The Enslaved Dataset: A Real-world Complex Ontology Alignment Benchmark using Wikibase

Lu Zhou
Cogan Shimizu
Pascal Hitzler
luzhou@ksu.edu
coganmshimizu@ksu.edu
hitzler@ksu.edu
Kansas State University
Manhattan, Kansas, USA

Alicia M. Sheill
Seila Gonzalez Estrecha
Catherine Foley
Duncan Tarr
Dean Rehberger
asheill@msu.edu
seilage@msu.edu
foleyc@msu.edu
tarrdunc@msu.edu
rehberge@msu.edu
Michigan State University
East Lansing, Michigan, USA

ABSTRACT

Ontology alignment has taken a critical place for helping heterogeneous resources to interoperate. It has been studied for over a decade, and over that time many alignment systems and methods have been developed by researchers to find simple 1:1 equivalence matches between two ontologies. However, very few alignment systems focus on finding complex correspondences. Even if the complex alignment systems are developed, the performance of finding complex relations still has a lot of room for improvement. One reason for this limitation may be that there are still few applicable alignment benchmarks that contain such complex relationships that can raise researchers' interests. In this paper, we propose a real-world dataset from the Enslaved project as a potential complex alignment benchmark. The benchmark consists of two resources, the Enslaved Ontology along with a Wikibase repository holding a large number of instance data from the Enslaved project, as well as a manually created reference alignment between them. The alignment was developed in consultation with domain experts in the digital humanities. The alignment not only includes simple 1:1 equivalence correspondences, but also more complex m:n equivalence and subsumption correspondences and are provided in both Expressive and Declarative Ontology Alignment Language (EDOAL) format and rule syntax. The Enslaved benchmark has been incorporated into the Ontology Alignment Evaluation Initiative (OAEI) 2020 and is completely free for public use to assist the researchers in developing and evaluating their complex alignment algorithms.

CCS CONCEPTS

• **Information systems** → **Information integration**; **Resource Description Framework (RDF)**; **Web Ontology Language (OWL)**; **Ontologies**; • **Computing methodologies** → **Ontology engineering**.

KEYWORDS

Ontology Alignment; Knowledge Graph; Wikibase; Benchmark

ACM Reference Format:

Lu Zhou, Cogan Shimizu, Pascal Hitzler, Alicia M. Sheill, Seila Gonzalez Estrecha, Catherine Foley, Duncan Tarr, and Dean Rehberger. 2020. The Enslaved Dataset: A Real-world Complex Ontology Alignment Benchmark using Wikibase. In *Proceedings of the 29th ACM International Conference on Information and Knowledge Management (CIKM '20), October 19–23, 2020, Virtual Event, Ireland*. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3340531.3412768>

1 INTRODUCTION

Ontology alignment is an important step in enabling computers to query and reason across the immense amounts of linked data on the semantic web. It has been considered a “silver bullet” for the semantic heterogeneity problem faced by computer systems. Ontology alignment is a difficult challenge as the ontologies, which are used as knowledge graph schemas, that underlie different linked data can vary significantly in terms of subject area coverage, level of abstraction, ontology modeling philosophy, and language. Due to the importance and difficulty of the ontology alignment problem, it has been an active area of research for over a decade [21].

Ideally, alignment systems should be able to uncover any entity relationships across two ontologies that can exist within a single ontology. Such relationships have a wide range of complexity, from basic 1:1 (1-to-1) equivalence, such as a Person in one ontology being equivalent to a Human in another ontology, to arbitrary m:n (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,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CIKM '20, October 19–23, 2020, Virtual Event, Ireland

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-6859-9/20/10...\$15.00

<https://doi.org/10.1145/3340531.3412768>

the majority of the research activities in the field of ontology alignment remain focused on the simplest end of this scale – finding 1:1 equivalence relations between ontologies. Part of the reason for this may be that there are still few widely used and accepted ontology alignment benchmarks that involve complex relations. Even though some benchmarks containing complex relations were proposed in Ontology Alignment Evaluation Initiative (OAEI) in 2018 [22], the performance of the alignment systems is still relatively poor when detecting these complex correspondences between two ontologies [1].

Wikibase is the powerful knowledge base software that drives Wikidata [24]. Wikidata is an immense, crowdsourced knowledge base with persistent data that is available for public use and consumption. It would be very difficult to have an ontology of everything, but Wikidata is probably close enough for this purpose. It contains millions of pieces of knowledge from many different domains in the world. In addition, Wikidata is crowdsourced and can act as a “common resource”.¹ People can export data to Wikidata so that it is publicly persistent in an open and transparent manner. Wikidata is an instance of Wikibase. Any organization can adapt it to their own needs, including setting up their own Wikibase repositories to host their data under different licenses, so that the other instances of Wikibase can be linked with the data on Wikidata. Therefore, it is crucial to be able to find alignments between domain or proprietary ontologies and this common resource. It is also a fact that some organizations have their own internal and proprietary knowledge graphs. They can apply their alignments to this public resource as an important tool to augment or induce new information into their own knowledge graph.

