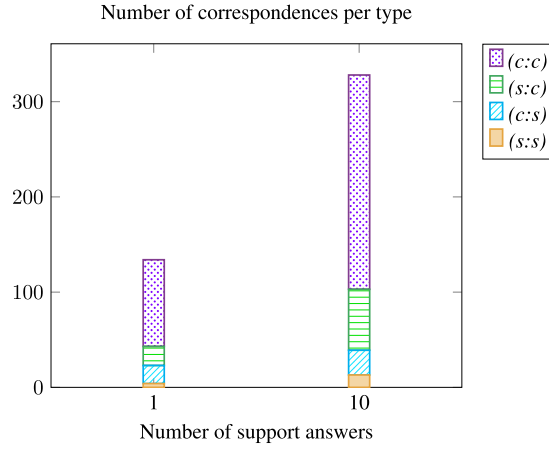


Fig. 19. Number of correspondence per type for the approach with 1 and 10 support answers on the taxon dataset.

knowledge bases and the uneven population of a knowledge base in itself. We guess that the more support answers the approach takes, the better its CQA Coverage will be when dealing with unevenly populated ontologies.

The uneven population of some knowledge bases leads to missing correspondences. For example, the entity *agronto:hasTaxonomicRank* is not represented for every instance of Agrovoc. *agrovoc:c_35661* which is the *Asplenium* genus taxon has no *agronto:hasTaxonomicRank* property. When this instance was used as a support instance, it could not lead to the detection of a correspondence involving its rank. When running our matching approach with only 1 support instance, using this instance would result in an empty set of correspondences for some CQAs. Consequently, the CQA Coverage is globally higher for the approach with 10 support answers.

The particularity of a dataset about species taxonomy is that two taxa are likely to share the same scientific name. Our exact label match strategy is therefore rather suited for such a dataset. In some cases, however, it introduced noise. For example, confusion was made between *wheat* the plant taxon, and *wheat* the consumable good, or between a *division*, part of an administrative structure, and the taxonomic rank *division*.

The Levenshtein-based string similarity brings the noise. For example, the correspondence $\langle \text{agrotaxon:GenusRank}, \exists \text{agronto:produces.}\{\text{agrovoc:c_8373}\}, \equiv \rangle$ whose target member represents all the agronomic taxa which produce wheat has been output. This is due to the string similarity between the Malay label of wheat “*Gandum*” and the English label “*Genus*” of the *agrotaxon:GenusRank* class. We could have chosen to compare labels in the same language together but sometimes, the language of a label was missing, and sometimes the scientific name was either tagged as English or Latin.

The total runtime over the 12 pairs of ontologies was **99,297 s (27.6 h)** for the approach with 1 support instance and **113,474 s (31.5 h)** for the approach with 10 support answers. The runtime per pair of ontologies is detailed in Table 5. Three factors explain the different runtime over the pairs of ontologies in Table 5.

Query difficulty Some CQA was long to run on large knowledge bases, in particular those involving the union of properties.

Percentage of source common instances The number of taxa instances can be different between knowledge bases. AgronomicTaxon and DBpedia share 22 instances. When AgronomicTaxon, which has only 32 instances is matched to DBpedia, finding a common instance between the two is rather easy because about 68% of its instances have an equivalent in DBpedia. The other way around is way harder because only 0.04% of DBpedia taxa instances have an equivalent in AgronomicTaxon.

Existence of instance links When no explicit instance links exist between two knowledge bases, all the source instances are explored and the exact label match is performed. This can take a lot of time according to the size of the target knowledge base.

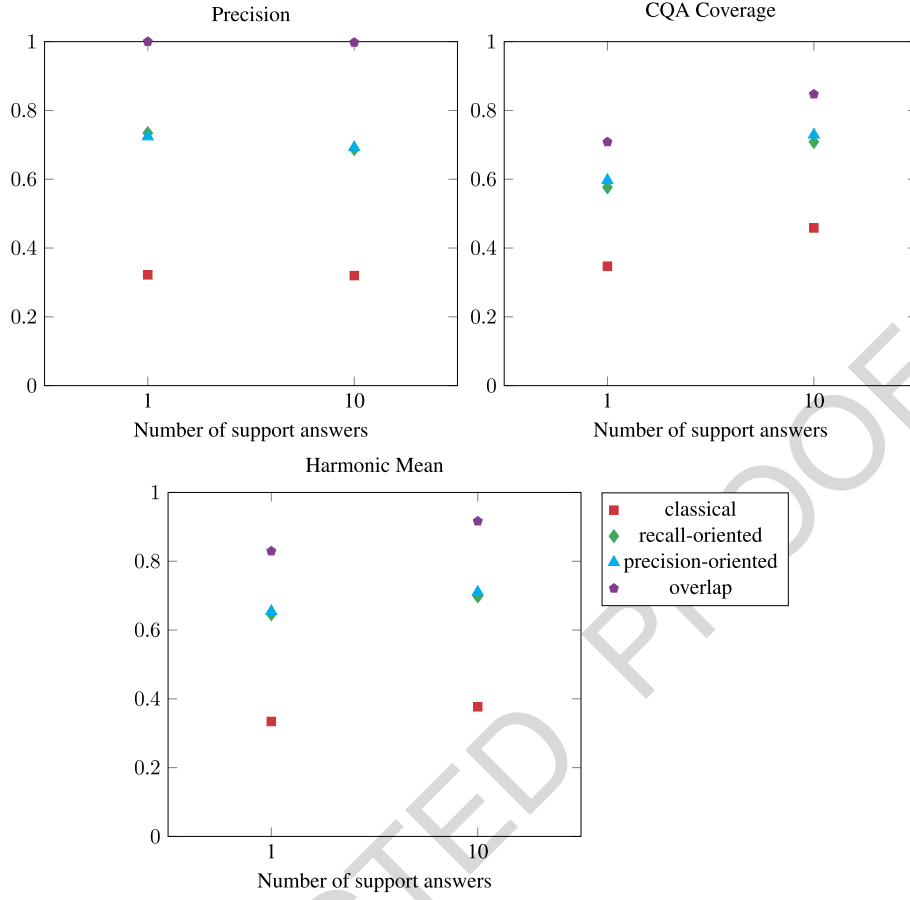


Fig. 20. Results of the approach with 1 and 10 support answers on the taxon dataset.

