

Introduction to Linked Data and its Lifecycle on the Web

Axel-Cyrille Ngonga Ngomo, Sören Auer, Jens Lehmann, Amrapali Zaveri

AKSW, Institut für Informatik, Universität Leipzig, Pf 100920, 04009 Leipzig
{lastname}@informatik.uni-leipzig.de
<http://aksw.org>

Abstract. With Linked Data, a very pragmatic approach towards achieving the vision of the Semantic Web has gained some traction in the last years. The term Linked Data refers to a set of best practices for publishing and interlinking structured data on the Web. While many standards, methods and technologies developed within by the Semantic Web community are applicable for Linked Data, there are also a number of specific characteristics of Linked Data, which have to be considered. In this article we introduce the main concepts of Linked Data. We present an overview of the Linked Data life-cycle and discuss individual approaches as well as the state-of-the-art with regard to extraction, authoring, linking, enrichment as well as quality of Linked Data. We conclude the chapter with a discussion of issues, limitations and further research and development challenges of Linked Data. This article is an updated version of a similar lecture given at Reasoning Web Summer School 2013.

1 Introduction

One of the biggest challenges in the area of intelligent information management is the exploitation of the Web as a platform for data and information integration as well as for search and querying. Just as we publish unstructured textual information on the Web as HTML pages and search such information by using keyword-based search engines, we are already able to easily publish structured information, reliably interlink this information with other data published on the Web and search the resulting data space by using more expressive querying beyond simple keyword searches. The Linked Data paradigm has evolved as a powerful enabler for the transition of the current document-oriented Web into a Web of interlinked Data and, ultimately, into the Semantic Web. The term Linked Data here refers to a set of best practices for publishing and connecting structured data on the Web. These best practices have been adopted by an increasing number of data providers over the past three years, leading to the creation of a global data space that contains many billions of assertions – the Web of Linked Data (cf. Figure 1).

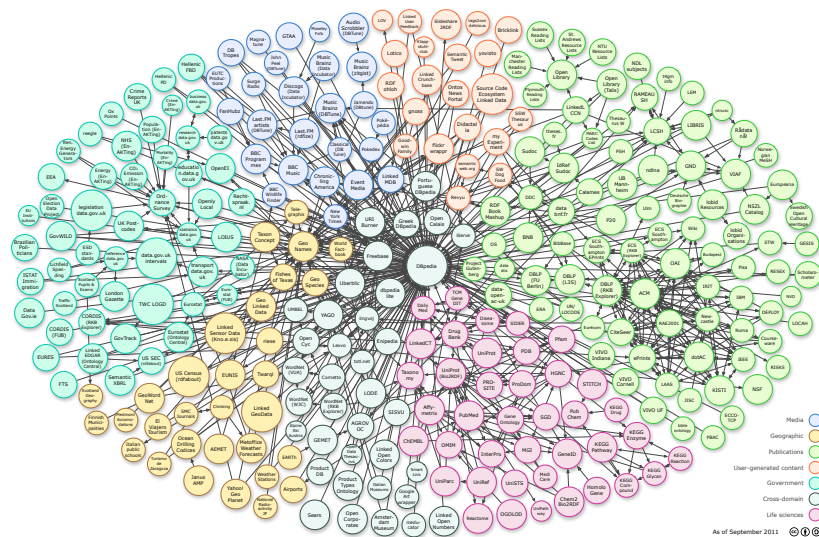


Fig.1: Overview of some of the main Linked Data knowledge bases and their interlinks available on the Web. (This overview is published regularly at <http://lod-cloud.net> and generated from the Linked Data packages described at the dataset metadata repository ckan.net.)

In this chapter we give an overview of recent development in the area of Linked Data management. The different stages in the linked data life-cycle [11] are depicted in Figure 2. Information represented in unstructured form or adhering to other structured or semi-structured representation formalisms must be mapped to the RDF data model (*Extraction*). Once there is a critical mass of RDF data, mechanisms have to be in place

to store, index and query this RDF data efficiently (*Storage & Querying*). Users must have the opportunity to create new structured information or to correct and extend existing ones (*Authoring*). If different data publishers provide information about the same or related entities, links between those different information assets have to be established (*Linking*). Since Linked Data primarily comprises instance data we observe a lack of classification, structure and schema information. This deficiency can be tackled by approaches for enriching data with higher-level structures in order to be able to aggregate and query the data more efficiently (*Enrichment*). As with the Document Web, the Data Web contains a variety of information of different quality. Hence, it is important to devise strategies for assessing the quality of data published on the Data Web (*Quality Analysis*). Once problems are detected, strategies for repairing these problems and supporting the evolution of Linked Data are required (*Evolution & Repair*). Last but not least, users have to be empowered to browse, search and explore the structure information available on the Data Web in a fast and user friendly manner (*Search, Browsing & Exploration*).

These different stages of the linked data life-cycle do not exist in isolation or are passed in a strict sequence, but support each other. Examples include the following:

- The detection of mappings on the schema level support instance-level matching and vice versa.
- Ontology schema mismatches between knowledge bases can be compensated for by learning which concepts of one are equivalent to which concepts of the other knowledge base.
- Feedback and input from end users can be taken as training input (i.e., as positive or negative examples) for machine-learning techniques to perform inductive reasoning on larger knowledge bases, whose results can again be assessed by end users for iterative refinement.
- Semantically-enriched knowledge bases improve the detection of inconsistencies and modelling problems, which in turn results in benefits for linking, fusion, and classification.
- The querying performance of the RDF data management directly affects all other components and the nature of queries issued by the components affects the RDF data management.

As a result of such interdependence, we envision the Web of Linked Data to realize an improvement cycle for knowledge bases, in which an improvement of a knowledge base with regard to one aspect (e.g., a new alignment with another interlinking hub) triggers a number of possible further improvements (e.g., additional instance matches).

The use of Linked Data offers a number of significant benefits:

- *Uniformity*. All datasets published as Linked Data share a uniform data model, the RDF statement data model. With this data model all information is represented in facts expressed as triples consisting of a subject, predicate and object. The elements used in subject, predicate or object positions are mainly globally unique identifiers (IRI/URI). Literals, i.e., typed data values, can be used at the object position.
- *De-referencability*. URIs are not just used for identifying entities, but since they can be used in the same way as URLs they also enable locating and retrieving resources describing and representing these entities on the Web.

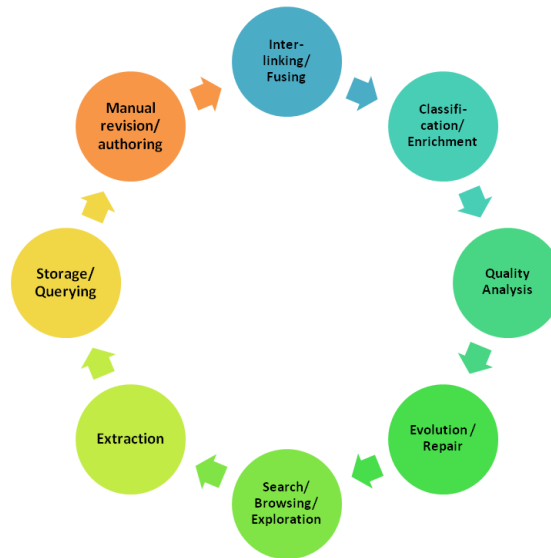


Fig. 2: The Linked Data life-cycle.

