

Towards a Benchmark for Knowledge Base Exchange

Bahar Ghadiri Bashardoost
University of Toronto
ghadiri@cs.toronto.edu

Renée J. Miller
Northeastern University
miller@northeastern.edu

Kelly Lyons
University of Toronto
klyons@cs.toronto.edu

ABSTRACT

As interest in supporting data exchange between heterogeneous knowledge bases (KBs) has increased, so has interest in benchmarking KB exchange systems. One important benchmark, LODIB (Linked Open Data Integration Benchmark), has been proposed that reflects the real and deep heterogeneity of knowledge bases in the Linked Open Data Cloud. In this position paper, we reflect on requirements for benchmarks of KB exchange systems and bring to bear important lessons learned from other data exchange systems. Specifically, we consider the importance, in data exchange, of explicitly modeling incompleteness, something that is at the heart of relational data exchange and that is solved using principled value invention methods. We also consider the incompleteness that arises naturally within knowledge bases and how that influences KB exchange. Instances of a single class in a KB may exhibit heterogeneity in their structure and modeling this is important in a KB exchange benchmark. Using these insights, we propose a set of new requirements for a KB exchange benchmark.

ACM Reference Format:

Bahar Ghadiri Bashardoost, Renée J. Miller, and Kelly Lyons. 2019. Towards a Benchmark for Knowledge Base Exchange. In *Proceedings of KDD's 1st International Workshop on Challenges and Experiences From Data Integration to Knowledge Graphs (DI2KG '19)*. ACM, New York, NY, USA, 5 pages.

1 INTRODUCTION

Machines cannot comprehend a large portion of data produced by humans. The grand vision of the Semantic Web is to be able to create a web-scale decentralized corpus of information and knowledge that can be processed automatically by machines [9]. Large-scale machine-readable KBs are one of the main enablers of the Semantic Web vision. These KBs contain rich information about domain-specific and general concepts (also known as classes or types), the relationships among them (also known as properties), and their instances (also known as individuals). KBs provide a means for machines to be able to make inferences and answer questions about the knowledge they store. To this end, it is very important that a KB is populated with relevant facts that model the domain that the KB tries to represent. However, populating an existing KB is not a trivial task. One powerful method for KB expansion is to share or exchange information between different, heterogeneous KBs.

Researchers have begun to formalize the problem of exchanging knowledge between a source and a target KB given a set of mapping

rules [2–5]. In addition, several tools have emerged which aim to automatically generate mapping rules between two KBs [25, 26, 28, 29], making KB exchange¹ more achievable than ever. Naturally, benchmarks are also beginning to be proposed to evaluate solutions to challenges that arise in practical KB exchange such as expressiveness of mapping languages used [30], or the efficiency and the quality of automatic mapping rule generation [27, 30]. Although these benchmarks propose an important set of criteria for evaluating various aspects of practical KB exchange, we believe that their evaluation criteria can be enriched to raise the bar for evaluating KB exchange.

Notably, our work considers the materialized exchange of data which can be distinguished from OBDA and OBDI. In Ontology-based Data Access (OBDA) the goal is typically to provide a virtual ontology (or KB) interface over databases expressed in different data models (also known as federation or mediation) [35]. In Ontology-based Data Integration (OBDI), also known as ontology merging, the goal is to create a single ontology (KB) that represents all information in a set of source KBs or DBs [19, 31, 34]. Solutions for OBDI mostly assume the heterogeneity is limited and can be reconciled with simple mappings like `same-as`, `subclass-of`, or `equivalent-class`. In contrast, both OBDA and KB exchange use much richer mapping rules that need to be represented using more expressive mapping languages. The benchmarking of KB exchange can inform the benchmarking of both OBDA and OBDI. In this position paper, we look at lessons learned from benchmarks for data exchange and mapping generation systems and from how KBs are being used in practice. We propose a new set of requirements for KB mapping generation tools that can be incorporated into benchmarks to broaden and deepen the tool evaluations. Having such benchmarks can attract more attention to this important area by highlighting the challenges that remain and improving the quality of research by pointing out new issues that are unique to KB exchange.