This paper seeks to take a step in that direction by proposing a complex alignment benchmark based on two knowledge graphs: the Enslaved knowledge graph, that was developed by ontology engineers and domain experts together for the Enslaved project, and the Wikibase repository storing historical slave trade data collected from different provenances. The Enslaved benchmark, including the reference alignment, can be considered to be objective as it was created for deployment and not for benchmarking. It is realistic, since it captures an application use case developed for the historical slave trade, and it is a valid ground truth alignment, as the reference alignment was developed together by historian domain experts and ontology engineers. Therefore, it is rather unique in nature and will inform complex ontology alignment research from a practical and applied perspective, rather than an artificial one. The main contributions of this paper are therefore the following:

- Introduction of two knowledge graphs to support data representation, sharing, integration, and discovery for the Enslaved project;
- Creation of alignment between these two knowledge graphs that include 1:1 and m:n correspondences. Given the creation steps and usage of the alignment, it is fair to say that the alignment will constitute a gold-standard reference;
- Publication of the benchmark alignment in both rule syntax and EDOAL format² at a persistent URL³ under a CC-BY 4.0

license, and it is also incorporated into the complex ontology alignment track in OAEI 2020;⁴

- Evaluation of the quality and validity of this benchmark by using a complex alignment system from OAEI and a discussion of the results in detail.

The rest of this paper is organized as follows. Section 2 discusses the few existing ontology alignment benchmarks that involve relationships other than 1:1 equivalence. Section 3 gives further background on the Enslaved project, including its Wikibase repository, knowledge graph schema, and property reification. Section 4 discusses the alignment between two resources, along with some descriptive statistics, an analysis of the types of correspondences constituting the alignment, and the performance of a complex ontology alignment system tested on the Enslaved benchmark to evaluate the quality of the benchmark. Section 5 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).⁵ 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, knowledge graph matching, and interactive matching, among others. Currently, most of these tracks involve the identification of 1:1 equivalence relationships, such as Person being equivalent to Human. A discussion at the last two Ontology Matching workshops⁶ made it clear that the community is interested in complex ontology alignment, but that lack of applicable benchmarks is hindering progress. In OAEI 2018, the complex ontology alignment track was proposed and organized for the first time [22]. The first version of the complex track is comprised of four benchmarks containing complex relationships from the conference, hydrography, ocean science, and plant taxonomy domains respectively. In OAEI 2019, in order to extend the functionality of the benchmarks and provide more scalability for researchers to explore algorithms that depend on the instance data, Thieblin et al. populated the Conference benchmark with some instances collected from the Extended Semantic Web Conference (ESWC), along with some synthetic data.⁷ and Zhou et al. also populated a large number of real-world instances that are currently used in the GeoLink Project⁸ as part of the GeoLink benchmark [28]. In addition, different evaluation strategies were applied in evaluating the performance of complex alignment systems on different benchmarks. More details of evaluations and results can be accessed on the OAEI website.^{9,10}

Wikidata is a free and open knowledge base that covers many interesting topics, with similar coverage to Wikipedia. There are several ways to access Wikidata; there are built-in tools, external

¹<https://en.wikipedia.org/wiki/Commons>

²<http://alignapi.gforge.inria.fr/edoal.html>

³<https://doi.org/10.6084/m9.figshare.12400976>

⁴<http://oaei.ontologymatching.org/2020/complex/index.html#popenslaved>

⁵<http://oaei.ontologymatching.org>

⁶<http://www.ontologymatching.org/>

⁷https://framagit.org/IRIT_UT2/conference-dataset-population

⁸<https://www.geolink.org/>

⁹<http://oaei.ontologymatching.org/2018/complex/index.html>

¹⁰<http://oaei.ontologymatching.org/2019/complex/index.html>

Figure 1: An Excerpt of the Wikibase Repository for the Enslaved Project

tools, or programming interfaces, such as Wikidata Query¹¹ and Reasonator¹² for searching and examining Wikidata items. Therefore, it has also been considered as a useful external knowledge base for many alignment tasks, particularly for instance matching or entity resolution. For example, Geiß et al. utilized the information on locations and places extracted from Wikidata as ground truth for their entity resolution task [8]. However, the Wikidata knowledge graph hasn't been widely used for ontology alignment tasks due to its young age relative to Wikipedia and DBpedia [16]. The knowledge graphs track in OAEI [12] executes the DBpedia Extraction Framework on several different Wikis from Fandom,¹³ which is one of the most popular Wiki Farms and generates several knowledge graphs for both the instance matching (i.e., entities derived from pages about the same real-world entity in different Wikis) and schema matching (i.e., classes and properties derived from different constructs in different Wikis).

The Enslaved benchmark we describe herein differs from current benchmarks in OAEI in some aspects. First, the Enslaved benchmark is a good reflection of real-life data since the Enslaved project comprises over 33 million triples currently from real-world datasets shared by different researchers and contributors from different institutions, while the conference benchmark only consists of synthetic instance data. Second, the Enslaved benchmark utilizes a Wikibase repository as a central storage repository to represent the knowledge in the historical slave trade domain. Anyone with any level of expertise can access the content in the same way they access Wikidata and make use of the knowledge graph. So, it greatly improves the availability of the benchmark. Furthermore, it is useful and important to align the domain ontology to the Wikidata schema, in order to further enrich other external knowledge graphs. Third,

the two knowledge graph schemas in the Enslaved benchmark are completely designed and modeled independently, while the two ontologies in the GeoLink benchmark were developed together for the same project, which may not be common occurrences. Therefore, it further improves the quality and generality of the benchmark and can be considered as a potential good benchmark for complex alignment research.