- *Coherence*. When an RDF triple contains URIs from different namespaces in subject and object position, this triple basically establishes a link between the entity identified by the subject (and described in the source dataset using namespace A) with the entity identified by the object (described in the target dataset using namespace B). Through the typed RDF links, data items are effectively interlinked.
- *Integrability*. Since all Linked Data sources share the RDF data model, which is based on a single mechanism for representing information, it is very easy to attain a syntactic and simple semantic integration of different Linked Data sets. A higher level semantic integration can be achieved by employing schema and instance matching techniques and expressing found matches again as alignments of RDF vocabularies and ontologies in terms of additional triple facts.
- *Timeliness*. Publishing and updating Linked Data is relatively simple thus facilitating a timely availability. In addition, once a Linked Data source is updated it is straightforward to access and use the updated data source, since time consuming and error prone extraction, transformation and loading is not required.

Representation \ degree of openness	Possibly closed	Open (cf. opendefinition.org)
<i>Structured data model</i> (i.e. XML, CSV, SQL etc.)	Data	Open Data
<i>RDF data model</i> (published as Linked Data)	Linked Data (LD)	Linked Open Data (LOD)

Table 1: Juxtaposition of the concepts Linked Data, Linked Open Data and Open Data.

The development of research approaches, standards, technology and tools for supporting the Linked Data lifecycle data is one of the main challenges. Developing adequate and pragmatic solutions to these problems can have a substantial impact on science, economy, culture and society in general. The publishing, integration and aggregation of statistical and economic data, for example, can help to obtain a more precise and timely picture of the state of our economy. In the domain of health care and life sciences making sense of the wealth of structured information already available on the Web can help to improve medical information systems and thus make health care more adequate and efficient. For the media and news industry, using structured background information from the Data Web for enriching and repurposing the quality content can facilitate the creation of new publishing products and services. Linked Data technologies can help to increase the flexibility, adaptability and efficiency of information management in organizations, be it companies, governments and public administrations or online communities. For end-users and society in general, the Data Web will help to obtain and integrate required information more efficiently and thus successfully manage the transition towards a knowledge-based economy and an information society.

Structure of this chapter. This chapter aims to explain the foundations of Linked Data and introducing the different aspects of the Linked Data lifecycle by highlighting particular approaches and providing references to related work and further reading. We start by briefly explaining the principles underlying the Linked Data paradigm in Section 2. The first aspect of the Linked Data lifecycle is the extraction of information from unstructured, semi-structured and structured sources and their representation according to the RDF data model (Section 3). We present the user friendly authoring and manual revision aspect of Linked Data with the example of Semantic Wikis in Section 4. The interlinking aspect is tackled in Section 5 and gives an overview on the LIMS framework. We describe how the instance data published and commonly found on the Data Web can be enriched with higher level structures in Section 6. We present an overview of the various data quality dimensions and metrics along with currently existing tools for data quality assessment of Linked Data in Section 7. Due to space limitations we omit a detailed discussion of the evolution as well as search, browsing and exploration aspects of the Linked Data lifecycle in this chapter. The chapter is concluded by several sections on promising applications of Linked Data and semantic technologies, in particular Open Governmental Data, Semantic Business Intelligence and Statistical and Economic Data. Overall, this is an updated version of a similar lecture given at Reasoning Web Summer School 2013 [13].

2 The Linked Data paradigm

In this section we introduce the basic principles of Linked Data. The section is partially based on the Section 2 from [79]. The term Linked Data refers to a set of best practices for publishing and interlinking structured data on the Web. These best practices were introduced by Tim Berners-Lee in his Web architecture note Linked Data¹ and have become known as the Linked Data principles. These principles are:

- Use URIs as names for things.
- Use HTTP URIs so that people can look up those names.
- When someone looks up a URI, provide useful information, using the standards (RDF, SPARQL).
- Include links to other URIs, so that they can discover more things.

The basic idea of Linked Data is to apply the general architecture of the World Wide Web [95] to the task of sharing structured data on global scale. The Document Web is built on the idea of setting hyperlinks between Web documents that may reside on different Web servers. It is built on a small set of simple standards: Uniform Resource Identifiers (URIs) and their extension Internationalized Resource Identifiers (IRIs) as globally unique identification mechanism [22], the Hypertext Transfer Protocol (HTTP) as universal access mechanism [56], and the Hypertext Markup Language (HTML) as a widely used content format [88]. Linked Data builds directly on Web architecture and applies this architecture to the task of sharing data on global scale.

2.1 Resource identification with IRIs

To publish data on the Web, the data items in a domain of interest must first be identified. These are the things whose properties and relationships will be described in the data, and may include Web documents as well as real-world entities and abstract concepts. As Linked Data builds directly on Web architecture, the Web architecture term *resource* is used to refer to these *things of interest*, which are in turn identified by HTTP URIs. Linked Data uses only HTTP URIs, avoiding other URI schemes such as URNs [136] and DOIs². The structure of HTTP URIs looks as follows:

[scheme:] [//authority] [path] [?query] [#fragment]

A URI for identifying Shakespeare’s ‘Othello’, for example, could look as follows:

`http://de.wikipedia.org/wiki/Othello#id`

HTTP URIs make good names for two reasons:

1. They provide a simple way to create globally unique names in a decentralized fashion, as every owner of a domain name or delegate of the domain name owner may create new URI references.
2. They serve not just as a name but also as a means of accessing information describing the identified entity.

¹ <http://www.w3.org/DesignIssues/LinkedData.html>.

² <http://www.doi.org/hb.html>

2.2 De-referencability

Any HTTP URI should be de-referencable, meaning that HTTP clients can look up the URI using the HTTP protocol and retrieve a description of the resource that is identified by the URI. This applies to URIs that are used to identify classic HTML documents, as well as URIs that are used in the Linked Data context to identify real-world objects and abstract concepts. Descriptions of resources are embodied in the form of Web documents. Descriptions that are intended to be read by humans are often represented as HTML. Descriptions that are intended for consumption by machines are represented as RDF data. Where URIs identify real-world objects, it is essential to not confuse the objects themselves with the Web documents that describe them. It is therefore common practice to use different URIs to identify the real-world object and the document that describes it, in order to be unambiguous. This practice allows separate statements to be made about an object and about a document that describes that object. For example, the creation year of a painting may be rather different to the creation year of an article about this painting. Being able to distinguish the two through use of different URIs is critical to the consistency of the Web of Data.

