

GeoLink Dataset: A Complex Alignment Benchmark from Real-world Ontology

Lu Zhou^a, Michelle Cheatham^b, Adila Krisnadhi^c, Pascal Hitzler^a

^a*DaSe Lab, Kansas State University, Manhattan, KS 66506, USA*

^b*Wright State University, Dayton, OH 45435, USA*

^c*Faculty of Computer Science, Universitas Indonesia, Depok, Jawa Barat 16424, Indonesia*

Abstract

Ontology alignment has been studied for over a decade, and over that time many alignment systems and methods have been developed by researchers in order to find simple 1-to-1 equivalence matches between two ontologies. However, very few alignment systems focus on finding complex correspondences. One reason for this limitation may be that there are no widely accepted alignment benchmarks that contain such complex relationships. In this paper, we propose a real-world dataset from the GeoLink project as a potential complex ontology alignment benchmark. The dataset consists of two ontologies, the GeoLink Base Ontology (GBO) and the GeoLink Modular Ontology (GMO), as well as a manually created reference alignment that was developed in consultation with domain experts from different institutions. The alignment includes 1:1, 1:n, and m:n equivalence and subsumption correspondences, and is available in both Expressive and Declarative Ontology Alignment Language (EDOAL) and rule syntax. The benchmark has been expanded from its original version to contain real-world instance data from seven geoscience data providers that has been published according to both ontologies. This allows it to be used by extensional alignment systems or those that require training data. This benchmark has been incorporated into the Ontology Alignment Evaluation Initiative (OAEI) complex track to help researchers test their automated alignment systems and algorithms. This paper also analyzes the challenges inherent in effectively generating, detecting, and evaluating complex ontology alignments and provides a road map for future work on this topic.

Keywords: Complex Ontology Alignment, Real-world Ontology, Ontology Population, Benchmark

1. Introduction

Ontology alignment is an important step in enabling computers to query and reason across the many linked datasets on the semantic web. This is a difficult challenge because the ontologies underlying different linked datasets can vary in terms of subject area coverage, level of abstraction, ontology modeling philosophy, and even language. Due to the importance and difficulty of the ontology alignment problem, it has been an active area of research for over a decade [1].

Email addresses: luzhou@ksu.edu (Lu Zhou), michelle.cheatham@wright.edu (Michelle Cheatham), adila@cs.ui.ac.id (Adila Krisnadhi), hitzler@ksu.edu (Pascal Hitzler)

Preprint submitted to Data Intelligence

December 6, 2019

Ideally, alignment systems should be able to uncover any entity relationship across two ontologies that can exist within a single ontology. Such relationships have a wide range of complexity, from basic 1-to-1 equivalence, such as a Person in one ontology being equivalent to a Human in another ontology, to arbitrary m-to-n relationships, such as a Professor with a hasRank property value of “Assistant” in one ontology being a subclass of the union of the Faculty and TenureTrack classes in another. Unfortunately, the majority of the research activities in the field of ontology alignment remains focus on the simplest end of this scale – finding 1-to-1 equivalence relations between ontologies. Part of the reason for this may be that there are no widely used and accepted ontology alignment benchmarks that involve complex relations.

This paper seeks to take a step in that direction by proposing a complex alignment benchmark based on two ontologies which were developed by domain experts jointly with the reference alignment, and which in fact were developed for deployment on major ocean science data repository platforms, i.e., without the actual intention to develop an alignment benchmark. For this reason, the benchmark, including the reference alignment, can be considered to be (a) objective, in that it was created for deployment and not for benchmarking, (b) realistic, in that it captures an application use case developed for deployment, and (c) a valid ground truth alignment, in that the two ontologies and the reference alignment were developed together, by domain experts. We argue that it is therefore of rather unique nature and will inform complex ontology alignment research from a practical and applied perspective, rather than artificial laboratory-like. The benchmark, coincidentally, as this was the requirement of the use case, has a particular focus on relationships involving properties, which is particularly interesting because those have been shown to be rather difficult to handle for current alignment approaches [2]. In addition, we have analyzed and categorized the mapping rules constituting the alignment. We found several which had not been classified or discussed previously, which we will present and discuss in our analysis.

The main contributions of this paper are therefore the following:

- Presentation of two ontologies to support data representation, sharing, integration, and discovery for the geoscience research domain.
- Creation of an alignment between these two ontologies that includes 1:1, 1:n, and m:n correspondences, and given the creation history and usage of the alignment, it is fair to say that the alignment constitutes a gold-standard reference.
- Publication of the benchmark alignment in both rule syntax and EDOAL format¹ at a persistent URL² under a CC-BY license.
- Population of the Abox information supported by data providers to extend the functionality of the benchmark in instance-based applications.
- Incorporation of the benchmark into the OAEI complex track³ in order to help researchers to test and improve their complex ontology alignment systems and algorithms.
- Discussion of the challenges related to the generation, detection, and evaluation of the complex ontology alignment and the potential methods for future work in this area.

¹<http://alignapi.gforge.inria.fr/edoal.html>

²<http://doi.org/10.6084/m9.figshare.5907172>

³<http://oaei.ontologymatching.org/2018/complex/index.html>

This paper is an extended version of the one presented at the International Semantic Web Conference 2018 [3]. The final three bullet points above represent new material.

The paper is organized as follows. Section 2 discusses the few existing ontology alignment benchmarks that involve relationships other than 1-to-1 equivalence and methods to detect them. Section 3 gives further background on the GeoLink modeling process, including why two different but related ontologies were developed. Section 4 discusses the alignment between the two GeoLink ontologies, along with some descriptive statistics and an analysis of the types of mapping rules constituting the alignment, and the instance data population process. Section 5 introduces the simplified version of the benchmark used in the OAEI complex track and presents the evaluation results. Section 6 discusses the challenges that we faced in the research and provides potential approaches to solve them. Section 7 concludes with a discussion of potential future work in this area.

2. Related Work

Most work associated with evaluating the performance of ontology alignment systems has been done in conjunction with the Ontology Alignment Evaluation Initiative (OAEI)⁴. These yearly events allow developers to test their alignment systems on various tracks that evaluate performance on different facets of the problem such as instance matching, large ontology matching, and interactive matching, among others. Currently, most of these tracks involve the identification of 1-to-1 equivalence relationships, such as a Participant being equivalent to an Attendee. In 2009, the OAEI ran an “oriented” matching track that challenged systems to find subsumption relationships such as a Book is a subclass of a Publication. However, this track was abandoned after one year. Some system developers complained that the quality of the reference alignment was low [4]. This frustrated system developers and limited participation. Discussions at the last two Ontology Matching workshops⁵ made it clear that the community is interested in complex alignment, but that lack of applicable benchmarks is hindering progress. Our proposed benchmark seeks to address this concern by providing a reference alignment as a benchmark, and by addressing the quality issue of the previous benchmark by the fact that the process leading to the reference alignment guarantees its high quality.

Related work is currently being undertaken by Thieblin and her colleagues [5], who are creating a complex alignment benchmark using the Conference track ontologies within the OAEI [6]. This work is partially completed, and at the time of this writing it covers three of the seven ontologies. The reference alignment we describe herein differs from the effort by Thieblin et al. in that the GeoLink ontologies and alignment constitute real-world datasets designed and used in a practical application by geoscientists, rather than being an artificial artifact designed solely for alignment benchmarking. Furthermore, data from seven geoscience repositories have been published according to the GeoLink schema and are available online⁶. This instance data can in the future be used by alignment systems that employ extensional matching techniques [7]. In contrast to this, significant instance data is currently not readily available for most of the OAEI conference track ontologies. With the increasing requirement of more complex ontology alignment and growing interest in generating complex correspondences in real-world datasets [3, 8],

⁴<http://oaei.ontologymatching.org>

⁵<http://www.ontologymatching.org/>

⁶<http://data.geolink.org>

the first version of the complex alignment track was introduced in OAEI 2018 [9]. Our GeoLink benchmark is one of the four benchmarks that contain complex correspondences in this track. The other three complex ontology alignment benchmarks are from different domains: conference, hydrography, and plant taxonomy. In addition, different evaluation strategies were applied in evaluating the performance of complex alignment systems on the different benchmarks. More details of evaluations and results can be accessed on the OAEI 2018 website⁷.