3 BACKGROUND

3.1 The Enslaved Project and Ontology

The Enslaved Ontology was developed as part of an ongoing project entitled Enslaved: Peoples of the Historical Slave Trade¹⁴ and funded by The Andrew W. Mellon Foundation where the focus is on tracking the movements and details of peoples in the historical slave trade. It further acts as an interchange format between a number of heterogeneous data formats among other projects in the digital humanities, because it has built a proof of concept for a slave data hub. At the heart of the project, the Enslaved Hub allows students, researchers, and the general public to search over numerous databases to reconstruct the lives of individuals who were part of the historical slave trade. The Enslaved project leverages Linked Open Data (LOD) techniques, including the use of Wikibase and a graph database, to create an innovative and compelling centralized Hub for engaging with historical slave trade data from a variety of sources. LOD is a method of exposing, sharing, and connecting data on the semantic web. Data from the different sources is standardized, aggregated, and formatted in such a way that it is machine-readable and is predicated on the relationship between data as developed with the Enslaved Ontology. The central notion of the Enslaved Ontology models *records* of historical agents [19].¹⁵

¹¹<https://query.wikidata.org/>

¹²<https://tools.wmflabs.org/reasonator/>

¹³<https://www.fandom.com/>

¹⁴<https://enslaved.org/>

¹⁵Comprehensive documentation can be found in [20]

The key observation is that the ontology is necessarily a secondary (or further) source and thus cannot purport to state ontological truth. As such, it models, instead, the observations that historians or record keepers have made over time.

The development of the ontology was a collaborative effort and was carried out using a modular ontology modeling approach based on ontology design patterns [4, 7, 13]. Such a methodology is designed to ensure high quality and reusability of the ontology, as well as cater to future expansions, both in terms of scope and in terms of granularity. This allows the Enslaved Ontology to adapt as needs evolve and the number of researchers and contributors increases. The modular ontology modeling approach and its rationale have been described in [15], and it is closely related to the eXtreme Design approach [3]. The modeling team included domain experts, data experts, software developers, and ontology engineers.

The primary purpose of the formal axiomatization is to disambiguate the model, i.e., we were striving for as complete an axiomatization as possible while avoiding ontological over-commitments. Each axiom was discussed in detail between the ontology engineers and the historians on the team. The axiomatization is expressed using the OWL 2 DL profile. The primary goal was not to do formal reasoning over the ontology, but it was authored in such a way as to not rule out such goals in the future (e.g. the use of reasoning for consistency checking) [14].

3.2 The Wikibase Repository and Wikidata Knowledge Graph Schema

Wikibase is a powerful, flexible, and customizable knowledge base software. Its primary components are the Wikibase Repository, an extension for storing and managing data. Wikibase makes collaboration easy for humans and machines alike, and its data model prioritizes language independence and knowledge diversity.

Wikidata is the largest website that is powered by Wikibase. It is an open knowledge base that was launched in 2012. Similar to all the other projects of Wikimedia, anyone can freely edit it.¹⁶ The main goal of Wikidata is to act as central storage for the structured data to provide support for Wikipedia. However, it has grown out of that, since it provides structured linked data about lots of interesting topics in the world, and it is licensed under Creative Commons CC-Zero, which is very close to the public domain and anyone can use it for any purpose. Wikibase is the software that Wikidata has utilized for such success. The Enslaved project uses its own installation of the Wikibase platform to a similar purpose, creating the Enslaved Hub, as mentioned in the previous section. For brevity, we will use the acronym, EKG for the Enslaved Wikibase Knowledge Graph Schema, and EWI for the Enslaved Project’s Wikibase installation.

Figure 1 shows the Enslaved Wikibase page for a Person named Maria that appears in the Enslaved benchmark. In the center, we can see the language and label of the entire descriptions of what it means. The important thing is that Maria could be ambiguous because multiple person records may have the same name or there could be other items which are called Maria. To make this item uniquely identified, an item identifier is used as a Q followed by a number, such as Q1534 in this case.

¹⁶https://www.wikidata.org/wiki/Wikidata:Main_Page

Figure 2: The Example of Enslaved Knowledge Graph for the Enslaved Project

The main part of any Enslaved Wikibase page is the statements section that can be seen in the center of Figure 1. For example, there is an object property in the EKG called *instance of* with the value of Class *Person*. It can be interpreted that an entity *Maria* connects to an entity *Person* by an edge. The edge is labeled as *instance of*. Properties in Wikibase have a P prefix followed by a number, such as instance of (P1), hasName (P20), and hasSex (P31). The references are used to point to specific sources that back up the data provided in a statement. For instance, the statement, “Maria is an instance of Person” which “is directly based on” Maranhão Plantation Inventories [11]; the latter statement allows an interested user to track the provenance of the information contained in the previous statement. This single example is just a small excerpt of EKG.