The Web is intended to be an information space that may be used by humans as well as by machines. Both should be able to retrieve representations of resources in a form that meets their needs, such as HTML for humans and RDF for machines. This can be achieved using an HTTP mechanism called content negotiation [56]. The basic idea of content negotiation is that HTTP clients send HTTP headers with each request to indicate what kinds of documents they prefer. Servers can inspect these headers and select an appropriate response. If the headers indicate that the client prefers HTML then the server will respond by sending an HTML document. If the client prefers RDF, then the server will send the client an RDF document.

There are two different strategies to make URIs that identify real-world objects de-referencable [179]. Both strategies ensure that objects and the documents that describe them are not confused and that humans as well as machines can retrieve appropriate representations.

303 URIs. Real-world objects can not be transmitted over the wire using the HTTP protocol. Thus, it is also not possible to directly de-reference URIs that identify real-world objects. Therefore, in the 303 URI strategy, instead of sending the object itself over the network, the server responds to the client with the HTTP response code 303 See Other and the URI of a Web document which describes the real-world object. This is called a *303 redirect*. In a second step, the client de-references this new URI and retrieves a Web document describing the real-world object.

Hash URIs. A widespread criticism of the 303 URI strategy is that it requires two HTTP requests to retrieve a single description of a real-world object. One option for avoiding these two requests is provided by the hash URI strategy. The hash URI strategy builds on the characteristic that URIs may contain a special part that is separated from the base part of the URI by a hash symbol (#). This special part is called the fragment identifier. When a client wants to retrieve a hash URI the HTTP protocol requires the fragment part to be stripped off before requesting the URI from the server. This means a URI that

includes a hash cannot be retrieved directly, and therefore does not necessarily identify a Web document. This enables such URIs to be used to identify real-world objects and abstract concepts, without creating ambiguity [179].

Both approaches have their advantages and disadvantages. Section 4.4. of the W3C Interest Group Note *Cool URIs for the Semantic Web* compares the two approaches [179]: Hash URIs have the advantage of reducing the number of necessary HTTP round-trips, which in turn reduces access latency. The downside of the hash URI approach is that the descriptions of all resources that share the same non-fragment URI part are always returned to the client together, irrespective of whether the client is interested in only one URI or all. If these descriptions consist of a large number of triples, the hash URI approach can lead to large amounts of data being unnecessarily transmitted to the client. 303 URIs, on the other hand, are very flexible because the redirection target can be configured separately for each resource. There could be one describing document for each resource, or one large document for all of them, or any combination in between. It is also possible to change the policy later on.

2.3 RDF Data Model

The RDF data model [1] represents information as sets of statements, which can be visualized as node-and-arc-labeled directed graphs. The data model is designed for the integrated representation of information that originates from multiple sources, is heterogeneously structured, and is represented using different schemata. RDF can be viewed as a *lingua franca*, capable of moderating between other data models that are used on the Web.

In RDF, information is represented in statements, called RDF triples. The three parts of each triple are called its subject, predicate, and object. A triple mimics the basic structure of a simple sentence, such as for example:

Burkhard Jung	is the mayor of	Leipzig
(subject)	(predicate)	(object)

The following is the formal definition of RDF triples as it can be found in the W3C RDF standard [1].

Definition 1 (RDF Triple). Assume there are pairwise disjoint infinite sets I , B , and L representing IRIs, blank nodes, and RDF literals, respectively. A triple $(v_1, v_2, v_3) \in (I \cup B) \times I \times (I \cup B \cup L)$ is called an *RDF triple*. In this tuple, v_1 is the subject, v_2 the predicate and v_3 the object. We call $T = I \cup B \cup L$ the set of *RDF terms*.

The main idea is to use IRIs as identifiers for entities in the subject, predicate and object positions in a triple. Data values can be represented in the object position as literals. Furthermore, the RDF data model also allows in subject and object positions the use of identifiers for unnamed entities (called blank nodes), which are not globally unique and can thus only be referenced locally. However, the use of blank nodes is discouraged in the Linked Data context as we discuss below. Our example fact sentence about Leipzig's mayor would now look as follows:

`<http://leipzig.de/id>`
`<http://example.org/p/hasMayor>`
`<http://Burkhard-Jung.de/id>` .
(subject) (predicate) (object)

This example shows that IRIs used within a triple can originate from different namespaces thus effectively facilitating the mixing and mashing of different RDF vocabularies and entities from different Linked Data knowledge bases. A triple having identifiers from different knowledge bases at subject and object position can be also viewed as an typed link between the entities identified by subject and object. The predicate then identifies the type of link. If we combine different triples we obtain an RDF graph.

Definition 2 (RDF Graph). A finite set of RDF triples is called RDF graph. The RDF graph itself represents an resource, which is located at a certain location on the Web and thus has an associated IRI, the graph IRI.

An example of an RDF graph is depicted in Figure 3. Each unique subject or object contained in the graph is visualized as a node (i.e. oval for resources and rectangle for literals). Predicates are visualized as labeled arcs connecting the respective nodes. There are a number of synonyms being used for RDF graphs, all meaning essentially the same but stressing different aspects of an RDF graph, such as *RDF document* (file perspective), *knowledge base* (collection of facts), *vocabulary* (shared terminology), *ontology* (shared logical conceptualization).

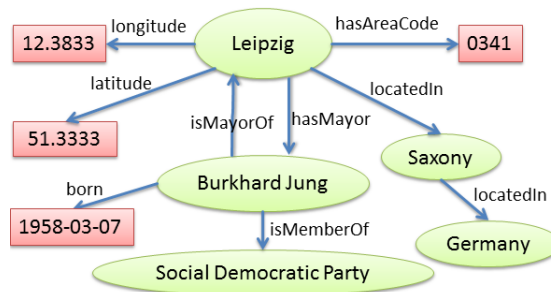


Fig. 3: Example RDF graph describing the city of Leipzig and its mayor.

Problematic RDF features in the Linked Data Context Besides the features mentioned above, the RDF Recommendation [1] also specifies some other features. In order to make it easier for clients to consume data only the subset of the RDF data model described above should be used. In particular, the following features are problematic when publishing RDF as Linked Data:

- *RDF reification* (for making statements about statements) should be avoided if possible, as reified statements are rather cumbersome to query with the SPARQL query

language. In many cases using reification to publish metadata about individual RDF statements can be avoided by attaching the respective metadata to the RDF document containing the relevant triples.

- *RDF collections* and *RDF containers* are also problematic if the data needs to be queried with SPARQL. Therefore, in cases where the relative ordering of items in a set is not significant, the use of multiple triples with the same predicate is recommended.
- The scope of *blank nodes* is limited to the document in which they appear, meaning it is not possible to create links to them from external documents. In addition, it is more difficult to merge data from different sources when blank nodes are used, as there is no URI to serve as a common key. Therefore, all resources in a data set should be named using IRI references.