5.6. Comparison on OAEI systems

Section 5.4 has presented a comparison of CANARD with existing systems in the OAEI Conference ontologies. This section provides a comparison of the complex systems participating in the OAEI Complex Track since its creation in 2018. These results are presented in Table 6. Results involving simple alignments are not reported here. With respect to the systems (some already introduced above), AROA [41] is based on association rule mining and implements the algorithm FP-growth to generate complex alignments. MatchaC [8] is the successor of AMLC and has introduced machine learning strategies.

Overall, results and participation are still modest (only 3 participants in all campaigns) and as explained in this paper, CANARD can only deal with populated datasets.

For the (non-populated) **Conference** dataset (manual evaluation on a subset of the original Conference dataset), only AMLC (and its successor MatchaC) was able to deal with the task, with results that have been similar over the years. Still, the performance is far from those obtained with simple alignments. AMLC maintained its F-measure over the campaigns. For the **Populated Conference**, introduced in 2019, CANARD achieved close results to AMLC in terms of Coverage and it maintains its performance over the campaigns (the detailed results are described in Section 5.3).

With respect to the **Hydrography** sub-track case, only AMLC can generate (few) correct complex correspondences, with fair results in terms of precision, to the detriment of recall. In **GeoLink**, AMLC, AROA, and CANARD were able to output correspondences, for the version of the dataset having instances, with a higher number of complex correspondences being found by AROA and CANARD (which report close results). In 2020, a fully populated

Table 5

Runtime (s) of our approach on each pair of ontologies. These measures are based on a single run

Source	Target			
	AgronomicTaxon	AgroVoc	DBpedia	TaxRef-LD
1 sup. inst.				
AgromaticTaxon	–	67	4	421
AgroVoc	747	–	238	27,776
DBpedia	50,542	2,733	–	2,477
TaxRef-LD	4,517	5,758	4,017	–
10 sup. inst.				
AgromaticTaxon	–	1,084	1,019	753
AgroVoc	1,173	–	220	29,351
DBpedia	52,214	4,813	–	5,062
TaxRef-LD	4,718	8,005	5,062	–

version has been introduced (**Populated GeoLink**), reporting as expected the same results as the previous version of GeoLink in 2019. In the **Populated Enslaved** sub-track, CANARD is outperformed by AMLC and AROA. AROA found the largest number of complex correspondences among the three systems, while the AMLC outputs the largest number of simple correspondences. For **Taxon**, CANARD is the only system that can deal with the high heterogeneity of the task and can retrieve correct complex correspondences (with a high performance considering at least one common instance in the coverage results). Overall systems still privilege precision and detriment of recall (except AMLC in 2018), leaving room for several improvements in the field.

Concerning OAEI 2023, MatchaC, LogMap, and LogMapLite have registered to participate. While LogMapLite and LogMap are dedicated to generating simple correspondences, only LogMap was able to generate nonempty (simple) alignments. MatchaC, the only system specifically designed to generate expressive correspondences in OAEI 2023, had some issues dealing with the datasets and was not able to generate valid alignments. Unfortunately, in 2023, several datasets have also been discontinued (Hydrography, GeoLink, Populated GeoLink, Populated Enslaved, and Taxon). While the last participation of CANARD in OAEI campaigns was in 2020, improvements so far have addressed runtime, as reported in this paper. We plan to come back to the campaigns with new improvements in the way expressive correspondences are generated.

5.7. Qualitative evaluation

In order to provide a more qualitative analysis of the generated alignments, we analyzed the correct correspondences uniquely identified by CANARD whereas missed by other systems. The alignments come from the outputs of the systems in the 2020 OAEI corresponding to Table 6, on the Populated Conference dataset. The choice of this dataset is motivated by the fact that we have reference competency questions and access to the evaluation system for providing such an analysis. The analysis involves 8 pairs of alignments: conference-confOf, conference-ekaw, confOf-conference, confOf-ekaw, edas-ekaw, ekaw-conference, ekaw-confOf, ekaw-edas. The number of correct correspondences in each pair found only by CANARD is presented in Table 7.

The majority of correspondences that CANARD can generate, and that AMLC is not able to, are related to properties in the target entity. They are correspondences between properties or correspondences between class constructors and restrictions that apply to properties. Such a correspondence is the correspondence between the property *reviewerOfPaper* and a property restriction of *contributes* with the domain of *Review*. Another similar correspondence is between the property *hasReview* and the inverse of *reviews* with range restriction to *Review*. The correct correspondences that CANARD can find and AMLC is not able to identify for the pair confOf-ekaw are listed in Table 8. It shows that all types of target entities (ent2_type) are related to properties since edoal:Relation relates to properties and edoal:AttributeDomainRestriction are restrictions applied to property domains.