We mapped the OWL classes and properties in the Enslaved ontology with the items and properties in the Enslaved Wikibase knowledge graph. To the authors’ knowledge, this is the first time an OWL ontology had been mapped onto a Wikibase installation. The Enslaved project team found that Wikibase was especially useful for organizing the historical slave trade data, as it had built-in tools that, for example, add qualifiers and references to every statement about the Enslaved data. Such features helped to connect time-bound statements to specific events and connect provenance information to each data point. The EWI stores the instance data, including all of the controlled vocabularies and multiple examples of people, events, and places. Through this process the exact manner in which people are connected to events, events are connected to places at specific periods of time, and how every piece of data is attached to provenance information can be examined. The work mapping the raw data onto the Enslaved ontology via Wikibase has proven that the fields developed for the Enslaved Hub can in fact represent diverse datasets.

3.3 Property Reification

Property reification is a classic strategy for adding context to a property. We mention this here, in particular, as it is frequently utilized in the Enslaved Ontology and Wikibase repository. Two such examples can be found in Figures 2 and 3.

Figure 3: The Example of Wikibase Repository for the Enslaved Project

In the abstract, property reification is an addition of a node between two other nodes, in a knowledge graph. For example, in Figure 2, the property of `hasName` is reified. That is, there is not a direct link from the `enslaved:Person` to the `xsd:string`. This is done in order to add further context. Within the confines and purposes of the Enslaved Ontology, a name is not a constant thing. Enslaved people frequently went by aliases, different family names, incorrect spellings in documents, and so on. By reifying the property of having a name to account for both historical records and the temporality of a name, we have a more accurate, historical model of that particular enslaved person.

In any Wikibase repository, reification is also very common. In order to track provenance in the background, many items are reified.

4 THE ENSLAVED COMPLEX ONTOLOGY ALIGNMENT BENCHMARK

In this section, we present the details of the Enslaved benchmark, including the process of dataset preparation, some descriptive statistics of the benchmark, the types of alignment rules and correspondence patterns with examples, the expression format of reference alignment, and the evaluation of the quality of the benchmark.

4.1 Dataset

The Enslaved benchmark consists of two knowledge graphs. The first knowledge graph utilizes the Enslaved ontology generated by the ontology engineers in consultation with domain experts from different institutions [19]. The Enslaved ontology serves as the underlying schema for the Enslaved knowledge graph in order to enable the historical slave trade data sharing and integration for different communities. The other one is the Enslaved Wikibase repository that is currently being built and used on Wikibase, which employs a completely different schema.

The Enslaved knowledge base currently contains over 33 million triples which are formatted according to the Enslaved Wikibase repository schema, and the number of triples is still growing as additional researchers and contributors start sharing their data for the integration. In order to utilize these two Enslaved knowledge graphs to establish a complex ontology alignment benchmark and facilitate the convenient storage and distribution for OAEI, we decided to pare down the size by only populating part of the instance

data into the benchmark. For each reference alignment between two knowledge graph schemas, we randomly selected up to 500 instances and populated them into both knowledge graphs with their underlying schema. If there is increasing demand of more instance data which are not related to reference alignment in the future, we can provide more data which can be found in the Enslaved Wiki pages.¹⁷

After finishing up the population of the instance data into the Enslaved knowledge graph and the Enslaved Wikibase knowledge graph, Table 1 shows the number of classes, properties, axioms, and instances in both resources respectively. Both of the knowledge graphs are comparable in size to the benchmarks currently used by the OAEI, which means that they are within the capabilities that most current ontology alignment systems to handle.

Table 1: The number of classes, properties, axioms and instances in two knowledge graph schema

Ontology	Classes	Properties	Axioms	Instances
Enslaved Knowledge Graph	43	75	67,613	13,763
Enslaved Wikibase Knowledge Graph	20	50	83,512	18,415

4.2 Simple and Complex Correspondences

There are two different types of correspondences, which are simple correspondence and complex correspondence [27]. Simple correspondence refers to basic 1:1 simple alignment between two ontologies, such as 1:1 class equivalence, property equivalence, and 1:1 class subsumption, property subsumption. Complex correspondence usually consists of more complex patterns compared to simple correspondence. It may comprise more than one class or property in both ontologies, such as 1:n equivalence, m:n equivalence, and m:n arbitrary relationship. With respect to the correspondence patterns, Zhou et al. list roughly 12 different types of simple and complex correspondence patterns [27]. In the Enslaved benchmark, there are three different types that emerge most frequently in ontology matching tasks, which are listed in Table 2. In the following, we explain the alignment types with a formal pattern and example for each. Some namespaces that are frequently used in the following examples are listed below.

```
@prefix ed:<https://lod.enslaved.org/entity/> .
@prefix ep:<https://lod.enslaved.org/prop/> .
@prefix eps:<https://lod.enslaved.org/prop/statement/> .
@prefix epq:<https://lod.enslaved.org/prop/qualifier/> .
@prefix wikibase:<http://wikiba.se/ontology#> .
@prefix enslaved:<https://enslaved.org/ontology/> .
```

- **Class Equivalence.** is simple 1:1 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: `enslaved:Person(x) ↔ ed:Q410(Person)(x)`

Note that `ed:Q410` has the label of `Person` in Wikidata.¹⁸

In order to better understand the example, we use both the

¹⁷https://lod.enslaved.org/wiki/Meta:Main_Page

¹⁸<https://lod.enslaved.org/wiki/Q410>

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

Pattern	Occurrences	Category
Class Equivalence	15	1:1
Typed Property Chain Equivalence	67	m:n
Typed Property Chain Subsumption	16	m:n

unique identifier, which is presented by a Q prefix followed by a number and the label in the alignment rules. However, in the real alignment rules, only the identifier is kept. This also applies to all the following examples.