While alignment systems capable of generating complex alignments are relatively rare, several approaches have been proposed in the literature. Ritze applied pattern-based [10] and linguistic analysis approaches [11] to detect the complex correspondences in a dataset. Jiang [12] accomplished the task of finding a complex alignment by defining knowledge rules and using a probabilistic framework to integrate a knowledge-based strategy with standard terminology-based and structure-based strategies.

Alignment systems that attempt to identify subsumption relations have sometimes used their own manually developed (and sometimes unpublished) reference alignments [13]. Other subsumption systems have evaluated the precision of their approach by manually validating relations produced by their system, while foregoing an assessment of recall [14]. Other related work has centered on developing a benchmark for compound alignments, which the authors define as mappings between class or property expressions involving more than two ontologies [15]. Their first step in this direction was to create a set of reference alignments containing relations of the form $\langle X, Y, Z, R, M \rangle$, where X, Y and Z are classes from three different ontologies and R is a relation between Y and Z that results in a class expression that is related to X by the relation M. For example, a DisabledVeteran (X) is equivalent to (M) the intersection (R) of Veteran (Y) and Disabled (Z). This benchmark is based on cross-products among the Open Biomedical Ontologies (OBO) Foundry⁸, which have been manually validated by at least two experts. The work presented herein differs from these approaches by considering a wider range of relationship types (beyond subsumption and the type of ternary relation described in [15]), as they naturally arose out of the application from which the reference alignment was taken.

3. The GeoLink Modeling Process

Benchmarks come in at least two varieties. On the one hand, there are artificial benchmarks that provide a kind of laboratory setting for evaluation. On the other hand, there are benchmarks created from data as it is used in realistic use cases or even deployed scenarios. Both of these types are important, and they cover different aspects of the spectrum, and may have different advantages. Artificial benchmarks can be made to be balanced, or to focus on certain aspects of a problem, and sometimes they can be used to test scalability issues more easily as different versions of the same benchmark set may be easily producible. Natural benchmarks, on the other hand, may expose issues arising in practice which may easily be overlooked by designers of artificial benchmarks, in particular in a young field such as complex ontology alignment. Natural benchmarks also may come with an independently verified gold standard baseline, as in our case.

The project that this benchmark arose from is called GeoLink [16] and was funded under the U.S. National Science Foundation's EarthCube initiative. This planned decade-long endeavor is a recognition that oftentimes the most innovative and useful discoveries come at the intersection

⁷<http://oaei.ontologymatching.org/2018/complex/index.html>

⁸<http://www.obofoundry.org/>

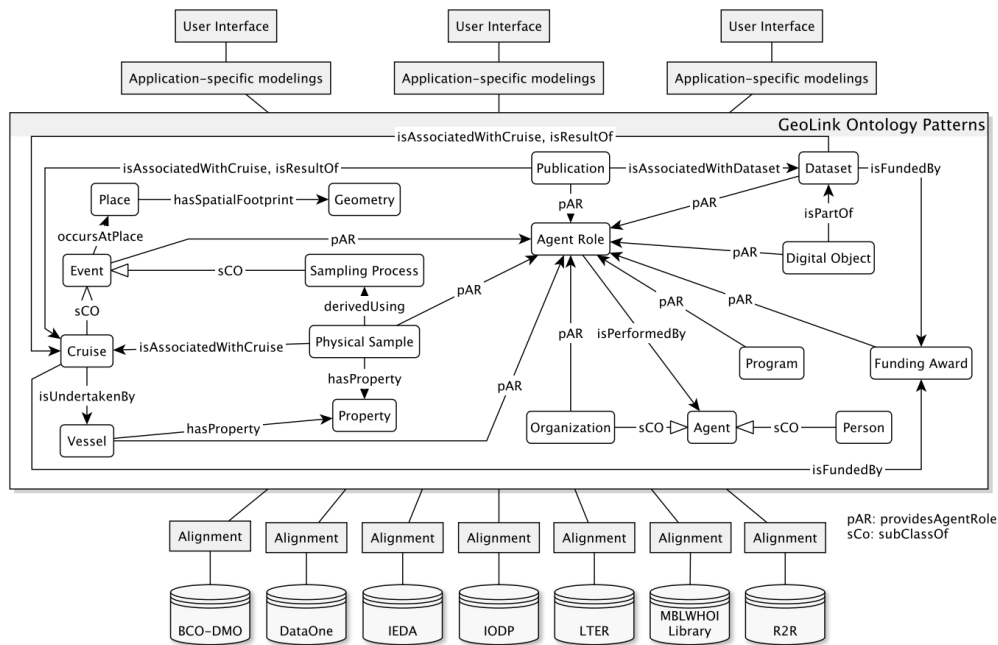


Figure 1: Intended usage of the GMO

of traditional fields of research. This is particularly true in the geosciences, which often bring together disparate groups of researchers such as geologists, meteorologists, climatologists, ecologists, archaeologists, and so on. For its part, GeoLink employs semantic web technologies to support data representation, sharing, integration, and discovery [17]. In particular, seven diverse geoscience datasets have been brought together into a single data repository.

At the beginning of the project, some providers' data resided in relational databases while others' had been published as RDF triples and exposed via a SPARQL endpoint. Because each provider had their own schema, the first step in the GeoLink project was to develop a unified schema according to which all data providers could publish their data [17]. Creating a unified schema for independently developed datasets is sometimes difficult, and the final product often ends up requiring providers to shoehorn their data into a schema that does not quite fit. GeoLink uses an approach that relies on ontology design patterns (ODP) in an attempt to avoid this issue [18]. An ODP represents a reusable solution to a recurring modeling problem and generally encodes a specific abstract notion, such as a process, event, agent, etc. These are frequently the small areas of semantic overlap that exist between datasets from different subfields of the same high-level domain. ODPs provide a structured and application-neutral representation of the key concepts within a domain. Throughout the first year of the project, geoscientists, data providers and ontologists worked together to identify and model the important concepts within the geosciences that recurred across two or more datasets. The results of this were what we call ontology modules, based on ODPs, and eventually they were stitched together to form the GeoLink Modular Ontology (GMO) [19].

As shown in Figure 1, the GMO allows data providers to publish only those aspects of their data modeled by the GMO according to that schema. Any data that the provider has that is




Figure 2: The Agent Role pattern

not covered by that schema can be published using the provider’s own schema, since no other providers have similar content. For example, in Figure 1, the provider R2R has data related mostly to the cruise and vessel modules in the lower left of the Figure 1, and so it publishes its related data using that terminology. R2R also has data not modeled by the GMO, and so it uses its own terminology when publishing that information. This freedom is intended to make the publishing process easier. However, some problems still remained.

Some of the patterns contain a rather complicated structure, mostly due to reification, which was employed to accommodate different perspectives (e.g., based on granularity) on the data. For example, many of the data providers have information about the sponsor of a project, and R2R has a native relation in their schema called `hasSponsor` with domain `Award` and range `Organization`. However, following best practices, it leads to a more versatile model if being a sponsor is recognized (and thus modeled) as a role which an agent (in this case an organization) can assume. Creating a distinct relation for each type of role on a project (sponsor, chief scientist, research assistant, etc.) is brittle, in the sense that if new roles will be added later, potentially due to the inclusion of a new dataset, then the schema will need to be edited by adding new vocabulary for new roles together with (possibly complex) role relationships. Another issue with using a relation such as `hasSponsor` is that a more fine-grained data repository may have additional temporal information related to the sponsor role, and then it is not clear how to add this temporal information to the `hasSponsor` model without punning. Essentially, `hasSponsor` should better be expressed as a ternary relation between award, organization, and the type of relation (in this case, being a sponsor) expressed using an individual which can be reused in all sponsor relationships. In terms of ODPs, this is realized by reusing the Agent Role pattern, shown in abstract form in Figure 2. This approach both allows new roles to be added easily (by subclassing `AgentRole`) and supports temporal queries if desired.