2.4 RDF serializations

The initial official W3C RDF standard [1] comprised a serialization of the RDF data model in XML called *RDF/XML*. Its rationale was to integrate RDF with the existing XML standard, so it could be used smoothly in conjunction with the existing XML technology landscape. However, *RDF/XML* turned out to be difficult to understand for the majority of potential users because it requires to be familiar with two data models (i.e., the tree-oriented XML data model as well as the statement oriented RDF datamodel) and interactions between them, since RDF statements are represented in XML. As a consequence, with *N-Triples*, *Turtle* and *N3* a family of alternative text-based RDF serializations was developed, whose members have the same origin, but balance differently between readability for humans and machines. Later in 2009, *RDFa* (RDF Annotations, [3]) was standardized by the W3C in order to simplify the integration of HTML and RDF and to allow the joint representation of structured and unstructured content within a single source HTML document. Another RDF serialization, which is particularly beneficial in the context of JavaScript web applications and mashups is the serialization of RDF in JSON. In the sequel we present each of these RDF serializations in some more detail. Figure 5 presents an example serialized in the most popular serializations.

N-Triples. This serialization format was developed specifically for RDF graphs. The goal was to create a serialization format which is very simple. N-Triples are easy to parse and generate by software. An N-Triples document consists of a set of triples, which are separated ‘.’ (lines 1-2, 3-4 and 5-6 in Figure 5 contain one triple each). URI components of a triple are written in full and enclosed by ‘<’ and ‘>’. Literals are enclosed in quotes, datatypes can be appended to a literal using ‘g’ (line 6), language tags using ‘@’ (line 4). They are a subset of *Notation 3* and *Turtle* but lack, for example, shortcuts such as CURIEs. This makes them less readable and more difficult to create manually. Another disadvantage is that N-triples use only the 7-bit US-ASCII character encoding instead of UTF-8.

Turtle. Turtle (Terse RDF Triple Language) is a subset of *Notation 3* (and is consequently compatible with *Notation 3*) and a superset of the minimal N-Triples format (cf. Figure 4). The goal was to use the essential parts of *Notation 3* for the serialization of

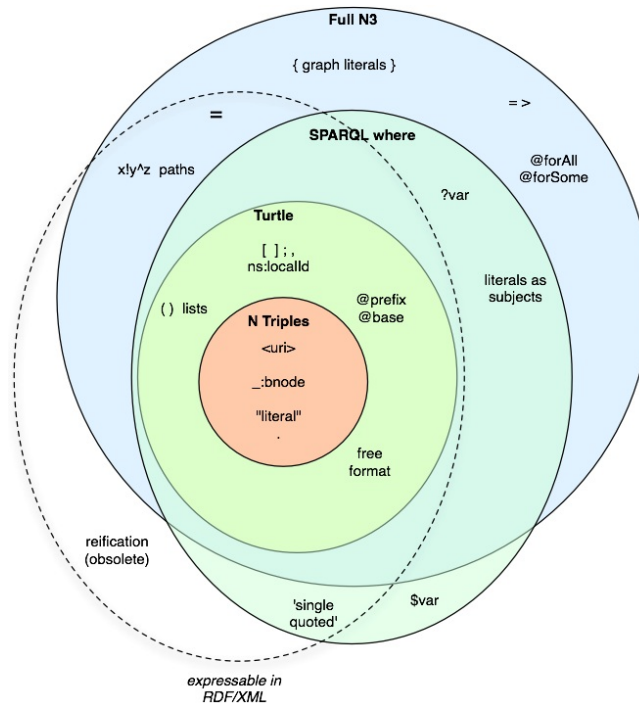


Fig. 4: Various textual RDF serializations as subsets of N3 (from [176]).

RDF models and omit everything else. Turtle became part of the SPARQL query language for expressing graph patterns. Compared to N-Triples, Turtle introduces a number of shortcuts, such as namespace definitions (lines 1-5 in Figure 5), the semicolon as a separator between triples sharing the same subject (which then does not have to be repeated in subsequent triples) and the comma as a separator between triples sharing the same subject and predicate. Turtle, just like Notation 3, is human-readable, and can handle the "%" character in URIs (required for encoding special characters) as well as IRIs due to its UTF-8 encoding.

Notation 3. N3 (Notation 3) was devised by Tim Berners-Lee and developed for the purpose of serializing RDF. The main aim was to create a very human-readable serialization. Hence, an RDF model serialized in N3 is much more compact than the same model in RDF/XML but still allows a great deal of expressiveness even going beyond the RDF data model in some aspects. Since the encoding for N3 files is UTF-8, the use of IRIs is natively supported by this format.

RDF/XML. The RDF/XML syntax [130] is standardized by the W3C and is widely used to publish Linked Data on the Web. However, the syntax is also viewed as difficult for humans to read and write, and therefore consideration should be given to using other serializations in data management and curation workflows that involve human

intervention, and to the provision of alternative serializations for consumers who may wish to eyeball the data. The MIME type that should be used for RDF/XML within HTTP content negotiation is `application/rdf+xml`.

RDFa. RDF in Attributes (RDFa, [3]) was developed for embedding RDF into XHTML pages. Since it is an extension to the XML based XHTML, UTF-8 and UTF-16 are used for encoding. The "%" character for URIs in triples can be used because RDFa tags are not used for a part of a RDF statement. Thus IRIs are usable, too. Because RDFa is embedded in XHTML, the overhead is higher compared to other serialization technologies and also reduces the readability. The basic idea of RDFa is enable an RDFa processor to extract RDF statements from an RDFa enriched HTML document. This is achieved by defining the scope of a certain resource description, for example, using the 'about' attribute (cf. line 10 in Figure 5). Within this scope, triples can now be extracted from links having an additional 'rel' attribute (line 13) or other tags having a 'property attribute' (lines 11 and 14).

JSON-LD. JavaScript Object Notation (JSON) was developed for easy data interchange between applications. JSON, although carrying JavaScript in its name and being a subset of JavaScript, meanwhile became a language independent format which can be used for exchanging all kinds of data structures and is widely supported in different programming languages. Compared to XML, JSON-LD requires less overhead with regard to parsing and serializing. JSON-LD has been developed by the JSON for Linking Data Community Group and been transferred to the RDF Working Group for review, improvement, and publication along the Recommendation track. JSON-LD's design goals are simplicity, compatibility, expressiveness, terseness, zero edits and one-pass processing. As a result, JSON-LD documents are basically standard attribute-value JSON documents with an additional context section (lines 2-7 in Figure 5) establishing mappings to RDF vocabularies. Text in JSON and, thus, also RDF resource identifiers are encoded in Unicode and hence can contain IRIs.