- **Typed Property Chain Equivalence.** A property chain is a classical complex pattern that was introduced by Ritze et al. [18]. 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 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:

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

Example:

enslaved:Person(x) $\wedge$ enslaved:hasSexRecord(x, y) $\wedge$
 enslaved:SexRecord(x, y) $\wedge$ enslaved:hasValue(y, z) $\wedge$
 enslaved:SexTypes(z) $\leftrightarrow$ ed:Q410(Person)(x) $\wedge$
 ep:P31(hasSex)(x, y) $\wedge$ wikibase:Statement(y) $\wedge$
 eps:P31(hasSex)(y, z) $\wedge$ ed:Q291(Sex)(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 essential $\top$). Also, for the left-to-right direction, we assume that $x_2, \dots, x_n$ are existentially quantified variables.

- **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:

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

Example:

enslaved:Person(x) $\wedge$ enslaved:hasNameRecord(x, y) $\wedge$
 enslaved:NameRecord(x, y) $\leftarrow$ ed:Q410(Person)(x) $\wedge$
 ep:P20(hasName)(x, y) $\wedge$ wikibase:Statement(y)

4.3 Format in EDOAL and Rule Syntax

Most ontology alignment benchmarks are formatted according to the format provided by the Alignment API [5]. The standard alignment format is not expressive enough to capture complex relations. Fortunately, the Alignment API also provides a format called EDOAL that can be used to express these types of complex 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 (e.g. [23, 27]) 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 ontologies, can be downloaded from <http://doi.org/10.6084/m9.figshare.12400976> under the Creative Commons CC-BY 4.0 license. We apply Hermit [9] reasoning 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 Enslaved project website¹⁹ and Enslaved data in Wikibase repository website²⁰ contains more detailed information, and corresponding documentation of the project which provides users with more insights about the resource, such as all entities and relationships between them. The complex ontology alignment track in OAEI 2020²¹ also introduces the detailed information of the Enslaved benchmark, including the benchmark download link and method to evaluate the performance.

4.4 Evaluation using Complex Alignment Systems

In order to examine the quality of this benchmark to see if it is within the capability of current complex ontology alignment systems to handle in OAEI, we apply the Association Rule-based Ontology Alignment System (AROA) [25, 26] on this benchmark since AROA participated in the evaluation of OAEI 2019 and achieved the best performance in terms of F-measure [2]. Table 3 lists the relaxed precision, recall, and F-measure [6] with different thresholds of minimum support and minimum confidence in association rule mining [10, 17]. Minimum support refers to an indication of how frequently the itemset appears in the dataset, while minimum confidence refers to an indication of how often the rule has been found to be true. From Table 3, we can find that the best precision is 0.94 when the minimum support with a value of 0.03 and the minimum confidence with a value of 0.5. And it reaches the best recall of 0.39 when the minimum support value is 0.01, and the minimum confidence value is 0.5. The best F-measure is 0.51, which is achieved when the minimum support and minimum confidence are 0.01 and 1.0 respectively. Overall, the higher the minimum support and minimum confidence, the higher the precision. The lower the minimum support and minimum confidence, the higher the recall. In terms of F-measure, for the same minimum support, the

¹⁹<http://enslaved.org/>

²⁰<https://lod.enslaved.org/>

²¹<http://oaei.ontologymatching.org/2020/complex/index.html>

Table 3: The Performance of AROA system on Enslaved Benchmark

MinSupport	MinConfidence	Relaxed Precision	Relaxed Recall	Relaxed F-measure
0.01	0.5	0.63	0.39	0.49
0.01	0.6	0.64	0.39	0.49
0.01	0.7	0.65	0.38	0.49
0.01	0.8	0.71	0.38	0.49
0.01	0.9	0.72	0.38	0.50
0.01	1.0	0.80	0.38	0.51
0.02	0.5	0.66	0.38	0.48
0.02	0.6	0.66	0.37	0.48
0.02	0.7	0.67	0.37	0.48
0.02	0.8	0.72	0.36	0.48
0.02	0.9	0.72	0.36	0.48
0.02	1.0	0.82	0.36	0.50
0.03	0.5	0.94	0.27	0.42
0.03	0.6	0.94	0.27	0.42
0.03	0.7	0.94	0.27	0.42
0.03	0.8	0.94	0.27	0.42
0.03	0.9	0.94	0.27	0.42
0.03	1.0	0.94	0.27	0.43
0.04	0.5	0.94	0.26	0.41
0.04	0.6	0.94	0.26	0.41
0.04	0.7	0.94	0.26	0.41
0.04	0.8	0.94	0.26	0.41
0.04	0.9	0.94	0.26	0.41
0.04	1.0	0.94	0.26	0.42
0.05	0.5	0.94	0.26	0.41
0.05	0.6	0.94	0.26	0.41
0.05	0.7	0.94	0.26	0.41
0.05	0.8	0.94	0.26	0.41
0.05	0.9	0.94	0.26	0.41
0.05	1.0	0.94	0.26	0.41