Unfortunately, while the data providers recognized the utility of this modeling approach, they found it cumbersome to map their data to it. Looking at their own schemas, they found nothing equivalent to `AgentRole`, and looking at the GMO, they found no obvious way to model the `Sponsor` field in their database. Additionally, reification led to the generation of blank nodes and the need to create and maintain many URIs. A simpler interface for the data providers was therefore requested.

To accommodate this, a second ontology, together with a manual alignment between this ontology and the GMO, was created to bridge the gap via an intermediate schema that is “flatter” than the patterns and closer to the data providers’ own schemas, but still easy to align to the GMO modules because it has been developed directly out of the GMO. This ontology is referred to as the GeoLink Base Ontology (GBO). The providers publish their data according to the GBO

and then SPARQL construct queries which encode the alignment can be used to map data to the GMO. From the very beginning, it was intended that the data integration process would be based on manual, and thus high-quality, mappings between different schemas. As a consequence, ontology alignment systems were not employed to make these mappings, not even to inform human decisions. All mappings were established as a collaborative effort between the data repository providers, the domain experts, and the ontology engineers involved in the modeling and deployment process. Because the GBO was manually engineered directly from the GMO in order to serve this particular purpose, the alignment is guaranteed to be precisely the one intended by the developers. I.e. the alignment is guaranteed to contain all of the relations necessary to solve this real-world alignment problem and no superfluous relations have been included. We argue that this characteristic makes the GeoLink ontologies a good example of a complex ontology alignment problem that can be used as a benchmark for systems that attempt to automate such alignment processes: While it is not a synthetic benchmark, it reflects complex alignment issues encountered in practice.

The example below illustrates the use of the GBO and its alignment to the GMO. In the GBO, there is a relation called `hasSponsor` with a domain that includes `Award` and range `Organization`. This mirrors many of the providers' existing schemas. Providers publish triples either directly according to the GMO schema (e.g., if they have temporal information), or according to the GBO schema.

```
x:award1 a          view:Award ;
         view:hasSponsor x:org1 .
x:org1   a          view:Organization .
```

In the latter case, the GBO-oriented triples are converted into the GMO schema using this SPARQL construct:

```
PREFIX view: <http://schema.geolink.org/dev/view#>
CONSTRUCT {
  ?x    a          :FundingAward ;
        :providesAgentRole _:bn1 .
  _:bn1 :isPerformedBy ?y ;
        a          :SponsorRole .
  ?y    a          :Organization .
} WHERE {
  ?x    a          view:Award ;
        view:hasSponsor ?y .
  ?y    a          view:Organization
}
```

Let us look at this by means of a schema diagram. In Figure 3, the three nodes and the two solid arrows indicate the graph pattern used to express the sponsoring organization role in the GMO. The dashed arrow is sometimes called a *shortcut* [20]. This shortcut (which is not part of the GMO) “flattens” this part of the GMO, and in the GBO, the `:SponsorRole` node is removed, but the shortcut is added (and `:FundingAward` and `:Organization` have been replaced by the local `view:Award` and `view:Organization`, respectively).

Note that there is no doubt here about the intended alignment between the corresponding parts of the GBO and the GMO: `view:Award` and `:FundingAward` should be mapped to each other

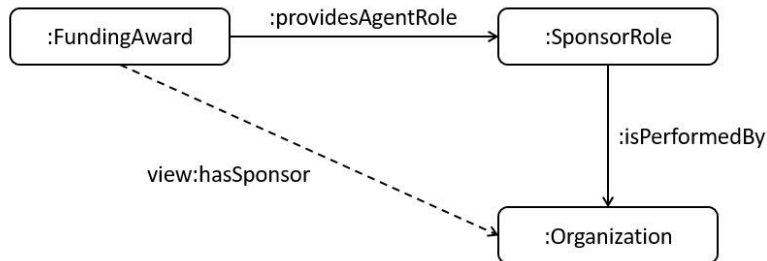


Figure 3: A schema diagram to explain an example alignment

(as equivalence), as should `view:Organization` and `:Organization`. It is also clear that the relation `view:hasSponsor` between an `view:Award` and an `view:Organization` should be aligned (as equivalence) to the concatenation of `:providesAgentRole` and `:isPerformedBy`, provided the entity shared by the two relation expressions is of type `:SponsorRole`, and the chain starts at a `:FundingAward` and ends at an `:Organization`. I.e. a complex alignment is required to express this very natural relationship between these two ontology snippets. Below we will give more examples of complex alignments arising from our setting, when we discuss the different alignment patterns we have identified. The example above is a “Typed Property Chain Equivalence” in our classification, and below we discuss this example further.

More information about the GMO and the project is available from [21] and from the project website⁹.

4. The GeoLink Complex Alignment Benchmark

4.1. Dataset

In order to prepare the GeoLink ontologies for use as a complex alignment benchmark, some changes to the namespaces were required. As we introduced in the previous section, several ODPs and modules were created to represent the frequently recurring concepts in the GeoLink datasets, and these were stitched together to form the GMO. During this process, the namespace of some entities was changed from one that reflected its originating pattern to the namespace of the GMO, which is `http://gmo#`. For example, the class `FundingAward` was originally in the `fundingaward` pattern, with the namespace `http://schema.geolink.org/1.0/pattern/fundingaward#`. After merging these modules, the namespace of the class `FundingAward` became `http://gmo#`. This has been applied to all entities except those that are imported from other ontologies, which retain their original namespace. For example, the namespace of the class `Instant`, which is imported from `http://www.w3.org/2006/time#`, remains unchanged. Additionally, the namespace of entities in the GBO has been changed from `http://schema.geolink.org/1.0/base/main#` to `http://gbo#`.

Table 1 shows the number of classes and properties in both ontologies. Both ontologies are comparable in size to ontologies currently used by the OAEL, meaning that they are within the capabilities of most current ontology alignment systems to handle.

⁹<http://www.geolink.org/>

Table 1: The number of classes, object properties, and data properties in both GeoLink ontologies

Ontology	Classes	Object Properties	Data Properties
GeoLink Base Ontology	40	149	49
GeoLink Modular Ontology	156	124	46

Table 2: The alignment pattern types found in the GeoLink complex alignment benchmark, along with the number of times each occurs and the type of relation.

Pattern	Occurrences	Category
Class Equivalence	10	1:1
Class Subsumption	2	1:1
Property Equivalence	7	1:1
Property Equivalence Inverse	5	1:1
Class Typecasting Equivalence	4	1:n
Class Typecasting Subsumption	1	1:n
Property Typecasting Subsumption	5	1:n
Property Typecasting Subsumption Inverse	5	1:n
Typed Property Chain Equivalence	26	m:n
Typed Property Chain Equivalence Inverse	17	m:n
Typed Property Chain Subsumption	17	m:n
Typed Property Chain Subsumption Inverse	12	m:n

4.2. Simple and Complex Correspondences

In order to understand the correspondences in the benchmark, we give the formal definition of simple and complex correspondences.

Simple Correspondence. Simple correspondence refers to a basic 1-to-1 simple mapping between two ontologies, in which the entities involved may be either classes or properties. This category not only includes 1-to-1 equivalence relations, but also 1-to-1 subsumption and 1-to-1 disjointness.

Complex Correspondence. Complex correspondence refers to more complex patterns, such as 1-to-n equivalence, 1-to-n subsumption, m-to-n equivalence, m-to-n subsumption, and m-to-n arbitrary relationship.

We have identified 12 different kinds of simple and complex correspondence patterns in the GeoLink complex alignment benchmark. Table 2 presents these different patterns and the corresponding number and category in the whole dataset. As the table shows, the alignment consists predominantly of complex relationships. In the following, we explain these alignment types, from simple 1-to-1 correspondence to complex m-to-n correspondence, with a formal pattern and example for each.