Table 6

Results of the complex track in OAEI for systems generating complex correspondences. R[P|F|R] refers to relaxed precision, f-measure, and recall, respectively. P in populated conference refers to (classical - not disjoint) precision and coverage to (classical - query F-measure) coverage. In taxon, (classical - overlap) precision and (classical - overlap) coverage. * Indicates tracks not available in that campaign, 1 indicates that different evaluation metrics have been applied and for which the results are not comparable (in 2018, the correspondences in taxon, have been evaluated in terms of precision and their ability to answer a set of queries over each dataset). - Means no results. For MatchaC the results are not fully comparable as the results have been taken from the author's paper

Matcher	Conference			Populated Conference		Hydrography			GeoLink			Populated GeoLink			Populated Enslaved			Taxon	
	P	F	R	P	Cov.	RP	RF	RR	RP	RF	RR	RP	RF	RR	RP	RF	RR	P	Cov.
2018																			
AMLC	.54	.42	.34	*	*	-	-	-	-	-	-	*	*	*	*	*	*	-	-
CANARD	-	-	-	*	*	-	-	-	*	*	*	*	*	*	*	*	*	.20 ¹	.13 ¹
2019																			
AMLC	.31	.34	.37	.30-.59	.46-.50	.45	.10	.05	.50	.32	.23	*	*	*	*	*	*	-	-
AROA	-	-	-	-	-	-	-	-	.87	.60	.46	*	*	*	*	*	*	-	-
CANARD	-	-	-	.21-.88	.40-.51	-	-	-	.89	.54	.39	*	*	*	*	*	*	.08-.91	.14-.36
2020																			
AMLC	.31	.34	.37	.23-.51	.26-.31	.45	.10	.05	.50	.32	.23	.50	.32	.23	.73	.40	.28	.19-.40	0
AROA	-	-	-	-	-	-	-	-	-	-	-	.87	.60	.46	.80	.51	.38	-	-
CANARD	-	-	-	.25-.88	.40-.50	-	-	-	-	-	-	.89	.54	.39	.42	.19	.13	.16-.57	.17-.36
2021																			
AMLC	.31	.34	.37	*	*	.49	.08	.04	.49	.30	.22	.49	.30	.22	.46	.18	.12	*	*
AROA	-	-	-	-	-	-	-	-	-	-	-	.87	.60	.46	.80	.38	.51	*	*
2022																			
MatchaC	.31	.34	.17	*	*	.49	.04	.08	.49	.22	.30	.49	.22	.30	.46	.12	.18	*	*
2023																			
MatchaC	-	-	-	-	-	*	*	*	*	*	*	*	*	*	*	*	*	*	*

Table 7

Number of correspondences that CANARD finds and AMLC does not find

Type	conf-confOf	conf-ekaw	confOf-conf	confOf-ekaw	edas-ekaw	ekaw-conf	ekaw-confOf	ekaw-edas
Simple	3	5	10	4	9	10	2	6
Complex	1	17	16	10	14	41	1	12

Table 8

EDOAL (correct) s:c correspondences found by CANARD for the pair confOf-ekaw. Prefix edoal refers to the namespace {<http://ns.inria.org/edoal/1.0/#>} and confof to the namespace <http://confOf#>

ent1_type	entity1	entity2	ent2_type	constructor2	relation
edoal:Class	confOf:Social_event	ekaw:partOfEvent	edoal:AttributeDomainRestriction	edoal:exists	=
edoal:Class	confOf:Conference	ekaw:partOfEvent	edoal:AttributeDomainRestriction	edoal:exists	=
edoal:Class	confOf:Conference	ekaw:partOf	edoal:AttributeDomainRestriction	edoal:exists	=
edoal:Class	confOf:Conference	ekaw:hasEvent	edoal:AttributeDomainRestriction	edoal:exists	=
edoal:Relation	confOf:location	ekaw:heldIn	edoal:Relation	edoal:compose	=
edoal:Relation	confOf:location	ekaw:locationOf	edoal:Relation	edoal:compose	=
edoal:Class	confOf:Poster	ekaw:reviewerOfPaper	edoal:AttributeDomainRestriction	edoal:exists	=
edoal:Class	confOf:Poster	ekaw:hasReviewer	edoal:AttributeDomainRestriction	edoal:exists	=
edoal:Class	confOf:Topic	ekaw:coversTopic	edoal:AttributeOccurenceRestriction	edoal:value	=
edoal:Class	confOf:Event	ekaw:partOfEvent	edoal:AttributeDomainRestriction	edoal:exists	=

5.8. Improvements in performance

As reported above, CANARD has a higher runtime in specific settings. To address this weakness, improvements in its implementation have been carried out.¹² The most expensive steps in the current implementation concern steps 4 to 7. The first issue is the text search done in Jena without a text index, which has to be mitigated with a text index configuration. The second issue relates to the many requests to the server and since the main line of communication is done through HTTP, some steps like the socket communication and HTTP request parsing slow the system. One example of such a case is in step 4 where the system looks for shared instances between the ontologies. When no instances are found, similar instances are queried using the SPARQL filter with a regex search that is slow without a full-text search index. As an exact string comparison is made in this step, the first improvement is the use of a map structure to store the triples in memory, and with this structure is possible to query similar instances by text in constant time without the HTTP overhead. With this structure, the majority of queries executed in Apache Jena can be replaced by the map lookup. This structure improves performance on steps 5, 6, and 7 as they depend on these functions to operate. The improvements come at the cost of increased memory usage as indexes need to be stored for subjects, predicates, and objects.

Another step with high running time is the subgraph query for binary CQs. This step needs to find similar paths in the ontology structure using an iterative path-finding algorithm. However, the populated ontologies can have imbalanced structures. For example, a Paper class can have thousands of instances that need to be verified in each step even for small paths (size 5 for example). This issue is still not addressed in terms of the number of comparisons done in path-finding. Still, since the similarity calculation performed in each step uses the indexed map it is faster than the original implementation.