best F-measure is usually achieved when the value of minimum confidence is 1.0. Figure 4 demonstrates the trend of the performance when the minimum confidence is set to 1.0. We can find that the variation of performance tends to be flat and steady after the minimum confidence with a value of 0.03. The reason is that the number of alignment rules generated is getting smaller, which it is reasonable to explain the higher precision, but with the lower recall. The results of more alignment systems will be available in the coming OAEI 2020. In this paper, we have not intended to focus on the improvement of the alignment algorithm. Instead, we would only like to prove that the Enslaved benchmark is within the capability of the current complex ontology alignment systems in OAEI. And based on the results, it also indicates that there is still much space for the improvement of the current alignment systems to detect more complex correspondences and solve the challenge of the knowledge graph and ontology integration problem. Thereby, the Enslaved benchmark can be considered as a useful potential resource to advance the development of the research in the complex ontology alignment field.

5 CONCLUSION

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

Figure 4: The Trend of the Performance When Minimum Confidence is 1.0

of Wikibase as a tool to represent the data, which is convenient for users with any level of expertise to use. Detecting alignments between ontologies and Wikibase knowledge graphs are helpful to solve many practical problems and enrich knowledge graphs by aligning common resources in Wikidata. In our benchmark, the alignments not only cover 1:1 simple correspondences but also contain m:n complex relations. All correspondences required to convert between the two ontologies (a key goal of ontology alignment) are guaranteed to be present. In addition, the alignment has been evaluated by domain experts from different organizations, and we also test the complex alignment systems on the benchmark to ensure high quality. Moreover, the alignments in both rule and EDOAL syntax have been published in FigShare and OAEI with an open-access license for reusability.

As future work in this area, we have put forth this benchmark into the complex track within the OAEI. We intent 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. Furthermore, we also plan to make use of Wikidata to generate more benchmarks for Multilingual ontology matching, instance matching, and knowledge graph matching tasks. We plan to generate and improve an automated alignment system to tackle the alignment problem set forth by this benchmark.

6 ACKNOWLEDGMENTS

This work was supported by The Andrew W. Mellon Foundation under Grant 1708-04732 and 1902-06575 through the project *Enslaved: Peoples of the Historic Slave Trade*.