N-Triples

```
1 <http://dbpedia.org/resource/Leipzig> <http://dbpedia.org/property/hasMayor>
2   <http://dbpedia.org/resource/Burkhard_Jung> .
3 <http://dbpedia.org/resource/Leipzig> <http://www.w3.org/2000/01/rdf-schema#label>
4   "Leipzig"@de .
5 <http://dbpedia.org/resource/Leipzig> <http://www.w3.org/2003/01/geo/wgs84_pos#lat>
6   "51.333332"^^<http://www.w3.org/2001/XMLSchema#float> .
```

Turtle

```
1 @prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
2 @prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .
3 @prefix dbp: <http://dbpedia.org/resource/> .
4 @prefix dbpp: <http://dbpedia.org/property/> .
5 @prefix geo: <http://www.w3.org/2003/01/geo/wgs84_pos#> .
6
7 dbp:Leipzig dbpp:hasMayor dbp:Burkhard_Jung ;
8   rdfs:label "Leipzig"@de ;
9   geo:lat "51.333332"^^xsd:float .
```

RDF/XML

```
1 <?xml version="1.0"?>
2 <rdf:RDF xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
3   xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
4   xmlns:dbpp="http://dbpedia.org/property/"
5   xmlns:geo="http://www.w3.org/2003/01/geo/wgs84_pos#">
6   <rdf:Description rdf:about="http://dbpedia.org/resource/Leipzig">
7     <property:hasMayor rdf:resource="http://dbpedia.org/resource/Burkhard_Jung" />
8     <rdfs:label xml:lang="de">Leipzig</rdfs:label>
9     <geo:lat rdf:datatype="http://www.w3.org/2001/XMLSchema#float">51.3333</geo:lat>
10   </rdf:Description>
11 </rdf:RDF>
```

RDfA

```
1 <?xml version="1.0" encoding="UTF-8"?>
2 <!DOCTYPE html PUBLIC "-//W3C//DTD XHTML+RDfA 1.0//EN"
3   "http://www.w3.org/MarkUp/DTD/xhtml-rdfa-1.dtd">
4 <html version="XHTML+RDfA 1.0" xml:lang="en" xmlns="http://www.w3.org/1999/xhtml"
5   xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
6   xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
7   xmlns:dbpp="http://dbpedia.org/property/"
8   xmlns:geo="http://www.w3.org/2003/01/geo/wgs84_pos#">
9   <head><title>Leipzig</title></head>
10   <body about="http://dbpedia.org/resource/Leipzig">
11     <h1 property="rdfs:label" xml:lang="de">Leipzig</h1>
12     <p>Leipzig is a city in Germany. Leipzig's mayor is
13       <a href="Burkhard_Jung" rel="dbpp:hasMayor">Burkhard Jung</a>. It is located
14       at latitude <span property="geo:lat" datatype="xsd:float">51.3333</span>.</p>
15   </body>
16 </html>
```

JSON-LD

```
1 {
2   "@context": {
3     "rdfs": "http://www.w3.org/2000/01/rdf-schema#",
4     "hasMayor": { "@id": "http://dbpedia.org/property/hasMayor", "@type": "@id" },
5     "Person": "http://xmlns.com/foaf/0.1/Person",
6     "lat": "http://www.w3.org/2003/01/geo/wgs84_pos#lat"
7   },
8   "@id": "http://dbpedia.org/resource/Leipzig",
9   "rdfs:label": "Leipzig",
10   "hasMayor": "http://dbpedia.org/resource/Burkhard_Jung",
11   "lat": { "@value": "51.3333", "@type": "http://www.w3.org/2001/XMLSchema#float"
12 }
```

Fig. 5: Different RDF serializations of three triples from Figure 3.

3 Extraction

Information represented in unstructured form or adhering to a different structured representation formalism must be mapped to the RDF data model in order to be used within the Linked Data life-cycle. In this section, we give an overview on some relevant approaches for extracting RDF from unstructured and structured sources.

3.1 From Unstructured Sources

The extraction of structured information from unstructured data sources (especially text) has been a central pillar of *natural language processing* (NLP) and *Information Extraction* (IE) for several decades. With respect to the extraction of RDF data from unstructured data, three sub-disciplines of NLP play a central role: *Named Entity Recognition* (NER) for the extraction of entity labels from text, *Keyword/Keyphrase Extraction* (KE) for the recognition of central topics and *Relationship Extraction* (RE, also called relation mining) for mining the properties which link the entities and keywords described in the data source. A noticeable additional task during the migration of these techniques to Linked Data is the extraction of suitable IRIs for the discovered entities and relations, a requirement that was not needed before. In this section, we give a short overview of approaches that implement the required NLP functionality. Then we present a framework that applies machine learning to boost the quality of the RDF extraction from unstructured data by merging the results of NLP tools.

Named Entity Recognition. The goal of NER is to discover instances of a predefined classes of entities (e.g., persons, locations, organizations) in text. NER tools and frameworks implement a broad spectrum of approaches, which can be subdivided into three main categories: dictionary-based, rule-based, and machine-learning approaches. The first systems for NER implemented dictionary-based approaches, which relied on a list of NEs and tried to identify these in text [199,6]. Following work that showed that these approaches did not perform well for NER tasks such as recognizing proper names [178], rule-based approaches were introduced. These approaches rely on hand-crafted rules [42,189] to recognize NEs. Most rule-based approaches combine dictionary and rule-based algorithms to extend the list of known entities. Nowadays, handcrafted rules for recognizing NEs are usually implemented when no training examples are available for the domain or language to process [141].

When training examples are available, the methods of choice are borrowed from supervised machine learning. Approaches such as Hidden Markov Models [213], Maximum Entropy Models [47] and Conditional Random Fields [57] have been applied to the NER task. Due to scarcity of large training corpora as necessitated by machine learning approaches, semi-supervised [162,140] and unsupervised machine learning approaches [142,53] have also been used for extracting NER from text. [140] gives an exhaustive overview of approaches for NER.

Keyphrase Extraction. Keyphrases/Keywords are multi-word units (MWUs) which capture the main topics of a document. The automatic detection of such MWUs has

been an important task of NLP for decades but due to the very ambiguous definition of what an appropriate keyword should be, current approaches to the extraction of keyphrases still display low F-scores [99]. From the point of view of the Semantic Web, the extraction of keyphrases is a very similar task to that of finding tags for a given document. Several categories of approaches have been adapted to enable KE, of which some originate from research areas such as summarization and information retrieval (IR). Still, according to [98], the majority of the approaches to KE implement combinations of statistical, rule-based or heuristic methods [60,157] on mostly document [128], keyphrase [192] or term cohesion features [161]. [99] gives an overview of current tools for KE.

Relation Extraction. The extraction of relations from unstructured data builds upon work for NER and KE to determine the entities between which relations might exist. Most tools for RE rely on pattern-based approaches. Some early work on pattern extraction relied on supervised machine learning [72]. Yet, such approaches demanded large amount of training data, making them difficult to adapt to new relations. The subsequent generation of approaches to RE aimed at bootstrapping patterns based on a small number of input patterns and instances. For example, [35] presents the Dual Iterative Pattern Relation Expansion (DIPRE) and applies it to the detection of relations between authors and titles of books. This approach relies on a small set of seed patterns to maximize the precision of the patterns for a given relation while minimizing their error rate of the same patterns. Snowball [4] extends DIPRE by a new approach to the generation of seed tuples. Newer approaches aim to either collect redundancy information from the whole Web [160] or Wikipedia [201,208] in an unsupervised manner or to use linguistic analysis [75,156] to harvest generic patterns for relations.