To evaluate the impact of these modifications, the base and improved version performances were compared in the populated Conference dataset with an exact label match approach. The base version was run in one alignment pair between CMT and Conference with 34 CQAs and takes approximately 6:18 hours to run with one threshold value. The improved version was run in 4 alignment pairs between CMT and Conference, Conference and ConfOf, ConfOf and Edas, and Edas and ekaw with 213 CQAs in total with 9 thresholds in the range from 0.1 to 0.9. The improved system runs the 4 pairs in 25 minutes and runs the CMT and conference pair in 52 seconds approximately.

¹²The improved version is available at <https://gitlab.irit.fr/melodi/ontology-matching/complex/canarde>.

6. Discussion

This section discusses the strengths and weaknesses of CANARD. First, even though the similarity metric in its current version is naive, the results of the approach are quite good (the query f-measure Harmonic Mean score of the baseline approach is 0.70). The approach is rather CQA Coverage-oriented, as it will try to output a correspondence for each source CQA. The values of the CQA Coverage are overall higher than the Precision values. The baseline achieves a classical CQA Coverage of 0.60 which means that 60% of the CQA have been covered with a strictly equivalent match by our approach while its classical Precision score is only 0.34. Using existing links gives better results than exact label matches. The use of CQA improves the Precision and CQA Coverage performance of the approach concerning queries. The counter-example exploration (similarity reassessment phase) increases significantly Precision, to the detriment of runtime. In comparison with all the other matching approaches evaluated, our approach has high CQA Coverage scores (Populated Conference and Taxon datasets). Overall, CANARD can deliver complex correspondences for all evaluated (populated) datasets in the OAEI, with a higher number of complex(s:c and c:c) correspondences. It would be interesting to compare our approach with extensional approaches such as [12,17,18,35] (whose implementation was not available) even though all of them are limited to (s:c) and (c:s) correspondences. The experiment on the Taxon dataset showed that our approach is one of the few that can perform on large knowledge bases. CANARD however depends on regularly populated knowledge bases and the quality instance links (what can explain the lower results in the Enslaved dataset with respect to the other systems).

In deep, the evaluation described in Section 5 helped answer the research questions:

What is the impact of the label similarity metric on the approach? The label similarity metric directly impacts the approach: the more constraining it is, the better the Precision but the worse the CQA Coverage. In the experiment of Section 5.3.1, we only changed the threshold of this metric. However, it would be interesting to investigate linguistic metrics and techniques in this phase.

Is one common instance per Competency Question for Alignment enough evidence to generate complex correspondences? In the experiments on the Populated Conference benchmark and the Taxon dataset, the approach based on only one common instance could generate complex correspondences. While in the Populated Conference dataset, the results with one support answer are slightly higher than with more support answers, in the Taxon dataset, they are lower. This can be explained by the irregular population of some Taxon dataset ontologies as well as the existence of inaccurate instance links. These aspects are also discussed in the next research question.

What is the impact of the number of support answers on the alignment quality? The impact of the number of support answers depends on the ontology population. In the experiment on the Taxon dataset, using 10 support answers instead of 1 improved the quality of the alignment. The reason is that the ontologies are not all regularly populated. The Precision score was about the same for 1 or 10 support answers while the CQA Coverage scores are about 12% higher with 10 support answers than with 1. In the Conference dataset which is regularly populated, using more support answers reduced the Precision score because noise was introduced. When dealing with many support answers, the noisy correspondences could be filtered out based on their frequency. For example, the formula $\exists \text{conference:has_the_last_name.}\{\text{"Benson"}\}$ only appears for one support instance of Person whereas *conference:Person* appears for all support answers. However, it was a choice in the approach design to not disregard “accidental” formulae (those that only appear for 1 answer and not in the other answers) because unevenly populated datasets may be faced with this problem. For example, in DBpedia, the taxonomic rank of a taxon can be represented in different ways: the label of a property (e.g., a taxon is the *dbo:genus* of another taxon or has a *dbp:genus* literal), a link to the rank instance (e.g., link to *dbp:Genus*), or the presence of a rank authority (e.g., *dbp:genusAuthority*). The problem is that all the genus instances do not share the same representation. It is possible that among the genus rank instances, only one is represented as a genus rank thanks to the *dbp:genusAuthority*. This may seem statistically accidental but it is relevant to our problem.

What is the impact of the quality of the instance links on the generated alignments quality? If the links are expressed and not erroneous, the generated alignment will have better Precision and CQA Coverage. If wrong links are used, as in the experiment with exact label matches, a lot of noise is introduced and the Precision of the

alignment decreases. The CQA Coverage score also decreases because the noise can prevent correct support answers from being found and all the output correspondences for a given correspondence can be erroneous. The quality of the instance links impacts the Precision and CQA Coverage scores of our approach. This highlights the need for effective instance-matching systems and the disambiguation of existing links.

Can Competency Questions for Alignment improve the Precision of generated alignments? Both the Precision and CQA Coverage scores are higher when the approach relies on CQA. The baseline and the *cqa+reassess* variants obtain a Precision score on average 15% above that of their generated query variants (*query* and *query+reassess*). The CQA Coverage also increases by an average of 14% because the CQA helps generate (*c:c*) correspondences that are relevant to the user (and to the evaluation). However, as part of the input CQA is used for the calculation of the CQA Coverage score, the evaluation is somewhat biased. In a user's need-oriented scenario, nonetheless, this evaluation makes sense: if users input their needs into a matcher, they may expect an output alignment that covers them well.