1.1 Benchmarking Data Exchange

Populating a structured information resource (target) using data translated from another structured resource (source) is one of the oldest problems in data management. Data exchange [16] is a prominent approach for solving this problem when both source and target adhere to a relational or nested relational data model (such as XML). The theory of data exchange was introduced by Fagin et al. [16]. This work laid the theoretical foundation of data exchange and identified important tasks in data exchange, including materializing a good target instance and answering queries. In data exchange, a set of mapping rules specify the relationship between a source and target schema. As defined by Fagin et al. [16], given a

¹The term *KB Exchange* was introduced by Arenas et al. [2] and we use it in this work to distinguish a specific setting of the data exchange problem in which the source and target are both KBs.

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 the author(s) 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.

DI2KG '19, August 05, 2019, Anchorage, AK

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

source schema, a target schema, an instance of the source schema, and a set of mapping rules, *data exchange* is the problem of creating an instance of the target that reflects the source instance as closely as possible. Fagin et al. defined when a target instance is a *solution* for data exchange and defined a class of good solutions (called universal solutions). They also defined a declarative mapping language (source-to-target tuple-generating-dependencies) for relational schemas that has a good balance between expressiveness and algorithmic properties and has since been widely adopted in tools and generalized to other data models [16]. An important issue in data exchange is *value invention* [18] which is required when the target models data not present in the source that plays a structural role in connecting other data [7].

The problem of data exchange assumes that mapping rules are given. However, in practice one important challenge that needs to be addressed is how to obtain these mapping rules. These rules can be defined manually, however the manual process is burdensome [33]. As a result, there is a large body of literature on systems that reduce the burden of creating these rules manually, by making the rule creation process as automatic as possible. Clio was an early mapping generation system for relational and XML schemas [23], and this has been followed by dozens of other research and industrial mapping systems [10]. Most systems use a set of correspondences between the source and target, and enrich them using implicit information that lies within the schemas (or instances) of the source and target to generate semantically correct mapping rules. Systems that use declarative mapping rules also translate them into executable programs (such as queries or scripts) that produce a single data exchange solution.

Over time, benchmarks have been developed to evaluate these systems. One of the first was STBenchmark [1] that proposed the use of mapping scenarios (or mapping patterns) that represent a set of transformations that should be supported by any mapping tool. More recently, iBench [6] permits the efficient creation of benchmarks with large and complex schemas and mappings. iBench provides control over a large set of mapping characteristics including the amount and complexity of the value invention required in a correct mapping solution. Bellahsene et al. [8] provide a survey of work on evaluating both schema matching and mapping discovery between schemas in different data models. An important (and computationally hard) issue in benchmarking data exchange is understanding if the solutions produced by mapping rules are always universal solutions [6].

1.2 Benchmarking KB Exchange

Arenas et al. [5] were one of the first to investigate data exchange among KBs. They identified the potential incompleteness of a source instance as one of the main challenges that needs to be addressed in KB exchange [2–5]. For an incomplete source instance, decisions on how to exchange the unknown entities of the source are not trivial. Another important challenge which arises in KB exchange is that both atomic values (RDF literals) and individuals (RDF IRIs) need to be exchanged.

Unlike in data exchange, there has been less work focused on automated generation of mapping rules for KB exchange [25, 26, 28, 29]. Additionally, a single declarative mapping constraint language

has not been widely adopted, so most systems generate executable mappings directly which produce a single KB exchange instance. Most of these tools generate mapping rules by enriching a set of given correspondences among two KBs, using clues that lie within the constraints of the KBs [25, 26, 28]. Buhmann et al. [12] stated that, in practice, there is a lack of KBs that contain high quality schema axioms and sufficient instance data adhering to the schema. Thus, it is important for KB mapping generation tools to rely on a small set of axioms which are commonly used in a large number of KBs to create mappings. Relying on other less used axioms (or axioms not part of a widely used schema language like RDFS) limits the applicability of a tool and the ability to compare it with others. Like many KB integration or alignment approaches [32, 34], we think it is reasonable for a mapping generation tool to require at most a small set of commonly used constructs, (for example, `rdfs:subClassOf`, `rdfs:domain`, `rdfs:range`, and `rdf:type`). Such constructs are present in most KBs (even those that are automatically created from relational or semi-structured data sets) and are part of the W3C standard OWL Web Ontology Language. We refer to tools that aid in the generation of KB mapping rules for KB exchange between two KBs as *KMGT*.