URI Disambiguation. One important problem for the integration of NER tools for Linked Data is the retrieval of IRIs for the entities to be manipulated. In most cases, the URIs can be extracted from generic knowledge bases such as DBpedia [118,112] by comparing the label found in the input data with the `rdfs:label` or `dc:title` of the entities found in the knowledge base. Furthermore, information such as the type of NEs can be used to filter the retrieved IRIs via a comparison of the `rdfs:label` of the `rdf:type` of the URIs with the name of class of the NEs. Still in many cases (e.g., Leipzig, Paris), several entities might bear the same label.

The FOX Framework

Several frameworks have been developed to implement the functionality above for the Data Web including OpenCalais³ and Alchemy⁴. Yet, these tools rely mostly on one approach to perform the different tasks at hand. In this section, we present the FOX (Federated knOwledge eXtraction) framework⁵. FOX differs from the state of the art

³ <http://www.opencalais.com>

⁴ <http://www.alchemyapi.com>

⁵ <http://aksw.org/projects/fox>

by making use of several NER algorithms at once and combining their results by using ensemble learning techniques. FOX’s NED is implemented by the AGDISTIS framework⁶, which combines graph search with the HITS algorithm to detect resources that match a set of strings.

Named Entity Recognition in FOX The basic intuition behind the FOX framework is that while manifold approaches have been devised for the purpose of NER, they mostly rely on one algorithm. However, it is known that each algorithm has intrinsic limitations by virtue of the model that it relies on. For example, polynomial classifiers cannot model decision planes that are non-polynomial. FOX makes use of the diversity of the algorithms available for NER approaches by combining their results within the ensemble learning setting. Formally, the idea is the following: Let $C_1 \dots C_n$ be classifiers that values between 0 and 1 when given a string s and an entity type t as input, i.e., $C_i(s, t) \in [0, 1]$. Then, we can always train a novel classifier C that performs at least as well as any of the C_i on a training dataset D . To this end, we consider the supervised ensemble learning problem where the input data for C is the vector $V = (v_i)_{i=1 \dots n}$ with $v_i = C_i(s, t)$ and the output must be the expected score for (s, t) . Learning the classifier C can be carried by using any supervised machine-learning approach that can deal with non-binary data.

The architecture of FOX consists of *three main layers* as shown in Figure 6. The *machine learning* layer implements interfaces for accommodating ensemble learning techniques such as simple veto algorithms but also neural networks. It consists of *two main modules*. The *training module* allows to load training data so as to enable FOX to learn the best combination of tools and categories for achieving superior recall and precision on the input training data. Depending on the training algorithm used, the user can choose to tune the system for either precision or recall. When using neural networks for example, the user can decide to apply a higher threshold for the output neurons, thus improving the precision but potentially limiting the recall. The *prediction module* allows to run FOX by loading the result of a training session and processing the input data according to the tool-category combination learned during the training phase. Note that the same learning approach can be applied to NER, KE, RE and URI lookup as they all can be modelled as classification tasks. The second layer of FOX is the *controller*, which coordinates the access to the modules that carry out the language processing. The controller is aware of each of the modules in its backend and carries out the initialisation of these modules once FOX is started. Furthermore, it collects the results from the backend modules and invokes the results of a training instance to merge the results of these tools. The final layer of FOX is the *tool layer*, wherein all NLP tools and services integrated in FOX can be found. It is important to notice that the tools per se are not trained during the learning phase of FOX. Rather, we learn of the models already loaded in the tools to allow for the best prediction of named entities in a given domain.

Currently, FOX includes Stanford NER, Illinois NER, Balie and OpenNLP. In an effort to quantify the gain in accuracy of FOX, we integrated the Waikato Environment for Knowledge Analysis (Weka) [74] and the implemented classifiers with default parameters: AdaBoostM1 (ABM1) [61] and Bagging (BG) [32] with J48 [165] as base

⁶ <http://aksw.org/projects/agdistis>

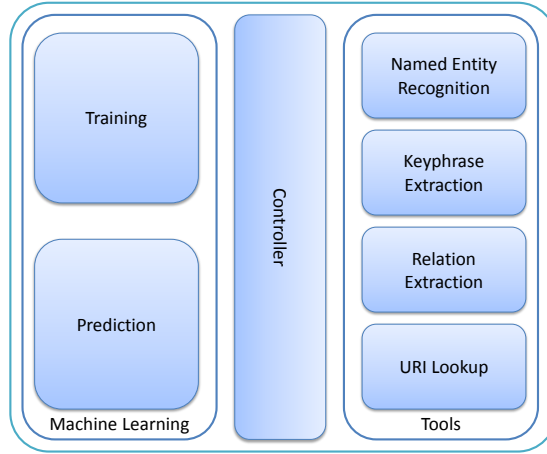


Fig. 6: FOX Architecture

classifier, Decision Table (DT) [101], Functional Trees (FT) [65,105], J48 [165], Logistic Model Trees (LMT) [105], Logistic Regression (Log) [106], Additive Logistic Regression (LogB) [63], Multilayer Perceptron (MLP), Naïve Bayes [96], Random Forest (RF) [33], Support Vector Machine (SVM) [38] and Sequential Minimal Optimization (SMO) [78]. In addition, we used voting at entity level (CVote) and a simple voting (Vote) approach [206] with equal weights for all NER tools. CVote selects the NER tool with the highest prediction performance for each type according to the evaluation and uses that particular tool for the given class. Vote as naive approach combines the results of the NER tools with the Majority Vote Rule [100] and was the baseline ensemble learning technique in our evaluation. In Table 2, the comparison of fifteen different types of classifiers in a ten-fold cross validation over a dataset extracted from online newspapers and containing 150 locations, 139 organizations and 178 persons is shown. Like in previous works [150], these results suggest that multi-layer perceptrons tend to perform well on this task. Interestingly, learning on this small dataset already pushes the overall F-measure of FOX to 95.23% while the best single algorithm achieves 91.01% F-measure.

Named Entity Disambiguation: AGDISTIS The goal of AGDISTIS is to detect correct resources from a KB K for a vector N of n a-priori determined named entities N_1, \dots, N_n extracted from a certain input text T . In general, several resources from a given knowledge base K can be considered as candidate resources for a given entity N_i . For the sake of simplicity and without loss of generality, we will assume that each of the entities can be mapped to m distinct candidate resources. Let C be the matrix which contains all candidate-entity mappings for a given set of entities. The entry C_{ij} stands for the j^{th} candidate resource for the i^{th} named entity. Let μ be a family of functions which maps each entity N_i to exactly one candidate C_{ij} . We call such functions *assign-*

Table 2: The News data set.