Does similarity reassessment based on counter-examples improve the quality of the generated alignments? When comparing the results of the baseline approach with the *cqa+reassess* variant which reassesses the similarity based on counter-examples, the CQA Coverage remains the same while the Precision is improved. The Precision of the *cqa+reassess* variant is between 8% and 15% higher than that of the baseline. The Precision of the *query+reassess* variant is between 6% and 17% higher than that of the *query* variant while its CQA Coverage is 3% lower.

What is the impact of the CQA on the type of output correspondence? Overly complex correspondences can be introduced in the alignment because of the way the approach uses the input CQA. We counted that about 14% of the (*c:c*) correspondences output by the baseline approach are overly complex, which means that they could be decomposed into simple correspondences. This comes from the translation of the input CQA into a DL formula without any analysis or decomposition of its elements. Moreover, the approach outputs more (*s:c*) and (*c:c*) correspondences than (*s:s*) and (*c:s*) which shows a tendency to output more complex than simple correspondences.

7. Related work

Classification of the approach CANARD is positioned using the characteristics in [27]. CANARD can generate (*s:s*), (*s:c*), and (*c:c*) correspondences depending on the shape of the input CQA. It focuses on correspondences with logical constructors. The approach relies on a path to find the correspondences for binary CQAs. For the unary CQAs, we classify CANARD as *no structure* because it does not explicitly rely on atomic or composite patterns. The source member form is fixed before the matching process by the CQA but the target member form is unfixed, therefore we classify it as *fixed to unfixed*. CANARD relies on ontology and instance-level evidence. CANARD fits in the formal resource-based because it relies on CQA and existing instance links, its implementation is string-based because of the label similarity metric chosen (see Section 5.1), and it is also graph-based and instance-based.

Comparison to other matching approaches The matching approaches generating expressive correspondences involve different techniques such as relying on templates (called patterns) and/or instance evidence. The approaches in [22,23] apply a set of matching conditions (label similarity, datatype compatibility, etc.) to detect correspondences that fit certain patterns. The approach of [24] uses the linguistic frames defined in FrameBase to find correspondences between object properties and the frames. KAOM [13] relies on *knowledge rules* which can be interpreted as probable axioms. In [38], a structural matching approach (FCM-Map) adopts the Formal Concept Analysis (FCA) method to find complex correspondence candidates. The approaches in [18,35] use statistical information based on the linked instances to find correspondences fitting a given pattern. The one in [20] uses a path-finding algorithm to find correspondences between two knowledge bases with common instances. The one in [12] iteratively constructs correspondences based on the information gained from matched instances between the two knowledge bases. [9] relies on lexical similarity and structural conditions to detect correspondence patterns, close to [22]. As introduced in Section 5.4, AROA [41] (Association Rule-based Ontology Alignment s) is based on association rule mining and implements the algorithm FP-growth to generate complex alignments. Generated alignments are filtered out

using simple and complex patterns. As CANARD, it also depends on populated datasets. More recently, in [2], the proposal combines an ontology fuzzification process with an embedding strategy. A fuzzy ontology has weights to describe the level of fuzzy membership of concepts and properties. The approach generates a fuzzy version of the ontology concepts, and later a graph embedding approach based on RDF2Vec, which traverses the graph in random walks, and generates sentences. It uses an embedding strategy to generate the final embeddings that are used to compare similarities between concepts. Then, a stable marriage-based alignment extraction algorithm is applied to establish correspondences. None of these approaches involve, however, the user before or during the matching process. As in [12,17,18,20,35,41], CANARD relies on common instances. Differently from them, it does not rely on correspondence patterns. Finally, CQA has not been adapted nor used for matching.

SPARQL CQA In our approach, CQA is used as basic pieces of information that will be transformed as source members of correspondences. Their formulation in a SPARQL query over the source ontology is a limitation of the approach as a user would need to be familiar with SPARQL and the source ontology. However, in the scenario where someone wants to publish and link a knowledge base he or she created on the LOD cloud, this person is already familiar with the source ontology and can reuse the CQ of their own ontology. In other cases, one could rely on question-answering systems that generate a SPARQL query from a question in natural language. This kind of system is evaluated in the QALD open challenge [33].

Generalisation process Ontology matching approaches relying on instances infer general statements, *i.e.*, they perform a generalisation.¹³ This is the principle of *machine learning* in general and methods such as *Formal Concept Analysis* [10] or *association rule mining* [1]. These generalisation processes however require a considerable amount of data (or instances). Approaches such as the ones from [12,17,18,35] rely on large amounts of common ontology instances for finding complex correspondences. Few exceptions in ontology matching rely on a few examples. For instance, the matcher of [36] relies on example instances given by a user. With this information, the generalization can be performed on a few examples. The idea behind our approach is to rely on a few examples to find general rules that would apply to more instances. In particular, the generalization phase of our approach is guided by the CQA labels. Thanks to that, only one instance is sufficient for finding a correspondence. This would apply to knowledge bases that represent different contexts or points of view but whose ontologies are overlapping.

8. Conclusions

This paper has presented a complex alignment generation approach based on CQAs. The CQA defines the knowledge needs of a user over two ontologies. The use of CQAs is both a strength of the approach as it allows for a generalization over a few instances and a limitation as it requires that the user can express her or his needs as SPARQL queries. It depends as well on the quality of the instance matches. The approach can be extended in several directions: one could consider exploring embeddings for similarity calculation or still sophisticated instance-based matching approaches and, alternatively, conditional or link keys (systems generating keys could also benefit from complex correspondences to improve their results); designing a purely T-Box strategy based on both linguistic and semantic properties of the ontologies and CQAs; or still dividing the problem into sub-tasks through ontology partitioning. Also, incoherence resolution systems for complex alignments are scarce. Last but not least, while it is assumed a dependency between CQA (in SPARQL) and correspondence expressiveness, the dependency of CQAs to SPARQL (and their generalisation) should be further investigated.