Class Equivalence. The first pattern is just simple 1-to-1 class equivalence. Classes C_1 and C_2 are from ontology O_1 and ontology O_2 , respectively.

Formal Pattern: $C_1(x) \leftrightarrow C_2(x)$

Example: $\text{Award}(x) \leftrightarrow \text{FundingAward}(x)$

Class Subsumption. This pattern is very similar to the first pattern. But, instead of class equivalence, this pattern describes simple 1-to-1 class subsumption.

Formal Pattern: $C_1(x) \rightarrow C_2(x)$

Example: $\text{GeoFeature}(x) \rightarrow \text{Place}(x)$

Property Equivalence. Property alignment is also an important part of ontology alignment research [20]. This pattern captures simple 1-to-1 property equivalence. Property p_1 and property p_2 are from ontology O_1 and ontology O_2 , respectively. The property can be either a data property or an object property.

Formal Pattern: $p_1(x, y) \leftrightarrow p_2(x, y)$

Example: $\text{hasAward}(x, y) \leftrightarrow \text{fundedBy}(x, y)$

Property Equivalence Inverse. This pattern is similar to the previous one, except that the domain and range values of a property are switched when it aligns to a property in another ontology.

Formal Pattern: $p_1(x, y) \leftrightarrow p_2(y, x)$

Example: $\text{isAwardOf}(x, y) \leftrightarrow \text{fundedBy}(y, x)$

Class Typecasting Equivalence. This pattern is more specific than the previous ones. The idea of typecasting, and why it is important in ontology modeling, is formally introduced and discussed in [20]. The pattern indicates that individuals of type C_1 in one ontology are cast into a subclass of C_2 in the other ontology. Note that punning is employed here – x is treated as an individual on the left-hand side of the rule and as a class on the right-hand side. For example, an instance of `PlaceType` in the GBO might be ‘ocean’. This is cast into a subclass of `Place` in the GMO. The reverse is also true: if the GMO has a subclass of `Place` called `Island`, then ‘island’ is an instance of the class `PlaceType` in the GBO.

Formal Pattern: $C_1(x) \leftrightarrow \text{rdfs:subClassOf}(x, C_2)$

Example: $\text{PlaceType}(x) \leftrightarrow \text{rdfs:subClassOf}(x, \text{Place})$

Class Typecasting Subsumption. This pattern is almost identical to the one above, except that the rule only holds in one direction. In the example, a `GeoFeatureType` (which comes from the General Bathymetric Chart of the Oceans¹⁰ vocabulary) is always a type of `Place`, but there are

¹⁰<https://www.gebco.net>

types of Places that are not GeoFeatureType.

Formal Pattern: $C_1(x) \rightarrow \text{rdfs:subClassOf}(x, C_2)$

Example: $\text{GeoFeatureType}(x) \rightarrow \text{rdfs:subClassOf}(x, \text{Place})$

Property Typecasting Subsumption. This pattern is similar in spirit to the Class Typecasting patterns mentioned above. However, in this case, a property is cast into a class assignment statement. In a sense, this alignment drops information, as y does not occur on the right-hand side.

Formal Pattern: $p_1(x, y) \rightarrow \text{rdf:type}(x, C_2)$

Example: $\text{hasPlaceType}(x, y) \rightarrow \text{rdf:type}(x, \text{Place})$

We note here that some rules that fall under this category are not exact translations of the underlying SPARQL queries, due to expressibility constraints in EDOAL (see Section 4.4 below). For instance, instead of the example above, which states that the `hasPlaceType` object property is subsumed by an `rdf:type` statement with the range value of `Place`, we would actually like to state the following, which reflects the SPARQL query:

Formal Pattern: $p_1(x, y) \leftrightarrow \text{rdf:type}(x, y) \wedge \text{rdfs:subClassOf}(y, C_2)$

Example: $\text{hasPlaceType}(x, y) \leftrightarrow \text{rdf:type}(x, y) \wedge \text{rdfs:subClassOf}(y, \text{Place})$

For instance, we would like a rule that implies that the GBO statement `hasPlaceType(Honolulu, Island)` is equivalent to stating that Honolulu is a type of Island *and* that Island is a subclass of Place in the GMO. In other words, one of the individuals occurring as a property filler on the GBO side is cast into a class on the GMO side. At the same time, the other property filler on the GBO side is asserted to be an instance of this class. However, this is not possible because the statement requires a variable (y), and that is not supported by the core EDOAL language. The EDOAL specification does mention a pattern language that might enable this type of statement, but it does not appear to be fully supported at this time.

Property Typecasting Subsumption Inverse. This pattern is the same as the one above, except that the property fillers are flipped.

Formal Pattern: $p_1(x, y) \rightarrow \text{rdf:type}(y, C_2)$

Example: $\text{isPlaceTypeOf}(x, y) \rightarrow \text{rdf:type}(y, \text{Place})$

Again, in some cases we would actually like to state the following, which cannot be fully expressed in EDOAL, to the best of our knowledge:

Formal Pattern: $p_1(x, y) \rightarrow \text{rdf:type}(y, x) \wedge \text{rdfs:subClassOf}(x, C_2)$

Example: $\text{isGeoFeatureTypeOf}(x, y) \rightarrow \text{rdf:type}(y, x) \wedge \text{rdfs:subClassOf}(x, \text{Place})$

Typed Property Chain Equivalence. A property chain is a classical complex pattern that was introduced in [10]. This pattern captures the situation related to the `hasSponsor` property discussed in detail in Section 3. The pattern applies when a property, together with a type restriction on one or both of its fillers, in one ontology has been used to “flatten” the structure of the other ontology by short-cutting a property chain in that ontology. The pattern also ensures that the types of the property fillers involved in the property chain are typed appropriately in the other ontology. The formal pattern and example are shown below. The classes D_i and property r are from ontology O_1 , and classes C_i and properties p_i are from ontology O_2 .

Formal Pattern:

$$D_1(x_1) \wedge r(x_1, x_{n+1}) \wedge D_2(x_{n+1}) \leftrightarrow C_1(x_1) \wedge p_1(x_1, x_2) \wedge C_2(x_2) \\ \wedge \cdots \wedge p_n(x_n, x_{n+1}) \wedge C_{n+1}(x_{n+1})$$

Example¹¹:

$$\text{Award}(x) \wedge \text{hasSponsor}(x, z) \leftrightarrow \text{FundingAward}(x) \wedge \text{providesAgentRole}(x, y) \\ \wedge \text{SponsorRole}(y) \wedge \text{performedBy}(y, z)$$

Note that in this and all following patterns, any of the D_i or C_i may be omitted (in which case they are essentially \top). Also, for the left-to-right direction, we assume that x_2, \dots, x_n are existentially quantified variables.

Typed Property Chain Equivalence Inverse. This pattern is the same as the one above, except that the property fillers are flipped.

Formal Pattern:

$$D_1(x_1) \wedge r(x_1, x_{n+1}) \wedge D_2(x_{n+1}) \leftrightarrow C_1(x_{n+1}) \wedge p_1(x_{n+1}, x_n) \wedge C_2(x_n) \\ \wedge \cdots \wedge p_n(x_2, x_1) \wedge C_{n+1}(x_1)$$

Example:

$$\text{Award}(z) \wedge \text{isSponsorOf}(x, z) \leftrightarrow \text{FundingAward}(z) \wedge \text{provideAgentRole}(z, y) \\ \wedge \text{SponsorRole}(y) \wedge \text{performedBy}(y, x)$$

Typed Property Chain Subsumption. This is identical to the Typed Property Chain Equivalence pattern except that the relationship only holds in one direction.