S	<i>Recall</i>	<i>Precision</i>	F_1	<i>error</i>	<i>MCC</i>
MLP	95.19	95.28	95.23	0.32	0.951
RF	95.15	95.28	95.21	0.32	0.951
ABM1.J48	94.82	95.18	95.00	0.33	0.948
SVM	94.86	95.09	94.97	0.33	0.948
BG.J48	94.76	94.93	94.84	0.34	0.947
J48	94.78	94.98	94.88	0.34	0.947
DT	94.63	94.95	94.79	0.34	0.946
LMT	94.68	94.95	94.82	0.34	0.946
FT	94.30	95.15	94.72	0.35	0.945
LogB	93.54	95.37	94.44	0.37	0.943
Log	94.05	94.75	94.40	0.37	0.942
SMO	94.01	94.37	94.19	0.39	0.940
Naïve Bayes	94.61	92.64	93.60	0.42	0.934
Stanford	92.36	91.01	91.68	0.53	0.914
CVote	92.02	90.84	91.42	0.54	0.911
Vote	89.98	82.97	85.92	0.94	0.857
Illinois	82.79	87.35	84.95	0.92	0.845
OpenNLP	71.42	90.47	79.57	1.13	0.797
Balie	77.68	82.05	79.8	1.21	0.792

ments. The output of an assignment is a vector of resources of length $|N|$ that is such that the i^{th} entry of the vector maps with N_i .

Let ψ be a function which computes the similarity between an assignment $\mu(C, N)$ and the vector of named entities N . The *coherence* function ϕ calculates the similarity of the knowledge base K and an assignment μ [169] to ensure the topical consistency of μ . The coherence function ϕ is implemented by the HITS algorithm, which calculates the most pertinent entities while the similarity function ψ is, e.g., string similarity. Given this formal model, the goal is to find the assignment μ^* with

$$\mu^* = \arg \max_{\mu} (\psi(\mu(C, N), N) + \phi(\mu(C, N), K)).$$

The formulation of the problem given above has been proven to be NP-hard [46]. Thus, for the sake of scalability, AGDISTIS computes an approximation μ^+ by using HITS, a fast graph algorithm which runs with an upper bound of $\Theta(k \cdot |V|^2)$ with k the number of iterations and $|V|$ the number of nodes in the graph. Furthermore, using HITS leverages 1) scalability, 2) well-researched behaviour and 3) the ability to explicate semantic authority.

For each named entity, candidates resources from the input knowledge base can be detected by the labels of the resources in the knowledge base with the named entity. Here, resources such as surface forms can also be used. Given a set of candidate nodes, we begin the computation of the optimal assignment by constructing a disambiguation graph G_d with search depth d . To this end, we regard the input knowledge base as a directed graph $G_K = (V, E)$ where the vertices V are resources of K , the edges E are properties of K and $x, y \in V, (x, y) \in E \Leftrightarrow \exists p : (x, p, y)$ is an RDF triple in K . Given

the set of candidates C , we begin by building an initial graph $G_0 = (V_0, E_0)$ where V_0 is the set of all resources in C and $E_0 = \emptyset$. Starting with G_0 we extend the graph in a breadth-first search manner. Therefore, we define the extension of a graph $G_i = (V_i, E_i)$ to a graph $\rho(G_i) = G_{i+1} = (V_{i+1}, E_{i+1})$ where $i = 0, \dots, d$ as follows:

$$V_{i+1} = V_i \cup \{y : \exists x \in V_i \wedge (x, y) \in E\} \quad (1)$$

$$E_{i+1} = \{(x, y) \in E : x, y \in V_{i+1}\} \quad (2)$$

We iterate the ρ operator d times on the input graph G_0 to compute the initial disambiguation graph G_d .

After constructing the disambiguation graph G_d we need to identify the correct candidate node for a given named entity. Using the graph-based HITS algorithm we calculate authoritative values x_a, y_a and hub values x_h, y_h for all $x, y \in V_d$. We initialize the authoritative and hub values as follows:

$$\forall x \in V_d, x_a = x_h = \frac{1}{|V_d|}. \quad (3)$$

Afterwards, we iterate k times the following equations:

$$x_a \leftarrow \sum_{(y,x) \in E_d} y_h, \quad y_h \leftarrow \sum_{(y,x) \in E_d} x_a. \quad (4)$$

We choose $k = 20$ iterations, which suffices to achieve convergence in general. Afterwards we identify the most authoritative candidate C_{ij} among the set of candidates C_i as correct disambiguation for a given named entity N_i . When using DBpedia as KB and C_{ij} is a redirect AGDISTIS uses the target resource. As can be seen, we calculate μ^+ solely by using algorithms with a polynomial time complexity. The evaluation of AGDISTIS shown in Table 3 on datasets by [45] shows clearly that AGDISTIS outperforms the state of the art in 3 out of 4 cases.

3.2 From Structured Sources

Structured knowledge, e.g. relational databases and XML, is the backbone of many (web) applications. Extracting or converting this knowledge to RDF is a long-standing research goal in the Semantic Web community. A conversion to RDF allows to integrate the data with other sources and perform queries over it. In this lecture, we focus on the conversion of relational databases to RDF (see Figure 7). In the first part, we summarize material from a recent relational database to RDF (RDB2RDF) project report. After that, we describe the mapping language R2RML, which is a language for expressing database to RDF conversion mappings. While we focus on relational data, we also want to note that extraction from CSV files is also highly important as illustrated in use cases in the financial [127] and health sector [211].

Triplify and RDB2RDF Survey report The table displayed in Figure 8 is taken from the Triplify WWW paper [8]. The survey report [177] furthermore contained a chart (see Figure 9) showing the reference framework for classifying the approaches and an extensive table classifying the approaches (see Figure 10). Another recent survey is [187].

The following criteria can be extracted:

Dataset	Approach	F ₁	Precision	Recall
AIDA/CO-NLL-TestB	TagMe 2	0.565	0.58	0.551
	DBPedia Spotlight	0.341	0.308	0.384
	AGDISTIS	0.596	0.642	0.556
AQUAINT	TagMe 2	0.457	0.412	0.514
	DBPedia Spotlight	0.26	0.178	0.48
	AGDISTIS	0.547	0.777	0.422
IITB	TagMe 2	0.408	0.416	0.4
	DBPedia Spotlight	0.46	0.434	0.489
	AGDISTIS	0.31	0.646	0.204
MSNBC	TagMe 2	0.466	0.431	0.508
	DBPedia Spotlight	0.331	0.317	0.347
	AGDISTIS	0.761	0.796	0.729

Table 3: Performance of AGDISTIS, DBpedia Spotlight and TagMe2 on four different datasets using micro F-measure