References

- [1] R. Agrawal, T. Imieliński and A. Swami, Mining association rules between sets of items in large databases, in: *Proceedings of the 1993 ACM SIGMOD Conference*, Washington DC, USA, May 26–28, 1993, Vol. 22, S. Jajodia and P. Buneman, eds, ACM, 1993, pp. 207–216. doi:10.1145/170035.170072.

¹³ ‘They infer general statements or concepts from specific cases’ (Oxford Dictionary, “Generalisation” Retrieved June 3, 2019, from <https://en.oxforddictionaries.com/definition/generalization>).

- [2] H. Akreimi, M.G. Ayadi and S. Zghal, A fuzzy OWL ontologies embedding for complex ontology alignments, in: *Discovery Science – 25th International Conference, DS 2022, Montpellier, France, October 10–12, 2022, Proceedings*, P. Poncelet and D. Ienco, eds, Lecture Notes in Computer Science, Vol. 13601, Springer, 2022, pp. 394–404.
- [3] S. Auer, C. Bizer, G. Kobilarov, J. Lehmann, R. Cyganiak and Z. Ives, DBpedia: A nucleus for a web of open data, in: *The Semantic Web: The 6th International Semantic Web Conference ISWC and the 2nd Asian Semantic Web Conference ASWC*, Busan, Korea, November 2007, K. Aberer, K.-S. Choi, N. Noy, D. Allemang, K.-I. Lee, L. Nixon, J. Golbeck, P. Mika, D. Maynard, R. Mizoguchi, G. Schreiber and P. Cudré-Mauroux, eds, LNCS, Vol. 4825, Springer, Berlin, Heidelberg, 2007, pp. 722–735. doi:[10.1007/978-3-540-76298-0_52](https://doi.org/10.1007/978-3-540-76298-0_52).
- [4] A. Borgida, On the relative expressiveness of description logics and predicate logics, *Artificial intelligence* **82**(1–2) (1996), 353–367. doi:[10.1016/0004-3702\(96\)00004-5](https://doi.org/10.1016/0004-3702(96)00004-5).
- [5] C. Caracciolo, A. Stellato, S. Rajbhandari, A. Morshed, G. Johannsen, J. Keizer and Y. Jaques, Thesaurus maintenance, alignment and publication as linked data: The AGROVOC use case, *International Journal of Metadata, Semantics and Ontologies* **7**(1) (2012), 65. doi:[10.1504/IJMSO.2012.048511](https://doi.org/10.1504/IJMSO.2012.048511).
- [6] M. Ehrig and J. Euzenat, Relaxed precision and recall for ontology matching, in: *Integrating Ontologies '05, Proceedings of the K-CAP 2005 Workshop on Integrating Ontologies*, Banff, Canada, October 2, 2005, CEUR Workshop Proceedings, Vol. 156, CEUR-WS.org, 2005, pp. 25–32.
- [7] J. Euzenat and P. Shvaiko, *Ontology Matching*, 2nd edn, Springer, Berlin, Heidelberg, 2013.
- [8] D. Faria, a. Contreiras, P. Cotovio, P. Eugenio and C. Pesquita, Matcha and matcha-dl results for oaei 2022, in: *Proceedings of the 17th International Workshop on Ontology Matching Co-Located with the 21th International Semantic Web Conference, OM@ISWC 2022*, Hangzhou, China, P. Shvaiko, J. Euzenat, E. Jiménez-Ruiz, O. Hassanzadeh and C. Trojahn, eds, 2022.
- [9] D. Faria, C. Pesquita, B.S. Balasubramani, T. Tervo, D. Carriço, R. Garrilha, F.M. Couto and I.F. Cruz, Results of AML participation in OAEI 2018, in: *Proceedings of the 13th International Workshop on Ontology Matching Co-Located with the 17th International Semantic Web Conference, OM@ISWC 2018*, Monterey, CA, USA, October 8, 2018, P. Shvaiko, J. Euzenat, E. Jiménez-Ruiz, M. Cheatham and O. Hassanzadeh, eds, CEUR Workshop Proceedings, Vol. 2288, CEUR-WS.org, 2018, pp. 125–131.
- [10] B. Ganter, G. Stumme and R. Wille (eds), *Formal Concept Analysis, Foundations and Applications*, Lecture Notes in Computer Science, Vol. 3626, Springer, 2005.
- [11] M. Grüninger and M.S. Fox, Methodology for the design and evaluation of ontologies. International joint conference on artificial intelligence, in: *Workshop on Basic Ontological Issues in Knowledge Sharing*, Vol. 15, 1995.
- [12] W. Hu, J. Chen, H. Zhang and Y. Qu, Learning complex mappings between ontologies, in: *The Semantic Web – Joint International Semantic Technology Conference, JIST 2011, Hangzhou, China, December 4–7, 2011, Proceedings*, J.Z. Pan, H. Chen, H.-G. Kim, J. Li, Z. Wu, I. Horrocks, R. Mizoguchi and Z. Wu, eds, Lecture Notes in Computer Science, Vol. 7185, Springer, 2011, pp. 350–357.