Formal Pattern:

$$D_1(x_1) \wedge r(x_1, x_{n+1}) \wedge D_2(x_{n+1}) \rightarrow C_1(x_1) \wedge p_1(x_1, x_2) \wedge C_2(x_2) \\ \wedge \cdots \wedge p_n(x_n, x_{n+1}) \wedge C_{n+1}(x_{n+1})$$

¹¹In contrast to the example discussed in Figure 3, we leave out `:Organization` and `view:Organization`, because it is possible, in principle, that a non-organization agent (e.g., an individual) may sponsor.

Example:

$$\text{Cruise}(x) \wedge \text{hasChiefScientist}(x, z) \rightarrow \text{Cruise}(x) \wedge \text{providesAgentRole}(x, y) \\ \wedge \text{AgentRole}(y) \wedge \text{performedBy}(y, z)$$

Typed Property Chain Subsumption Inverse. This pattern is the same as the one above, except that the property fillers are flipped.

Formal Pattern:

$$D_1(x_1) \wedge r(x_1, x_{n+1}) \wedge D_2(x_{n+1}) \rightarrow C_1(x_{n+1}) \wedge p_1(x_{n+1}, x_n) \wedge C_2(x_n) \\ \wedge \dots \wedge p_n(x_2, x_1) \wedge C_{n+1}(x_1)$$

Example:

$$\text{Cruise}(z) \wedge \text{isChiefScientistOf}(x, z) \rightarrow \text{Cruise}(z) \wedge \text{providesAgentRole}(z, y) \\ \wedge \text{AgentRole}(y) \wedge \text{performedBy}(y, x)$$

In [10], four alignment types were identified, some of which are subsumed by ours. We do not at all claim that our classification above is exhaustive, but we consider it a refinement of the ones listed in [10]. We conjecture that there are many additional important types of relevance to other use cases. Mapping out the space of complex alignment types is, in our understanding, helpful for further research into complex alignment algorithms.

4.3. Instance Data Population

Instance-based ontology mapping algorithms have been shown to be effective in several practical use cases [22]. The basic idea of instance-based mapping is to query the instance data of the two entities or constructs in two ontologies and calculate the overlap of the common instances, as assessed by some coreference resolution method. In order to extend the functionality of our benchmark and provide more scalability for researchers to explore algorithms that depend on the instance data, we have included the same instance data published according to both the GBO and the GMO in the GeoLink dataset.

4.3.1. Instance Data Information

The GeoLink knowledge base aims at helping users to query and reason over some of the most prominent geoscience metadata repositories in the United States. These include:

- Rolling Deck to Repository (R2R)¹²
- Biological and Chemical Oceanography Data Management Office (BCO-DMO)¹³
- International Ocean Discovery Program (IODP)¹⁴
- Marine Biological Laboratory Woods Hole Oceanographic Institution (MBLWHOI) Library¹⁵

¹²<http://www.rvdata.us/>

¹³<https://www.bco-dmo.org/>

¹⁴<https://www.iodp.org/>

¹⁵<http://www.mblwhoilibrary.org/>

- System for Earth Sample Registration (SESAR)¹⁶
- Data Observation Network for Earth (DataONE)¹⁷
- American Geophysical Union (AGU), the National Geochemical Database (NGDB)¹⁸
- United States Antarctic Program (USAP)¹⁹

Owing to these data providers, the GeoLink knowledge base contains over 48 million triples, which are formatted according to the GBO schema. As explained in Section 3, the data providers had difficulty publishing directly to the GMO schema, so the simpler (i.e. “flatter”) GBO schema was developed and they published their data according to that. In order to enable instance-based matching systems to utilize our benchmark and evaluate their performance, we have used SPARQL construct queries based on the reference alignment to expand the GeoLink ABox to include the GMO as well as the GBO tags.

4.3.2. Population Approach

As mentioned previously, the Geolink knowledge base contains over 48 million triples. In order to facilitate the convenient storage and distribution of the benchmark, we decided to pare down the size by only populating part of the instance data into the benchmark for future OAEI usage. For each reference mapping between the two ontologies, we randomly selected up to 500 instances from the GBO in the SPARQL construct queries. For usage of OAEI benchmark, currently we only published the instance data that are related to the classes and properties in the reference alignment. If there is increasing demand of other instance data which are not related to the reference alignment in the future, we are also willing to provide more instance data which can be found in the GeoLink Website²⁰.

As an example, referring to the property equivalence correspondence: $\text{hasAward}(x, y) \leftrightarrow \text{fundedBy}(x, y)$. This mapping means that the property `hasAward` in the GBO and the property `fundedBy` should be mapped to each other as equivalence. Therefore, instances that are related by the `hasAward` property in the GBO should be also related by the `fundedBy` in the GMO. The corresponding SPARQL construct query is:

```
PREFIX gbo: <http://gbo#>
PREFIX gmo: <http://gmo#>
CONSTRUCT {?x    gmo:fundedBy    ?y .}
           WHERE {?x    gbo:hasAward    ?y .}
LIMIT 500
```

This SPARQL construct query looks for the triples that have `hasAward` as the property in the dataset and creates a new graph corresponding to the `fundedBy` relation with the same x and y values. This example illustrates the usage of the SPARQL construct query.

¹⁶<http://www.geosamples.org/>

¹⁷<https://www.dataone.org/>

¹⁸<https://minerals.usgs.gov/science/natl-geochemical-db/>

¹⁹<http://www.usap-dc.org/>

²⁰<http://www.geolink.org/>

```
x:professor1 gbo:hasAward x:NSF
```

If this is a triple according to the GBO schema (ignoring the namespace “x” of the individual), the SPARQL construct query creates another one for the GMO, which is shown below:

```
x:professor1 gmo:fundedBy x:NSF
```

Besides this relatively simple mapping, our GeoLink benchmark contains more complex relations that involve reification, which lead to the generation of blank nodes. For an example we refer to the typed property chain equivalence correspondence

$$\text{Award}(x) \wedge \text{hasSponsor}(x, z) \leftrightarrow \text{FundingAward}(x) \wedge \text{providesAgentRole}(x, y) \wedge \text{SponsorRole}(y) \wedge \text{performedBy}(y, z)$$

The GBO uses a “flattened” structure for the property `hasSponsor`. Compared to the corresponding structure in the GMO, it is a shortcut of the property chain that involves the property `providesAgentRole` and `performedBy`. The SPARQL construct query for this mapping is shown below. It acquires up to 500 instances that satisfy this relation. A blank node, which is of type `SponsorRole`, is generated to maintain the property reification.

```
PREFIX gbo: <http://gbo#>
PREFIX gmo: <http://gmo#>
CONSTRUCT {
  ?x a gmo:FundingAward ;
     gmo:providesAgentRole _:bn1 .
  _:bn1 a gmo:SponsorRole ;
        gmo:performedBy ?z .
} WHERE {
  ?x a gbo:Award ;
     gbo:hasSponsor ?z .
}
LIMIT 500
```

We utilize the Jena API [23] to generate the blank node when it is needed by a SPARQL construct query. Then, we leverage the OWL API to insert the assertions into the ontologies and finally finish the population process.

4.3.3. Population Result

After finishing up the population of the instance data into the GBO and the GMO, the total number of individuals in the GBO and GMO are 10897 and 11419, respectively. In addition, the number of axioms in the two ontologies are 18336 and 56318, respectively. However, among 67 reference mappings between the GBO and the GMO, there are 19 mappings that lack any applicable instance data currently, because the data providers do not have any more instance data within their repositories. Therefore, the data providers can not publish the related instance data into the GeoLink knowledge base at this stage. We introduce and discuss some potential methods to rectify this in Section 6. In the meantime, the instance data that is currently present in the knowledge base is sufficient for the detection of most of the complex mappings within the benchmark by automated ontology alignment systems that depend on instances.