Benchmarks such as STBenchmark and metadata generation tools like iBench evaluate tools using a set of parameterized mapping scenarios (or mapping patterns). To evaluate *KMGTS*, DTS-Benchmark [27], proposed a set of patterns when the source and target are both KBs. These patterns were later extended in LODIB (linked open data integration benchmark) [30]. Since a single mapping language for KB exchange is not yet widely adopted, LODIB is mainly designed to benchmark the expressive power of mapping languages. This benchmark has been also used to evaluate *KMGTS*. The scenarios proposed in DTSBenchmark and LODIB represent an important set of transformations that should be supported by *KMGTS*. In this paper, we identify additional scenarios that we think should be supported by any mapping generation tool and any language that aims to express KB mapping rules. Our goal is to push the research on KB exchange to the next level.

2 REQUIREMENTS FOR *KMGTS*

We now describe some challenges that every *KMGT* should address in order for it to be of use in practical applications of KB exchange. Any benchmark that aims to evaluate *KMGTS* should include scenarios that can be used to test *KMGTS*' output in the presence of these challenges. Thus, in addition to describing these challenges, we include example scenarios which can be incorporated in benchmarks. Note that the challenges discussed in our paper augment and do not replace other benchmarks.

2.1 Dealing with Incomplete Correspondences

In data exchange, mapping rules are created by enriching a given set of correspondences with implicit information which lies within the schemas, the data, or which is provided by users (through crowdsourcing or visual interfaces). The algorithms proposed for generating these rules often assume that the set of given correspondences is *incomplete* [11, 20, 22, 24]. This is a realistic assumption since these correspondences are usually the output of an ontology alignment

(a.k.a., schema matching) process. Assuming that an alignment process can automatically produce *all* correspondences between two KBs with high precision is usually unrealistic [14]. When the set of correspondences is incomplete, a mapping generation tool aims to produce mapping rules based on understanding the different ways that corresponding schema elements can be associated with each other. To the best of our knowledge, current *KMG*Ts assume the input correspondences are complete. If there is no correspondence to a target element (for example a property P), current *KMG*Ts do not populate the element with any data. We argue instead that if there is source data that could be used to populate the element (for example, a source property path between source elements that are matched by correspondences to the target domain and range of P), then a useful function of a *KMG*T is to suggest a (perhaps ranked) list of alternative ways to populate the unmatched target element.

Unfortunately, the LODIB benchmark does not consider this case since as mentioned, it is designed to benchmark the expressive power of languages that represent the mapping rules. DTSBenchmark is designed to benchmark the *KMG*Ts, however the scenarios proposed in this benchmark always contain complete sets of correspondences. An example of a scenario that tests the ability of a *KMG*T to handle incomplete correspondences is given below.

Scenario 2.1. Missing Correspondence to a Target Property:

Figure 1.a represents the RDFS layer of a source and target KB and a single correspondence which has been identified between them. We argue that a *KMG*T should create mapping rules that not only translate the source Person instances into the target Person, but also suggest possible ways of populating the unmatched related property in the target. For example, a *KMG*T should suggest a set of mapping rules that express that if a Person is supervised by another Person in the source, these two Persons are related in the target. Or if a Person P_1 hasWorkedOn Project J , and a Person P_2 is a contributor to Project J in the source, then P_1 should be related to P_2 in the target. □

A *KMG*T should be able to suggest a set of mapping rules that are consistent with the schemas and correspondences (and perhaps rank them if additional information like data examples are available [33]). Indeed, a lesson learned from data exchange is that an important role of *KMG*Ts is in systematically enumerating possible mappings, something humans do poorly [15].

2.2 KB Value Invention

Of course, sometimes an unmatched target element cannot be populated with source data. This occurs often in data exchange. In order to materialize a target instance, sometimes values need to be filled in for the undetermined elements. For instance, Clio [23] creates oids (using Skolem functions), which are unique identifiers (also called a *labeled nulls*), when source data is matched to multiple relations connected by unmatched target foreign keys. These labeled nulls help capture some of the structural characteristics of the target data. Similarly, when dealing with KBs, sometimes a value needs to be created in order to preserve the associations between the resources of the target KB. As discussed by Arenas et al. [3], *blank nodes* in knowledge graphs can play the same role that *labeled nulls* play in relational and nested relational data exchange (see [17] for more

information on blank nodes). Thus similar to relational data exchange, blank nodes should be used if an association between two transferred values must be represented and the mapping rules must have enough information to guide the generation of these blank nodes. Unfortunately, in the benchmarks proposed so far, there is no scenario that clearly evaluates value invention in *KMG*Ts. We propose two scenarios in this direction.