- [13] S. Jiang, D. Lowd, S. Kafle and D. Dou, Ontology matching with knowledge rules, in: *Transactions on Large-Scale Data and Knowledge-Centered Systems XXVIII*, Vol. 28, Springer, 2016, pp. 75–95. doi:[10.1007/978-3-662-53455-7_4](https://doi.org/10.1007/978-3-662-53455-7_4).
- [14] V.I. Levenshtein, Binary codes capable of correcting deletions, insertions, and reversals, *Soviet physics doklady* **10**(8) (1966), 707–710.
- [15] A. Maedche, B. Motik, N. Silva and R. Volz, MAFRA – a mapping framework for distributed ontologies, in: *Knowledge Engineering and Knowledge Management. Ontologies and the Semantic Web, 13th International Conference, EKAW 2002, Sigüenza, Spain, October 1–4, 2002, Proceedings*, A. Gómez-Pérez and V.R. Benjamins, eds, Lecture Notes in Computer Science, Vol. 2473, Springer, 2002, pp. 235–250.
- [16] F. Michel, O. Gargominy, S. Tercerie and C. Faron-Zucker, A model to represent nomenclatural and taxonomic information as linked data. Application to the French taxonomic register, TAXREF, in: *Proceedings of the 2nd International Workshop on Semantics for Biodiversity (S4BioDiv 2017) Co-Located with 16th International Semantic Web Conference (ISWC 2017)*, Vol. 1933, Vienna, Austria, October 2017, A. Algergawy, N. Karam, F. Klan and C. Jonquet, eds, CEUR-WS.org, 2017.
- [17] R. Parundekar, C.A. Knoblock and J.L. Ambite, Linking and building ontologies of linked data, in: *The Semantic Web – ISWC 2010–9th International Semantic Web Conference, ISWC 2010, Shanghai, China, November 7–11, 2010, Revised Selected Papers, Part I*, P.F. Patel-Schneider, Y. Pan, P. Hitzler, P. Mika, L. Zhang, J.Z. Pan, I. Horrocks and B. Glimm, eds, Lecture Notes in Computer Science, Vol. 6496, Springer, 2010, pp. 598–614.
- [18] R. Parundekar, C.A. Knoblock and J.L. Ambite, Discovering concept coverings in ontologies of linked data sources, in: *The Semantic Web – ISWC 2012–11th International Semantic Web Conference – ISWC 2012 – 11th International Semantic Web Conference, Boston, MA, USA, November 11–15, 2012, Proceedings, Part I*, P. Cudré-Mauroux, J. Heflin, E. Sirin, T. Tudorache, J. Euzenat, M. Hauswirth, J.X. Parreira, J. Hendler, G. Schreiber, A. Bernstein and E. Blomqvist, eds, Lecture Notes in Computer Science, Vol. 7649, Springer, 2012, pp. 427–443.
- [19] B. Pereira Nunes, A.A.M. Caraballo, M.A. Casanova, K.K. Breitman and L.A.P.P. Leme, Complex matching of RDF datatype properties, in: *Proceedings of the 6th International Workshop on Ontology Matching*, Bonn, Germany, October 24, 2011, P. Shvaiko, J. Euzenat, T. Heath, C. Quix, M. Mao and I.F. Cruz, eds, CEUR Workshop Proceedings, Vol. 814, CEUR-WS.org, 2011.
- [20] H. Qin, D. Dou and P. LePendu, Discovering executable semantic mappings between ontologies, in: *On the Move to Meaningful Internet Systems 2007: CoopIS, DOA, ODBASE, GADA, and IS, OTM Confederated International Conferences CoopIS, DOA, ODBASE, GADA, and IS 2007, Vilamoura, Portugal, November 25–30, 2007, Proceedings, Part I*, R. Meersman and Z. Tari, eds, Lecture Notes in Computer Science, Vol. 4803, Springer, 2007, pp. 832–849.
- [21] Y. Ren, A. Parvizi, C. Mellish, J.Z. Pan, K. van Deemter and R. Stevens, Towards competency question-driven ontology authoring, in: *The Semantic Web: Trends and Challenges*, Vol. 8465, Springer, 2014, pp. 752–767. doi:[10.1007/978-3-319-07443-6_50](https://doi.org/10.1007/978-3-319-07443-6_50).
- [22] D. Ritze, C. Meilicke, O.Š. Zamazal and H. Stuckenschmidt, A pattern-based ontology matching approach for detecting complex correspondences, in: *Proceedings of the 4th International Workshop on Ontology Matching (OM-2009) Collocated with the 8th International Semantic Web Conference (ISWC-2009)*, Chantilly, USA, October 25, 2009, P. Shvaiko, J. Euzenat, F. Giunchiglia, H. Stuckenschmidt, N.F. Noy and A. Rosenthal, eds, CEUR Workshop Proceedings, Vol. 551, CEUR-WS.org, 2009.