REFERENCES

- [1] Alsayed Algergawy, Michelle Cheatham, Daniel Faria, Alfio Ferrara, Irini Fundulaki, Ian Harrow, Sven Hertling, Ernesto Jiménez-Ruiz, Naouel Karam, Abderrahmane Khia, Patrick Lambrix, Huanyu Li, Stefano Montanelli, Heiko Paulheim, Catia Pesquita, Tzanina Saveta, Daniela Schmidt, Pavel Shvaiko, Andrea Splendiani, Élodie Thiéblin, Cássia Trojahn, Jana Vatasacinová, Ondrej Zamazal, and Lu Zhou. 2018. Results of the Ontology Alignment Evaluation Initiative 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 (CEUR Workshop Proceedings)*, Pavel Shvaiko, Jérôme Euzenat, Ernesto Jiménez-Ruiz, Michelle Cheatham, and Oktie Hassanzadeh (Eds.), Vol. 2288. CEUR-WS.org, 76–116. http://ceur-ws.org/Vol-2288/oaie18_paper0.pdf
- [2] Alsayed Algergawy, Daniel Faria, Alfio Ferrara, Irini Fundulaki, Ian Harrow, Sven Hertling, Ernesto Jiménez-Ruiz, Naouel Karam, Abderrahmane Khia, Patrick Lambrix, Huanyu Li, Stefano Montanelli, Heiko Paulheim, Catia Pesquita, Tzanina Saveta, Pavel Shvaiko, Andrea Splendiani, Élodie Thiéblin, Cássia Trojahn, Jana Vatasacinová, Ondrej Zamazal, and Lu Zhou. 2019. Results of the Ontology Alignment Evaluation Initiative 2019. 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 (CEUR Workshop Proceedings)*, Pavel Shvaiko, Jérôme Euzenat, Ernesto Jiménez-Ruiz, Oktie Hassanzadeh, and Cássia Trojahn (Eds.), Vol. 2536. CEUR-WS.org, 46–85. http://ceur-ws.org/Vol-2536/oaie19_paper0.pdf
- [3] Eva Blomqvist, Karl Hammar, and Valentina Presutti. 2016. Engineering Ontologies with Patterns – The eXtreme Design Methodology. In *Ontology Engineering with Ontology Design Patterns – Foundations and Applications*, Pascal Hitzler, Aldo Gangemi, Krzysztof Janowicz, Adila Krisnadhi, and Valentina Presutti (Eds.), Studies on the Semantic Web, Vol. 25. IOS Press, 23–50.
- [4] Eva Blomqvist and Kurt Sandkuhl. 2005. Patterns in Ontology Engineering: Classification of Ontology Patterns. In *ICEIS 2005, Proceedings of the Seventh International Conference on Enterprise Information Systems, Miami, USA, May 25-28, 2005*, Chin-Sheng Chen, Joaquim Filipe, Isabel Seruca, and José Cordeiro (Eds.), 413–416.
- [5] Jérôme David, Jérôme Euzenat, François Scharffe, and Cássia Trojahn dos Santos. 2011. The Alignment API 4.0. *Semantic Web* 2, 1 (2011), 3–10. <https://doi.org/10.3233/SW-2011-0028>
- [6] Marc Ehrig and Jérôme Euzenat. 2005. 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)*, Benjamin Ashpole, Marc Ehrig, Jérôme Euzenat, and Heiner Stuckenschmidt (Eds.), Vol. 156. CEUR-WS.org. <http://ceur-ws.org/Vol-156/paper5.pdf>
- [7] Aldo Gangemi. 2005. Ontology Design Patterns for Semantic Web Content. In *The Semantic Web – ISWC 2005, 4th International Semantic Web Conference, ISWC 2005, Galway, Ireland, November 6-10, 2005, Proceedings (Lecture Notes in Computer Science)*, Yolanda Gil, Enrico Motta, V. Richard Benjamins, and Mark A. Musen (Eds.), Vol. 3729. Springer, 262–276.
- [8] Johanna Geiß, Andreas Spitz, Jannik Strötgen, and Michael Gertz. 2015. The Wikipedia location network: overcoming borders and oceans. In *Proceedings of the 9th Workshop on Geographic Information Retrieval, GIR 2015, Paris, France, November 26-27, 2015*, Ross S. Purves and Christopher B. Jones (Eds.). ACM, 2:1–2:3. <https://doi.org/10.1145/2837689.2837694>
- [9] Birte Glimm, Ian Horrocks, Boris Motik, Giorgos Stoilos, and Zhe Wang. 2014. HermiT: An OWL 2 Reasoner. *J. Autom. Reasoning* 53, 3 (2014), 245–269. <https://doi.org/10.1007/s10817-014-9305-1>
- [10] Jiawei Han, Jian Pei, Yiwen Yin, and Runying Mao. 2004. Mining Frequent Patterns without Candidate Generation: A Frequent-Pattern Tree Approach. *Data Min. Knowl. Discov.* 8, 1 (2004), 53–87. <https://doi.org/10.1023/B:DAMI.0000005258.31418.83>
- [11] Walter Hawthorne. 2020. Maranhão Inventories Slave Database. *Journal of Slavery and Data Preservation* 1, 1 (2020). <https://jsdp.enslaved.org/fullDataArticle/40-59-4>
- [12] Sven Hertling and Heiko Paulheim. 2020. The Knowledge Graph Track at OAIE - Gold Standards, Baselines, and the Golden Hammer Bias. In *The Semantic Web - 17th International Conference, ESWC 2020, Heraklion, Crete, Greece, May 31-June 4, 2020, Proceedings (Lecture Notes in Computer Science)*, Andreas Harth, Sabrina Kirrane, Axel-Cyrille Ngonga Ngomo, Heiko Paulheim, Anisa Rula, Anna Lisa Gentile, Peter Haase, and Michael Cochez (Eds.), Vol. 12123. Springer, 343–359. https://doi.org/10.1007/978-3-030-49461-2_20
- [13] Pascal Hitzler, Aldo Gangemi, Krzysztof Janowicz, Adila Krisnadhi, and Valentina Presutti (Eds.). 2016. *Ontology Engineering with Ontology Design Patterns – Foundations and Applications*. Studies on the Semantic Web, Vol. 25. IOS Press.