Scenario 2.2. Value Invention 1:

Consider Figure 1.b. An office is not a contact so the IRIs of offices are not exchanged into contacts. However, the *KMG*T must be able to create a blank node which associates address and phone properly. That is, in this scenario a *KMG*T should be able to capture a source-to-target dependency which expresses that for each office in the source that has an address a , and a phone p , a and p are associated with each other in the target. □

The following example illustrates why supporting the scenario above is not trivial for *KMG*Ts. Imagine a *KMG*T produces the following queries as mapping rules to transfer data from a source to a target. The where clauses of these queries select data from the source and the construct clauses create the facts in the target KB using the target structure and the values selected from the source. These two queries should not be considered correct in a benchmark evaluation since they do not guarantee that for a single office in the source with phone *416-345* and address *Toronto* that a single blank node is created in the target associated with these same values.

```

construct{
?-:a a trgt:Contact.
?-:a trgt:phone ?trgtPhone}
where{
?o a src:Office.
?o src:phone ?srcPhone.
bind(?srcPhone as ?trgtPhone)}
&
construct{
?-:b a trgt:Contact.
?-:b trgt:address ?trgtAddress}
where{
?o a src:Office.
?o src:Address ?srcAddress.
bind(?srcAddress as ?trgtAddress)}
```

Scenario 2.3. Value Invention 2:

Consider Figure 1.c. The difference between this scenario and scenario 2.2 is that the blank nodes which need to be created are for a concept (Address) which does not have any data property which participates in a correspondence. However, appropriate blank nodes need to be created to correctly associate individual people and their countries in the target. □

2.3 Dealing with an Open-World Assumption

Traditionally, the data exchange problem is defined over a setting where the source instance is assumed to be complete and has a single interpretation. However, KBs follow an *open-world assumption* and are by nature incomplete. The mapping rules which are generated automatically must be able to handle *unknown* values in the source KB properly. Consider Figure 1.b and imagine that a *KMG*T produced the following query as a mapping rule.

```

1 construct{
2 ?-:p a trgt:Contact.
3 ?-:p trgt:phone ?trgtPhone.
4 ?-:p trgt:address ?trgtAddress}
5 where{
6 ?o a src:Office.
7 ?o src:Address ?srcAddress.
8 ?o src:phone ?srcPhone.
9 bind(?srcAddress as ?trgtAddress)
10 bind(?srcPhone as ?trgtPhone)}
```

It is straightforward to see that the mapping rule expressed by the SPARQL query above is not enough to transfer all instances of `src:Office`. For example the query above can not be used to transfer information of an `Office` which does not have a `Phone`. One way to fix the problem above is to break the above query into two queries (see Scenario 2.2). However, as we have seen in Section 2.2, breaking queries into separate smaller queries is not trivial and can result in other problems. Another approach when dealing with missing values of KBs is to use SPARQL's optional keyword [36] and this approach can be adapted in KB exchange. Whatever the approach, it is important that a *KMGT* create a correct target instance even when the source is incomplete and that a benchmark include scenarios to test this. An example scenario for testing the ability of a *KMGT* to deal with an incomplete source instance is given below.