4.4. Format in Rule syntax and EDOAL format

As mentioned previously, SPARQL construct queries are used to convert data published by the data providers according to the GBO into the schema described in the GMO, because the GMO employs modeling practices that enhance extensibility and facilitate reasoning. However, most ontology alignment benchmarks are not formatted in SPARQL but rather according to the format provided by the Alignment API [13]. The standard alignment format is not expressive enough to capture complex relations. However, the Alignment API also provides a format called EDOAL that can be used to express these types of relations. This format can be read and manipulated programmatically using the Alignment API and is therefore very convenient for ontology alignment researchers. In addition, EDOAL is already accepted by the ontology alignment community. It has been used by others when proposing new alignment benchmarks [15] and [6], and we continue that approach here. Because EDOAL can be difficult for humans to parse quickly, we have also expressed the alignments in using a naive rule syntax. The rule presentation is not intended for programmatic manipulation, but rather to make it easier for humans to read and understand the alignments. Both versions of the alignment, along with the GBO and GMO ontologies, can be downloaded from <http://doi.org/10.6084/m9.figshare.5907172> under a CC-BY License. We applied the Hermit [24] reasoner to the ontologies independently to check satisfiability, since some EDOAL mappings which are part of our benchmark do not seem to be expressible in OWL DL. The GeoLink website²¹ contains detailed documentation of the dataset and provides users with more insights about the resource, such as all entities, patterns, and relationships between them in both ontologies.

5. OAEI Evaluation

5.1. Simplified Version of Benchmark

The version of the GeoLink alignment benchmark used for the first version of the complex alignment track in OAEI 2018 was slightly simplified compared to the one discussed in Section 4. Some relatively complex relations involving class typecasting have been removed due to a concern that many automated alignment systems would not consider these as potential mappings. One example is $\text{PlaceType}(x) \leftrightarrow \text{rdfs:subClassOf}(x, \text{Place})$. This mapping expresses that the set of individuals of the class `PlaceType` in the GBO is equivalent to a subclass of the class `Place` on the GMO side. This is probably a challenge for current automated alignment systems to detect because it involves entities that are not in either the source or target ontology but are rather a construct of the language (e.g. `rdfs:subClassOf`). In addition, we also removed correspondences that involve the inverse relationship, because at the time the reference alignment was created, an evaluation methodology had not yet been finalized for alignment systems on this task. In particular, our thinking was that if an alignment system managed to find a mapping for either a relation or its inverse (e.g. `isGeoFeatureTypeOf`), but not the other (`hasGeoFeatureType`), then it should not be penalized. Even though using semantic precision and recall [25] as the evaluation metric will probably resolve this issue, GeoLink ontologies fail to be expressed in OWL DL, which makes us decide to leave these mappings that involve inverse relations out of the benchmark for the OAEI 2018. After these two modifications, 67 correspondences including simple and complex relations remained in the simplified version of the benchmark. Table 3 presents the remaining patterns and their corresponding number and category in the simplified the version.

²¹<http://schema.geolink.org/>

Table 3: The alignment pattern types found in the GeoLink complex alignment benchmark in OAEI 2018, along with the number of times each occurs and the type of relation.

Pattern	Occurrences	Category
Class Equivalence	10	1:1
Class Subsumption	2	1:1
Property Equivalence	7	1:1
Property Typecasting Subsumption	5	1:n
Typed Property Chain Equivalence	26	m:n
Typed Property Chain Subsumption	17	m:n

5.2. Evaluation Results

There are three subtasks related to the evaluation of complex ontology alignment systems in OAEI 2018:

1. **Entity Identification:** For each entity in the source ontology, the alignment systems will be asked to list all of the entities that are related in some way in the target ontology. For example, referring to the example we used above,

$$\text{Award}(x) \wedge \text{hasSponsor}(x, z) \leftrightarrow \text{FundingAward}(x) \wedge \text{providesAgentRole}(x, y) \\ \wedge \text{SponsorRole}(y) \wedge \text{performedBy}(y, z)$$

the expected output from an alignment system is that the property `hasSponsor` in the GBO is related to `FundingAward`, `providesAgentRole`, `SponsorRole`, `performedBy` in the GMO and `Award` in the GBO.

2. **Relationship Identification:** Given a dictionary containing entities from the source ontology paired with all related entities, determine the expression that specifies the nature of the relation. So, in terms of the example above in this subtask, an alignment system needs to eventually determine the relationship between two sides is equivalence.
3. **Full Complex Alignment Identification;** A combination of the two former step to determine the complex alignment that exist between the source and target ontology.

All three subtasks were evaluated based on standard precision, recall and F-measure. There were 16 ontology alignment systems that participated in this year’s OAEI. Unfortunately, none of the alignment systems were capable of producing results for subtasks 2 and 3 on the GeoLink benchmark. The Table 4 shows the results of the systems that can produce results on subtask 1. There were seven such systems. The performance of these systems are shown in Table 4. Among the alignments produced by these systems, all correspondences identified between the GBO and the GMO were 1-to-1 equivalences. The precision of most of the systems was relatively high, which means that traditional ontology alignment systems can handle the simple relations in this real-world ontology alignment task fairly well. But, it is not surprising that the low recall reflects that current ontology alignment systems are not capable of identifying more complex relations, a situation that we hope will change in future years.

Table 4: The Performance of Complex Ontology Alignment Systems in OAEI on Subtask 1

System	Precision	F-measure	Recall
ALOD2Vec	0.78	0.19	0.11
DOME	0.44	0.17	0.11
LogMap	0.85	0.18	0.10
LogMapKG	0.85	0.18	0.10
LogMapLt	0.73	0.19	0.11
POMAP++	0.90	0.17	0.09
Xmap	0.39	0.15	0.09

6. Discussion

This work creates a complex ontology alignment benchmark in real-world ontologies and evaluates the performance of traditional ontology alignment systems. It can be a stepping stone towards deeper understanding and discovery in this area. It is clear that there are still some challenges in the generation, detection, and evaluation of the complex correspondences between real-world ontologies. This section outlines the challenges that we faced during our research and presents some possible methods to solve them.

- Challenge 1:** The first challenge is how to identify the complex mappings between ontologies, no matter if they are real-world or artificial ontologies, in order to create new benchmarks. So far the process of generating a consensual complex ontology alignment is time-intensive and somewhat tedious, because it requires the ontologists to design or understand the ontologies in the best practice way, and also necessitates that multiple domain experts help the ontologists with the verification of the ground truth manually. This issue could potentially be resolved in the future through creating automated recommendation systems to select and rank the possible entities and relationships from one ontology to another one, which will effectively help people in interactive reference alignment generation. One possible method based on logical RDF compression has been introduced in the paper [8]. We are currently working on this alignment system. And we hypothesize that it will be able to help the researchers to pick the possible mappings between two ontologies effectively.
- Challenge 2:** The second challenge is how to generate and populate the instance data for the entities in the source and target ontologies. In our GeoLink benchmark, even though there are over 48 million triples provided by the data providers, some entities, like the object property “hasContact” in the GMO, still lacks any corresponding individuals because none of the GeoLink data providers currently use this property. (Note: the GeoLink ontologies were also designed for possible future extension. Therefore, some entities will not be used until the data providers acquire the corresponding dataset in the future.) But the alignment exists between the two ontologies no matter whether the instance data exists or not. Therefore, we still decide to keep these alignment in our reference alignment. The lack of instance data may have a negative impact on the performance of automated complex ontology alignment systems that require instance data to support their algorithms. Similarly, significant instance data is not readily available for most of the artificial benchmarks

in OAEI. It is a challenge to supply a large amount of instance data for these benchmarks. One potential method to solve this issue is to first locate useful real-world datasets online based on the domain of the benchmark and then populate the most suitable instance data into the ontologies. For example, it might be possible to incorporate additional geoscience data repositories into the GeoLink Knowledge base to enrich our instance data. However, the amount of real-world instance data may be limited due to a lack of datasets relevant to the domain. In such cases, an artificial population process may be needed to enrich the first step, because the performance of some instance-based ontology alignment systems relies on statistical analysis and computational similarity measures that require a large number of instances [7, 22]. One possible approach might be to use the techniques described in [26].