- [14] Pascal Hitzler and Adila Krisnadhi. 2016. On the Roles of Logical Axiomatizations for Ontologies. In *Ontology Engineering with Ontology Design Patterns – Foundations and Applications*, Pascal Hitzler, Aldo Gangemi, Krzysztof Janowicz, Adila Krisnadhi, and Valentina Presutti (Eds.), Studies on the Semantic Web, Vol. 25. IOS Press, 73–80.
- [15] Adila Krisnadhi and Pascal Hitzler. 2016. Modeling With Ontology Design Patterns: Chess Games As a Worked Example. In *Ontology Engineering with Ontology Design Patterns – Foundations and Applications*, Pascal Hitzler, Aldo Gangemi, Krzysztof Janowicz, Adila Krisnadhi, and Valentina Presutti (Eds.), Studies on the Semantic Web, Vol. 25. IOS Press, 3–21.
- [16] Jens Lehmann, Robert Isele, Max Jakob, Anja Jentzsch, Dimitris Kontokostas, Pablo N. Mendes, Sebastian Hellmann, Mohamed Morsey, Patrick van Kleef, Sören Auer, and Christian Bizer. 2015. DBpedia - A large-scale, multilingual knowledge base extracted from Wikipedia. *Semantic Web* 6, 2 (2015), 167–195. <https://doi.org/10.3233/SW-140134>
- [17] Gregory Piatetsky-Shapiro. 1991. Discovery, Analysis, and Presentation of Strong Rules. In *Knowledge Discovery in Databases*, Gregory Piatetsky-Shapiro and William J. Frawley (Eds.). AAAI/MIT Press, 229–248.
- [18] Dominique Ritze, Christian Meilicke, Ondrej Sváb-Zamazal, and Heiner Stuckenschmidt. 2009. 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 (CEUR Workshop Proceedings)*, Pavel Shvaiko, Jérôme Euzenat, Fausto Giunchiglia, Heiner Stuckenschmidt, Natalya Fridman Noy, and Arnon Rosenthal (Eds.), Vol. 551. CEUR-WS.org. http://ceur-ws.org/Vol-551/om2009_Tpaper3.pdf
- [19] Cogan Shimizu, Pascal Hitzler, Quinn Hirt, Dean Rehberger, Seila Gonzalez Estrecha, Catherine Foley, Alicia M Sheill, Walter Hawthorne, Jeff Mixer, Ethan Watrall, et al. 2020. The enslaved ontology: Peoples of the historic slave trade. *Journal of Web Semantics* (2020), 100567.
- [20] Cogan Shimizu, Pascal Hitzler, Quinn Hirt, Alicia Sheill, Seila Gonzalez, Catherine Foley, Dean Rehberger, Ethan Watrall, Walter Hawthorne, Duncan Tarr, Ryan Carty, and Jeff Mixer. 2019. *The Enslaved Ontology 1.0: Peoples of the Historic Slave Trade*. Technical Report. Enslaved: Peoples of the Historic Slave Trade. Available from <http://docs.enslaved.org>.
- [21] Pavel Shvaiko and Jérôme Euzenat. 2013. Ontology Matching: State of the Art and Future Challenges. *IEEE Trans. Knowl. Data Eng.* 25, 1 (2013), 158–176. <https://doi.org/10.1109/TKDE.2011.253>
- [22] Élodie Thiéblin, Michelle Cheatham, Cássia Trojahn dos Santos, Ondrej Zamazal, and Lu Zhou. 2018. The First Version of the OAIE 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 (CEUR Workshop Proceedings)*, Marieke van Erp, Medha Atre, Vanessa López, Kavitha Srinivas, and Carolina Fortuna (Eds.), Vol. 2180. CEUR-WS.org. <http://ceur-ws.org/Vol-2180/paper-67.pdf>
- [23] Élodie Thiéblin, Ollivier Haemmerlé, Nathalie Hernandez, and Cássia Trojahn dos Santos. 2017. 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 (CEUR Workshop Proceedings)*, Pavel Shvaiko, Jérôme Euzenat, Ernesto Jiménez-Ruiz, Michelle Cheatham, and Oktie Hassanzadeh (Eds.), Vol. 2032. CEUR-WS.org, 217–218. http://ceur-ws.org/Vol-2032/om2017_poster6.pdf
- [24] Denny Vrandečić and Markus Krötzsch. 2014. Wikidata: a free collaborative knowledgebase. *Commun. ACM* 57, 10 (2014), 78–85. <https://doi.org/10.1145/2629489>
- [25] Lu Zhou, Michelle Cheatham, and Pascal Hitzler. 2019. AROA Results for 2019 OAIE. 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 (CEUR Workshop Proceedings)*, Pavel Shvaiko, Jérôme Euzenat, Ernesto Jiménez-Ruiz, Oktie Hassanzadeh, and Cássia Trojahn (Eds.), Vol. 2536. CEUR-WS.org, 107–113. http://ceur-ws.org/Vol-2536/oaie19_paper4.pdf
- [26] Lu Zhou, Michelle Cheatham, and Pascal Hitzler. 2019. Towards Association Rule-Based Complex Ontology Alignment. In *Semantic Technology - 9th Joint International Conference, JIST 2019, Hangzhou, China, November 25-27, 2019, Proceedings (Lecture Notes in Computer Science)*, Xin Wang, Francesca Alessandra Lisi, Guohui Xiao, and Elena Botoeva (Eds.), Vol. 12032. Springer, 287–303. https://doi.org/10.1007/978-3-030-41407-8_19
- [27] Lu Zhou, Michelle Cheatham, Adila Krisnadhi, and Pascal Hitzler. 2018. 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 (Lecture Notes in Computer Science)*, Denny Vrandečić, Kalina Bontcheva, Mari Carmen Suárez-Figueroa, Valentina Presutti, Irene Celino, Marta Sabou, Lucie-Aimée Kaffee, and Elena Simperl (Eds.), Vol. 11137. Springer, 273–288. https://doi.org/10.1007/978-3-030-00668-6_17
- [28] Lu Zhou, Michelle Cheatham, Adila Krisnadhi, and Pascal Hitzler. 2019. GeoLink Data Set: A Complex Alignment Benchmark from Real-world Ontology. *Data Intelligence* (2019), 1–26.