Scenario 2.4. Missing Values in the Source:

Consider Figure 1.b. Also, assume that the source instance contains the following facts: `Office(office1)`, `Office(office2)`, `phone(office1, '416-345')`, `phone(office2, '416-456')`, `address(office1, 'Toronto')`. In this case, it is expected that the mapping rule can materialize a correct instance of the target such as the following (where `c1` and `c2` are blank nodes) `Contact(-:c1)`, `Contact(-:c2)`, `phone(-:c1, '416-345')`, `phone(-:c2, '416-456')`, `address(-:c1, 'Toronto')`. It is important that all data from the source including `office2` with an unknown address be mapped to the target. □

2.4 Dealing with Cycles

A cycle can be used in a KB to model the relationships between individuals of the same type. For instance, Figure 1.a depicts many cycles each modeling a relationship between two people. This type of cycle is very common. For instance, any KB created from a social network contains a large number of these cycles, since social networks model various relationships between instances of a specific type. We believe that an algorithm that generates mapping rules should be able to handle this type of cycle. Scenario 2.5 can be used to test whether a *KMGT* can correctly map a cycle. We have said that an important role of *KMGTS* is in systematically enumerating possible mappings. Cycles pose a challenge as the possible mappings become infinite.

Scenario 2.5. Property path with same domain and range:

Consider Figure 1.a. It is expected that a *KMGT* can generate mapping rules based on cycles of source and target KBs. For instance a mapping rule for the setting introduced in this scenario might express something like ²:

$$\forall(x, y), x \neq y,$$

$$\text{where } (x, y) \in [(((hasSupervisor^* + (hasWorkedOn.contributor)^*) \cdot (hasSupervisor^* + (hasWorkedOn.contributor)^* + hasSupervisor^{-*}))^*)]_{src}$$

$$\text{then } (x, y) \in [related]_{tgt}$$

What is important is for a *KMGT* to consider alternatives to help a mapping designer to arrive at a semantically correct mapping. □

Not all mapping languages will contain recursion so a *KMGT* may create mappings that only traverse a cycle a fix number of times. Nonetheless, a benchmark should include scenarios containing cycles and where different mappings are desired, some that include

²Notation borrowed from Kostylev et al. [21].

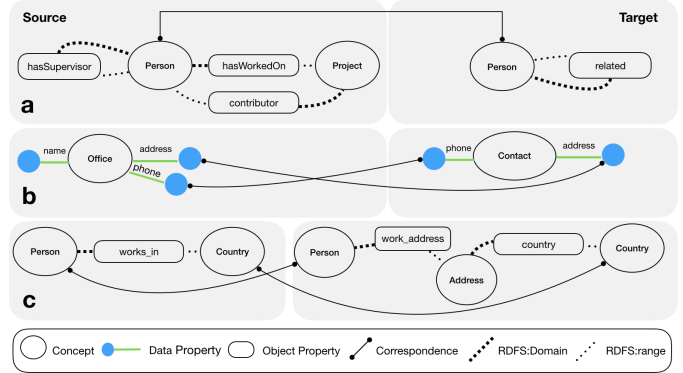


Figure 1: RDFS layer of Source and target KBs and the correspondences between them

a single traversal (e.g., only an immediate supervisor) and some that include more (e.g., all supervisors and their supervisors).

3 VISION FOR KB EXCHANGE BENCHMARKS

Data exchange between heterogeneous data sources is important for sharing data among organizations but, until recently, much of the data exchange research has focused on relational and nested relational data models. Initiatives such as ontology based data access (OBDA) [35] aim to facilitate the exchange of data between a relational source and a target KB. New exciting work is emerging on the theory and practice of KB exchange, where a KB can be used to expand the knowledge contained in another KB. We believe that expanding the current set of benchmarks in a principled way can enhance research in both mapping generation and KB exchange.

Our position paper enumerates a few important benchmarking issues, but of course is not complete. One issue that we did not include is dealing with multiple inheritance. Suppose that in a source KB two concepts `Employee` and `Student` are sub-classes of another concept `Person`. Now assume that instance *i* is both an `Employee` and a `Student` and instance *j* is only an `Employee`. Then the mapping rules generated should be able to transfer *i* and *j* to the target while respecting the fact that *i* and *j* share some properties while differing on other properties.

Benchmarking KB exchange presents many important new research challenges. Testing whether a set of mappings is equivalent can already be undecidable for data exchange and is more complex for KB exchange [3]. For a benchmark, we must develop new scalable ways of testing if mappings (or KB exchange solutions) for a benchmark are correct and sufficient. Duo et al. [13] argue that when it comes to the problem of exchanging data between two KBs, even checking that a given set of mapping rules is consistent with each other and with the axioms of the source and the target KBs is not a trivial task. The consistency of a set of mapping rules generated by a *KMGT* can be a powerful measure of quality that benchmarks can report. Thus, we also suggest new approaches need to be proposed to be able to efficiently evaluate the consistency of the mapping rules generated by a *KMGT*.