- **Challenge 3:** The third challenge that we experienced in our research was presenting the complex alignment in EDOAL format and converting between EDOAL and OWL DL. Referring to the example of property typecasting subsumption correspondence in Section 4, we were actually trying to state the following mapping, as expressed in the rule: $\text{hasPlaceType}(x, y) \leftrightarrow \text{rdf:type}(x, y) \wedge \text{rdfs:subClassOf}(y, \text{Place})$. This is currently not supported by the core EDOAL language, because EDOAL is not good at dealing with mapping individuals. Instead of calling this as a mapping, we would probably rather say it as a mapping rule that describes the context of converting datasets from one ontology to another one. But, it seems that it falls outside of the capability of the current automated matching algorithms to detect it directly, as it defines a transformation between entities that are not listed in the ontologies.

A related problem stems from the inexpressibility of some mappings from the reference alignment in OWL DL. This came up because we originally planned to apply semantic precision and recall [25] as the evaluation metrics to compute the performance of ontology alignment systems on this benchmark, which require a reasoner to test the entailed axioms and therefore need the alignment present in OWL DL. Unfortunately, only 24 of 67 EDOAL expressions in the GeoLink alignment can be expressed in this language. In particular, many mappings that involve typed property chains are valid in EDOAL but not in OWL DL. For example, $\text{Award}(x) \wedge \text{hasEndDate}(x, z) \leftrightarrow \text{FundingAward}(x) \wedge \text{endsOnDate}(x, y) \wedge \text{time:Instant}(y) \wedge \text{time:inXSDDate}(y, z)$. This means that `hasEndDate` in the GBO is equivalent to the concatenation of `endsOnDate` and `inXSDDate` with some additional domain and range restrictions. While this type of concatenation should be unproblematic in terms of semantics, it involves concatenation of an object property with a datatype property, which is not allowed in OWL DL. We are not aware of any good solution to these two issues that we describe here.

- **Challenge 4:** The last and most difficult challenge is how to correctly and accurately evaluate the performance of complex ontology alignment systems. By far, classical precision and recall are the most widely used evaluation metrics to assess performance in the majority of existing work on ontology alignment. However, several complications arise in the use of these metrics when the alignments contain complex relations due to their all-or-nothing syntactic comparisons of individual mappings, which do not distinguish between correspondences that are formally incorrect but closely related to the correct correspon-

dences and those that are completely incorrect. For example, this is a mapping in the GeoLink reference alignment:

$$\text{Award}(x) \wedge \text{hasSponsor}(x, z) \leftrightarrow \text{FundingAward}(x) \wedge \text{providesAgentRole}(x, y) \\ \wedge \text{SponsorRole}(y) \wedge \text{performedBy}(y, z)$$

And here are two different mappings that were generated by two different hypothetical complex ontology alignment systems. The first mapping is:

$$\text{Award}(x) \wedge \text{hasSponsor}(x, z) \leftrightarrow \text{FundingAward}(x) \wedge \text{providesAgentRole}(x, y) \\ \wedge \text{AgentRole}(y) \wedge \text{performedBy}(y, z)$$

and the second one is:

$$\text{Award}(x) \wedge \text{hasSponsor}(x, z) \rightarrow \text{Program}(x) \wedge \text{providesAgentRole}(x, y) \\ \wedge \text{DataManagerRole}(y) \wedge \text{performedBy}(y, z)$$

The first mapping is formally incorrect compared to the reference alignment, but it is very closely related to it because `SponsorRole` is a subclass of `AgentRole`. Conversely, the second one is completely incorrect, as it contains incorrect domain and range restrictions of `providesAgentRole` and the relationship between the two sides indicates a subsumption rather than an equivalence relation. Some variations of the traditional precision and recall metrics have been proposed to mitigate the limitations of the basic approach, but these do not resolve all of the issues. For instance, semantic precision and recall [25] compare correspondences based on their semantic meaning rather than their syntactic representation. This is done by applying a reasoner to determine when one mapping is logically equivalent to another. Even though the semantic approaches solve an important problem for evaluating alignments with complex correspondences, they still have several limitations. One is that the reasoning takes a significant amount of time, particularly for large ontologies. Furthermore, such reasoning is not possible at all if the merged ontology is not in OWL DL, like the example introduced in Challenge 3 in our GeoLink benchmark. Therefore, a new evaluation metric will need to be designed to conquer this challenge. This new evaluation metric will need to have more detailed and accurate penalties for different kinds of closeness of entities and relationship comparisons to avoid the all-or-nothing problem in order to provide more nuanced results that can assist researchers in improving their algorithms.

7. Conclusion and Future Work

Complex ontology alignment has been discussed for a long time, but relatively little work has been done to advance the state of the art. The lack of an available complex ontology alignment benchmark may be a primary reason for the slow speed of the development. In addition, most current ontology alignment benchmarks have been created by humans for the sole purpose of evaluating ontology alignment systems, and they may not always represent real-world cases. In this paper, we have proposed a complex ontology alignment benchmark based on the real-world GeoLink knowledge base. The two ontologies and the reference alignment were designed and created by ontologists and geoscience domain experts to support data representation, sharing,

integration and discovery. We take advantage of these ontologies to create a complex ontology alignment benchmark. In our benchmark, the alignments not only cover 1:1 simple correspondences, but also contain 1:n and m:n complex relations. All correspondences required to convert between the two ontologies (a key goal of ontology alignment) are guaranteed to be present, because one ontology was consciously created from the other, with SPARQL queries to mitigate each change. In addition, the alignment has been evaluated by domain experts from different organizations to ensure the high quality. Moreover, instance data has been published according to both ontologies, which is important in order to support use of the benchmark by extensional alignment systems. Furthermore, the ontologies and alignments in both rule syntax and EDOAL format have been published in FigShare with an open access license for reusability and can be accessed in OAEI 2018 website as well²². The evaluation results of the automated ontology alignment systems that participated in OAEI 2018 are also presented in this paper.

We discuss four challenges in this paper, which we plan to explore in our future work on this topic. Besides this, with respect to the maintenance of the benchmark, our intention is to remain actively involved for years to come in the OAEI complex alignment benchmarking track, and to also develop corresponding alignment methods. We thus have an intrinsic interest in keeping the benchmark maintained and usable, which would, e.g., mean that we are prepared to transfer it to a new benchmarking framework if required in the future. At the same time, based on participants' feedback, we will modify the reference alignment if necessary to perfect the benchmark by making it more convenient to use. This may involve, for example, making the alignment available in additional formats.

Acknowledgment

We would like to thank all of the members of the GeoLink project team. In addition, we extend our gratitude to Jerome Euzenat for providing advice regarding the conversion of rules to EDOAL. Moreover, we would like to express our appreciation to Daniel Faria for helping us to prepare the dataset for evaluation on the SEALS platform²³.

References

- [1] P. Shvaiko, J. Euzenat, Ontology matching: State of the art and future challenges, *IEEE Trans. Knowl. Data Eng.* 25 (2013) 158–176.
- [2] M. Cheatham, P. Hitzler, The properties of property alignment, in: *Proceedings of the 9th International Workshop on Ontology Matching collocated with the 13th International Semantic Web Conference (ISWC 2014)*, Riva del Garda, Trentino, Italy, October 20, 2014., 2014, pp. 13–24. URL: http://ceur-ws.org/Vol-1317/om2014_Tpaper2.pdf.
- [3] L. Zhou, M. Cheatham, A. Krisnadhi, P. Hitzler, A complex alignment benchmark: Geolink dataset, in: *The Semantic Web - ISWC 2018 - 17th International Semantic Web Conference*, Monterey, CA, USA, October 8-12, 2018, *Proceedings, Part II*, 2018, pp. 273–288. URL: https://doi.org/10.1007/978-3-030-00668-6_17. doi:10.1007/978-3-030-00668-6_17.
- [4] J. David, AROMA results for OAEI 2009, 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, 2009. URL: http://ceur-ws.org/Vol-551/oei09_paper3.pdf.