- [23] D. Ritze, J. Völker, C. Meilicke and O.Š. Zamazal, Linguistic analysis for complex ontology matching, in: *Proceedings of the 5th International Workshop on Ontology Matching (OM-2010)*, Shanghai, China, November 7, 2010, P. Shvaiko, J. Euzenat, F. Giunchiglia, H. Stuckenschmidt, M. Mao and I.F. Cruz, eds, CEUR Workshop Proceedings, Vol. 689, CEUR-WS.org, 2010.
- [24] J. Rouces, G. de Melo and K. Hose, Complex schema mapping and linking data: Beyond binary predicates, in: *Proceedings of the Workshop on Linked Data on the Web, LDOW 2016, Co-Located with 25th International World Wide Web Conference (WWW 2016)*, S. Auer, T. Berners-Lee, C. Bizer and T. Heath, eds, CEUR Workshop Proceedings, Vol. 1593, CEUR-WS.org, 2016.
- [25] C. Roussey, J.-P. Chanet, V. Cellier and F. Amarger, Agronomic taxon, in: *Proceedings of the 2nd International Workshop on Open Data, WOD 2013*, Paris, France, June 3, 2013, V. Christophides and D. Vodislav, eds, ACM, 2013, pp. 5:1–5:4.
- [26] E. Thiéblin, O. Haemmerlé, N. Hernandez and C. Trojahn, Task-oriented complex ontology alignment: Two alignment evaluation sets, in: *The Semantic Web – 15th International Conference, ESWC 2018, Heraklion, Crete, Greece, June 3–7, 2018, Proceedings*, A. Gangemi, R. Navigli, M.-E. Vidal, P. Hitzler, R. Troncy, L. Hollink, A. Tordai and M. Alam, eds, Lecture Notes in Computer Science, Vol. 10843, Springer, 2018, pp. 655–670.
- [27] É. Thiéblin, O. Haemmerlé, N. Hernandez and C. Trojahn, Survey on complex ontology matching, *Semantic Web* **11**(4) (2020), 689–727. doi:[10.3233/SW-190366](https://doi.org/10.3233/SW-190366).
- [28] E. Thiéblin, O. Haemmerlé and C. Trojahn, Complex matching based on competency questions for alignment: A first sketch, in: *Proceedings of the 13th International Workshop on Ontology Matching Co-Located with the 17th International Semantic Web Conference, OM@ISWC 2018*, Monterey, CA, USA, October 8, 2018, CEUR Workshop Proceedings, Vol. 2288, CEUR-WS.org, 2018, pp. 66–70.
- [29] É. Thiéblin, O. Haemmerlé and C. Trojahn, Generating expressive correspondences: An approach based on user knowledge needs and a-box relation discovery, in: *The Semantic Web – ISWC 2020 – 19th International Semantic Web Conference, Semantic Web – ISWC 2020 – 19th International Semantic Web Conference, Athens, Greece, November 2–6, 2020, Proceedings, Part I*, J.Z. Pan, V.A.M. Tamma, C. d'Amato, K. Janowicz, B. Fu, A. Polleres, O. Seneviratne and L. Kagal, eds, Lecture Notes in Computer Science, Vol. 12506, Springer, 2020, pp. 565–583.
- [30] É. Thiéblin, O. Haemmerlé and C. Trojahn, Automatic evaluation of complex alignments: An instance-based approach, *Semantic Web* **12**(5) (2021), 767–787.
- [31] E. Thiéblin, N. Hernandez, C. Roussey and C. Trojahn, Cross-querying LOD data sets using complex alignments: An experiment using agronomictaxon, agrovoc, dbpedia and TAXREF-LD, *International Journal of Metadata, Semantics and Ontologies* **13**(2) (2018), 104–119. doi:[10.1504/IJMSO.2018.098387](https://doi.org/10.1504/IJMSO.2018.098387).
- [32] É. Thiéblin and C. Trojahn, Conference v3.0: A populated version of the conference dataset, in: *ISWC Poster Track*, 2019.
- [33] C. Unger, C. Forascu, V. López, A.-C. Ngonga Ngomo, E. Cabrio, P. Cimiano and S. Walter, Question answering over linked data (QALD-4), in: *Linda Cappellato, Nicola Ferro, Martin Halvey, and Wessel Kraaij, Editors, Working Notes for CLEF 2014 Conference*, Sheffield, UK, September 15–18, 2014, CEUR Workshop Proceedings, Vol. 1180, CEUR-WS.org, 2014, pp. 1172–1180.
- [34] P.R.S. Visser, D.M. Jones, T.J.M. Bench-Capon and M.J.R. Shave, An analysis of ontology mismatches: Heterogeneity versus interoperability, in: *AAAI 1997 Spring Symposium on Ontological Engineering*, Stanford CA, USA, 1997, pp. 164–172.
- [35] B. Walshe, R. Brennan and D. O'Sullivan, Bayes-recce: A Bayesian model for detecting restriction class correspondences in linked open data knowledge bases, *Int. J. Semant. Web Inf. Syst.* **12**(2) (2016), 25–52. doi:[10.4018/IJSWIS.2016040102](https://doi.org/10.4018/IJSWIS.2016040102).
- [36] B. Wu and C.A. Knoblock, An iterative approach to synthesize data transformation programs, in: *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015*, Buenos Aires, Argentina, July 25–31, 2015, Q. Yang and M.J. Wooldridge, eds, AAAI Press, 2015, pp. 1726–1732.
- [37] O.Š. Zamazal and V. Svátek, The ten-year OntoFarm and its fertilization within the onto-sphere, *Web Semantics: Science, Services and Agents on the World Wide Web* **43** (2017), 46–53. doi:[10.1016/j.websem.2017.01.001](https://doi.org/10.1016/j.websem.2017.01.001).
- [38] M. Zhao, S. Zhang, W. Li and G. Chen, Matching biomedical ontologies based on formal concept analysis, *J. Biomed. Semant.* **9**(1) (2018), 11:1–11:27. doi:[10.1186/s13326-018-0178-9](https://doi.org/10.1186/s13326-018-0178-9).
- [39] W. Zheng, L. Zou, W. Peng, X. Yan, S. Song and D. Zhao, Semantic sparql similarity search over rdf knowledge graphs, *Proceedings of the VLDB Endowment* **9**(11) (2016), 840–851. doi:[10.14778/2983200.2983201](https://doi.org/10.14778/2983200.2983201).
- [40] L. Zhou, M. Cheatham and P. Hitzler, AROA results for 2019 OAEI, in: *Proceedings of the 14th International Workshop on Ontology Matching Co-Located with the 18th International Semantic Web Conference (ISWC 2019)*, Auckland, New Zealand, October 26, 2019, 2019, pp. 107–113.
- [41] L. Zhou and P. Hitzler, AROA results for OAEI 2020, in: *Proceedings of the 15th International Workshop on Ontology Matching Co-Located with the 19th International Semantic Web Conference*, 2020, pp. 161–167.