ACKNOWLEDGMENTS

This research was supported in part by a Natural Sciences and Engineering Research Council (NSERC) Strategic Partnership Grant.

REFERENCES

- [1] Bogdan Alexe, Wang-Chiew Tan, and Yannis Velegarakis. 2008. STBenchmark: towards a benchmark for mapping systems. *Proceedings of the VLDB Endowment* 1, 1 (2008), 230–244.
- [2] Marcelo Arenas, Elena Botoeva, and Diego Calvanese. 2011. Knowledge base exchange. In *24th International Workshop on Description Logics*. 4.
- [3] Marcelo Arenas, Elena Botoeva, Diego Calvanese, and Vladislav Ryzhikov. 2016. Knowledge base exchange: The case of OWL 2 QL. *Artificial Intelligence* 238 (2016), 11–62.
- [4] Marcelo Arenas, Elena Botoeva, Diego Calvanese, Vladislav Ryzhikov, and Evgeny Sherkhonov. 2012. Exchanging description logic knowledge bases. AAAI Press.
- [5] Marcelo Arenas, Jorge Pérez, and Juan Reutter. 2013. Data exchange beyond complete data. *Journal of the ACM (JACM)* 60, 4 (2013), 28.
- [6] Patricia C Arocena, Boris Glavic, Radu Ciucanu, and Renée J Miller. 2015. The iBench integration metadata generator. *Proceedings of the VLDB Endowment* 9, 3 (2015), 108–119.
- [7] Patricia C. Arocena, Boris Glavic, and Renée J. Miller. 2013. Value invention in data exchange. In *SIGMOD*. 157–168.
- [8] Zohra Bellahsene, Angela Bonifati, Fabien Duchateau, and Yannis Velegarakis. 2011. On evaluating schema matching and mapping. In *Schema matching and mapping*. Springer, 253–291.
- [9] Tim Berners-Lee, James Hendler, and Ora Lassila. 2001. The semantic web. *Scientific american* 284, 5 (2001), 28–37.
- [10] Angela Bonifati, Giansalvatore Mecca, Paolo Papotti, and Yannis Velegarakis. 2011. Discovery and correctness of schema mapping transformations. In *Schema matching and mapping*. Springer, 111–147.
- [11] A. Bonifati, G. Mecca, A. Pappalardo, S. Raunich, and G. Summa. 2008. Schema Mapping Verification: The Spicy Way. In *Proceedings of the 11th International Conference on Extending Database Technology: Advances in Database Technology (EDBT '08)*. ACM, New York, NY, USA, 85–96.
- [12] Lorenz Bühmann and Jens Lehmann. 2012. Universal OWL axiom enrichment for large knowledge bases. In *International Conference on Knowledge Engineering and Knowledge Management*. Springer, 57–71.
- [13] Dejing Dou, Drew McDermott, and Peishen Qi. 2005. Ontology translation on the semantic web. In *Journal on data semantics II*. Springer, 35–57.
- [14] Jérôme Euzenat and Pavel Shvaiko. 2013. *Ontology Matching* (2nd ed.). Springer Publishing Company, Incorporated.
- [15] Ronald Fagin, Laura Haas, Mauricio Hernández, Renée Miller, Lucian Popa, and Yannis Velegarakis. 2009. Clio: Schema mapping creation and data exchange. *Conceptual modeling: foundations and applications* (2009), 198–236.
- [16] Ronald Fagin, Phokion G Kolaitis, Renée J Miller, and Lucian Popa. 2005. Data exchange: semantics and query answering. *Theoretical Computer Science* 336, 1 (2005), 89–124.
- [17] Aidan Hogan, Marcelo Arenas, Alejandro Mallea, and Axel Polleres. 2014. Everything you always wanted to know about blank nodes. *Journal of Web Semantics* 27 (2014), 42–69.
- [18] Richard Hull and Masatoshi Yoshikawa. 1990. ILOG: Declarative Creation and Manipulation of Object Identifiers. In *VLDB*. 455–468.
- [19] Ernesto Jiménez-Ruiz and Bernardo Cuenca Grau. 2011. Logmap: Logic-based and scalable ontology matching. In *International Semantic Web Conference*. Springer, 273–288.