²²<http://oei.ontologymatching.org/2018/complex/index.html#geolink>

²³<http://oei.ontologymatching.org/2018/seals.html> used in the OAEI

- [5] É. Thiéblin, O. Haemmerlé, N. Hernandez, 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, 2018*, pp. 655–670. URL: https://doi.org/10.1007/978-3-319-93417-4_42. doi:10.1007/978-3-319-93417-4_42.
- [6] É. Thiéblin, O. Haemmerlé, N. Hernandez, C. T. dos Santos, Towards a complex alignment evaluation dataset, in: *Proceedings of the 12th International Workshop on Ontology Matching co-located with the 16th International Semantic Web Conference (ISWC 2017), Vienna, Austria, October 21, 2017., 2017*, pp. 217–218. URL: http://ceur-ws.org/Vol-2032/om2017_poster6.pdf.
- [7] É. Thiéblin, O. Haemmerlé, C. Trojahn, CANARD complex matching system: results of the 2018 OAEI evaluation campaign, 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., 2018*, pp. 138–143. URL: http://ceur-ws.org/Vol-2288/oei18_paper4.pdf.
- [8] L. Zhou, A journey from simple to complex alignment on real-world ontologies, in: *Proceedings of the Doctoral Consortium at ISWC 2018 co-located with 17th International Semantic Web Conference (ISWC 2018), Monterey, USA, October 8th - to - 12th, 2018., 2018*, pp. 93–101. URL: <http://ceur-ws.org/Vol-2181/paper-11.pdf>.
- [9] É. Thiéblin, M. Cheatham, C. T. dos Santos, O. Zamazal, L. Zhou, The first version of the OAEI complex alignment benchmark, in: *Proceedings of the ISWC 2018 Posters & Demonstrations, Industry and Blue Sky Ideas Tracks co-located with 17th International Semantic Web Conference (ISWC 2018), Monterey, USA, October 8th - to - 12th, 2018., 2018*. URL: <http://ceur-ws.org/Vol-2180/paper-67.pdf>.
- [10] D. Ritze, C. Meilicke, O. Sváb-Zamazal, 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, 2009*. URL: http://ceur-ws.org/Vol-551/om2009_Tpaper3.pdf.
- [11] D. Ritze, J. Völker, C. Meilicke, O. Sváb-Zamazal, Linguistic analysis for complex ontology matching, in: *Proceedings of the 5th International Workshop on Ontology Matching (OM-2010), Shanghai, China, November 7, 2010, 2010*. URL: http://ceur-ws.org/Vol-689/om2010_Tpaper1.pdf.
- [12] S. Jiang, D. Lowd, S. Kafle, D. Dou, Ontology matching with knowledge rules, *T. Large-Scale Data- and Knowledge-Centered Systems* 28 (2016) 75–95.
- [13] P. Jain, P. Hitzler, A. P. Sheth, K. Verma, P. Z. Yeh, Ontology alignment for linked open 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, 2010*, pp. 402–417. URL: https://doi.org/10.1007/978-3-642-17746-0_26. doi:10.1007/978-3-642-17746-0_26.
- [14] F. M. Suchanek, S. Abiteboul, P. Senellart, PARIS: probabilistic alignment of relations, instances, and schema, *PVLDB* 5 (2011) 157–168.
- [15] C. Pesquita, M. Cheatham, D. Faria, J. Barros, E. Santos, F. M. Couto, Building reference alignments for compound matching of multiple ontologies using OBO cross-products, in: *Proceedings of the 9th International Workshop on Ontology Matching collocated with the 13th International Semantic Web Conference (ISWC 2014), Riva del Garda, Trentino, Italy, October 20, 2014., 2014*, pp. 172–173. URL: http://ceur-ws.org/Vol-1317/om2014_poster2.pdf.
- [16] J. You, Geoscientists aim to magnify specialized web searching, 2015.
- [17] A. A. Krisnadh, Y. Hu, K. Janowicz, P. Hitzler, R. A. Arko, S. Carbotte, C. Chandler, M. Cheatham, D. Fils, T. Finin, P. Ji, M. B. Jones, N. Karima, K. A. Lehnert, A. Mickle, T. Narock, M. O'Brien, L. Raymond, A. Shepherd, M. Schildhauer, P. Wiebe, The geolink framework for pattern-based linked data integration, in: *Proceedings of the ISWC 2015 Posters & Demonstrations Track co-located with the 14th International Semantic Web Conference (ISWC-2015), Bethlehem, PA, USA, October 11, 2015., 2015*. URL: http://ceur-ws.org/Vol-1486/paper_99.pdf.
- [18] P. Hitzler, A. Gangemi, K. Janowicz, A. Krisnadh, V. Presutti (Eds.), *Ontology Engineering with Ontology Design Patterns - Foundations and Applications*, volume 25 of *Studies on the Semantic Web*, IOS Press, 2016.
- [19] A. Krisnadh, Y. Hu, K. Janowicz, P. Hitzler, R. A. Arko, S. Carbotte, C. Chandler, M. Cheatham, D. Fils, T. W. Finin, P. Ji, M. B. Jones, N. Karima, K. A. Lehnert, A. Mickle, T. W. Narock, M. O'Brien, L. Raymond, A. Shepherd, M. Schildhauer, P. Wiebe, The geolink modular oceanography ontology, in: *The Semantic Web - ISWC 2015 - 14th International Semantic Web Conference, Bethlehem, PA, USA, October 11-15, 2015, Proceedings, Part II, 2015*, pp. 301–309. URL: https://doi.org/10.1007/978-3-319-25010-6_19. doi:10.1007/978-3-319-25010-6_19.
- [20] A. A. Krisnadh, P. Hitzler, K. Janowicz, On the capabilities and limitations of OWL regarding typecasting and ontology design pattern views, in: *Ontology Engineering - 12th International Experiences and Directions Workshop on OWL, OWLED 2015, co-located with ISWC 2015, Bethlehem, PA, USA, October 9-10, 2015, Revised Selected Papers, 2015*, pp. 105–116. URL: https://doi.org/10.1007/978-3-319-33245-1_11.

doi:10.1007/978-3-319-33245-1_11.

- [21] A. Krisnadhi, *Ontology Pattern-Based Data Integration*, Ph.D. thesis, Wright State University, 2015.
- [22] A. Isaac, L. van der Meij, S. Schlobach, S. Wang. An empirical study of instance-based ontology matching, in: *The Semantic Web, 6th International Semantic Web Conference, 2nd Asian Semantic Web Conference, ISWC 2007 + ASWC 2007*, Busan, Korea, November 11-15, 2007., 2007, pp. 253–266. URL: https://doi.org/10.1007/978-3-540-76298-0_19. doi:10.1007/978-3-540-76298-0_19.
- [23] J. J. Carroll, I. Dickinson, C. Dollin, D. Reynolds, A. Seaborne, K. Wilkinson, Jena: implementing the semantic web recommendations, in: *Proceedings of the 13th international conference on World Wide Web - Alternate Track Papers & Posters, WWW 2004*, New York, NY, USA, May 17-20, 2004, 2004, pp. 74–83. URL: <https://doi.org/10.1145/1013367.1013381>. doi:10.1145/1013367.1013381.
- [24] B. Glimm, I. Horrocks, B. Motik, G. Stoilos, Z. Wang, Hermit: An OWL 2 reasoner, *J. Autom. Reasoning* 53 (2014) 245–269.
- [25] J. Euzenat, Semantic precision and recall for ontology alignment evaluation, in: *IJCAI 2007, Proceedings of the 20th International Joint Conference on Artificial Intelligence*, Hyderabad, India, January 6-12, 2007, 2007, pp. 348–353. URL: <http://ijcai.org/Proceedings/07/Papers/054.pdf>.
- [26] B. A. C. Schopman, S. Wang, A. Isaac, S. Schlobach, Instance-based ontology matching by instance enrichment, *J. Data Semantics* 1 (2012) 219–236.