- [20] A. Kimmig, A. Memory, R. J. Miller, and L. Getoor. 2017. A Collective, Probabilistic Approach to Schema Mapping. In *2017 IEEE 33rd International Conference on Data Engineering (ICDE)*. 921–932.
- [21] Egor V Kostylev, Juan L Reutter, Miguel Romero, and Domagoj Vrgoč. 2015. SPARQL with property paths. In *International Semantic Web Conference*. Springer, 3–18.
- [22] Renée J Miller, Laura M Haas, and Mauricio A Hernández. 2000. Schema Mapping as Query Discovery. In *VLDB*, Vol. 2000. 77–88.
- [23] Lucian Popa, Yannis Velegarakis, Mauricio A. Hernández, Renée J. Miller, and Ronald Fagin. 2002. Translating Web Data. In *Proceedings of the 28th International Conference on Very Large Data Bases (VLDB '02)*. VLDB Endowment, 598–609.
- [24] Li Qian, Michael J Cafarella, and HV Jagadish. 2012. Sample-driven schema mapping. In *Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data*. ACM, 73–84.
- [25] Han Qin, Dejing Dou, and Paea LePendu. 2007. Discovering executable semantic mappings between ontologies. In *On the Move to Meaningful Internet Systems 2007: CoopIS, DOA, ODBASE, GADA, and IS* (2007), 832–849.
- [26] Carlos R Rivero, Inma Hernández, David Ruiz, and Rafael Corchuelo. 2011. Generating SPARQL executable mappings to integrate ontologies. In *International Conference on Conceptual Modeling*. Springer, 118–131.
- [27] Carlos R Rivero, Inma Hernández, David Ruiz, and Rafael Corchuelo. 2011. On benchmarking data translation systems for semantic-web ontologies. In *Proceedings of the 20th ACM international conference on Information and knowledge management*. ACM, 1613–1618.
- [28] Carlos R Rivero, Inma Hernández, David Ruiz, and Rafael Corchuelo. 2013. Exchanging data amongst linked data applications. *Knowledge and information systems* 37, 3 (2013), 693–729.
- [29] Carlos R Rivero, Inma Hernández, David Ruiz, and Rafael Corchuelo. 2016. Mapping RDF knowledge bases using exchange samples. *Knowledge-Based Systems* 93 (2016), 47–66.
- [30] Carlos R Rivero, Andreas Schultz, Christian Bizer, and David Ruiz Cortés. 2012. Benchmarking the performance of linked data translation systems. In *LDOW 2012: WWW2012 Workshop on Linked Data on the Web (2012)*. CEUR-WS.
- [31] Gerd Stumme and Alexander Maedche. 2001. FCA-Merge: Bottom-up merging of ontologies. In *IJCAI*, Vol. 1. 225–230.
- [32] Fabian M. Suchanek, Serge Abiteboul, and Pierre Senellart. 2011. PARIS: Probabilistic Alignment of Relations, Instances, and Schema. *Proc. VLDB Endow* 5, 3 (Nov. 2011), 157–168.
- [33] Balder ten Cate, Phokion G. Kolaitis, and Wang-Chiew Tan. 2013. Schema Mappings and Data Examples. In *Proceedings of the 16th International Conference on Extending Database Technology (EDBT '13)*. ACM, New York, NY, USA, 777–780.
- [34] Octavian Udrea, Lise Getoor, and Renée J. Miller. 2007. Leveraging Data and Structure in Ontology Integration. In *Proceedings of the 2007 ACM SIGMOD International Conference on Management of Data (SIGMOD '07)*. ACM, New York, NY, USA, 449–460.
- [35] Guohui Xiao, Diego Calvanese, Roman Kontchakov, Domenico Lembo, Antonella Poggi, Riccardo Rosati, and Michael Zakharyashev. 2018. Ontology-based data access: A survey. *IJCAI*.
- [36] Guohui Xiao, Roman Kontchakov, Benjamin Cogrel, Diego Calvanese, and Elena Botoeva. 2018. Efficient Handling of SPARQL optional for OBDA. In *International Semantic Web Conference*. Springer, 354–373.