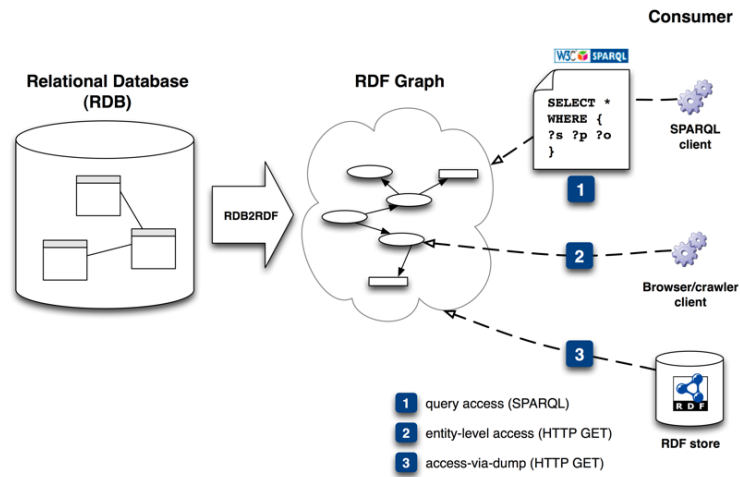


Fig. 7: Illustration of RDB to RDF conversion.

Source: <http://www.w3.org/2001/sw/rdb2rdf/use-cases/>.

Approach	Automation (a)	Domain or database semantics-driven (b)	Access paradigm (c)	Mapping language (d)	Domain reliance (e)
Dartgrid [17]	Manual	Domain	SPARQL	Visual Tool	dependent
Hu et al. [11]	Auto	Both	ETL	intern	dependent
Tirmizi et al. [16]	Auto	DB	ETL	FOL	general
Li et al. [12]	Semi	DB	ETL	n/a	general
DB2OWL[10]	Semi	DB	SPARQL	R2O	general/dependent
RDBToOnto [6]	Semi	DB+M	ETL	Visual Tool	general
Sahoo et al. [15]	Manual	Domain	ETL	XSLT	dependent
R2O[13]	Manual	DB+M	SPARQL	R2O	dependent
D2RQ[4]	Auto	DB+M	LD, SPARQL	D2RQ	general
Virtuoso RDF View [5, 9]	Semi	DB+M	SPARQL	own	general
Triplify	Manual	Domain	LD	SQL	general

Table 4: An integrated overview of mapping approaches. Criteria for classification were merged, some removed, fields were completed, when missing. DB+M means that the semi-automatic approach can later be customized manually

Fig. 8: Table comparing relevant approaches from [8].

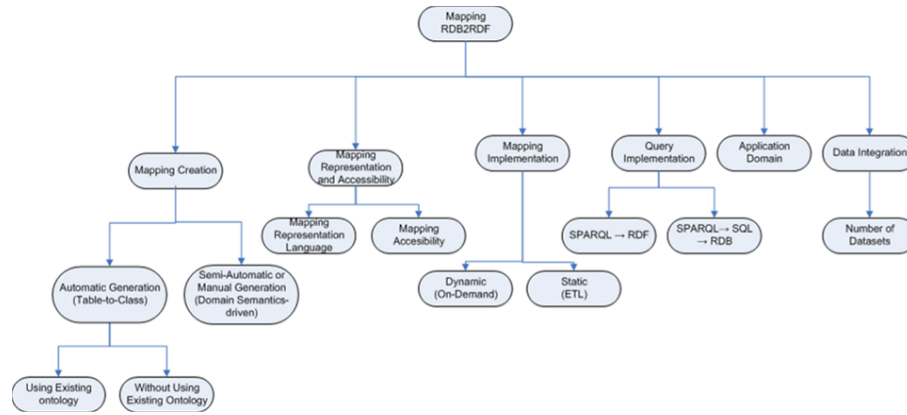


Fig. 9: Reference framework by [177].

Automation Degree. Degree of mapping creation automation.

Values: Manual, Automatic, Semi-Automatic.

Domain or Database Semantics Driven. Some approaches are tailored to model a domain, sometimes with the help of existing ontologies, while others attempt to extract domain information primarily from the given database schema with few other resources used (domain or database semantics-driven). The latter often results in a table-to-class, column-to-predicate mapping. Some approaches also use a (semi) automatic approach based on the database, but allow manual customization to model domain semantics.

Values: Domain, DB (database), DB+M (database and later manual customisation), Both (Domain and DB)

Access Paradigm. Resulting access paradigm (ETL [extract transform load], Linked Data, SPARQL access). Note that the access paradigm also determines whether the resulting RDF model updates automatically. ETL means a one time conversion, while Linked Data and SPARQL always process queries versus the original database.

Values: SPARQL, ETL, LD

PROJECTS	MAPPING CREATION	MAPPING REPRESENTATION AND ACCESSIBILITY		MAPPING IMPLEMENTATION	QUERY IMPLEMENTATION	APPLICATION DOMAIN	DATA INTEGRATION	
	Automatic (Table-to-Class) or Manual/Semi-Automatic (Domain Semantics-driven)	Representation Language	Mapping Access	Static (ETL) or Dynamic	SPARQL \rightarrow RDF or SPARQL \rightarrow SQL \rightarrow RDB		Yes/No	Number of Datasets
1. Hu et al, 2007	Automatic (with use of existing ontology)	First Order Logic formulas or Horn Clauses	Files	None Specified	None Specified	Genetic	Enables (through contextual mappings)	Potentially Multiple
2. Kashyap et al, 2007	Manual/Semi-Automatic (Domain Semantics-driven)	Mediator Framework Classes	Mapping mediator	Dynamic	SPARQL \rightarrow SQL \rightarrow RDB	Life Sciences	Enables	Potentially Multiple
3. DB2OWL (Cullot et al, 2007)	Automatic (Table to Class)	R2O language	R2O mapping document		SPARQL \rightarrow SQL \rightarrow RDB	Genetic	Enables	Potentially Multiple
4. Timizi et al 2008	Automatic (Table to Class, SQL-DDL to RDF)	First Order Logic	None specified	Static	None Specified	Genetic	No	None
5. SOAM (Li et al, 2005)	Automatic (Table to Class) with user input	Logic Rules	Implemented as part of system	Static	Potentially SPARQL (on generated populated ontology)	Genetic (Case Study: Economics)	No	None
6. Sahoo et al, 2008	Manual/Semi-Automatic (Domain Semantics-driven)	XPath expressions	XSLT document	Static	SPARQL	Life Sciences	Yes	Test included five (Gene, Biological Pathway)
7. Byrne, 2008	Manual/Semi-Automatic (Domain Semantics-driven)	SKOS vocabulary	RDF document	Static	SPARQL	Cultural Heritage	No	None
8. Green et al, 2008	Manual/Semi-Automatic (Domain Semantics-driven)	D2RQ language	D2RQ mapping file	Dynamic	SPARQL \rightarrow SQL \rightarrow RDB	Ordnance Survey	Yes	Multiple
9. Virtuoso RDF View (Blakeley, 2007)	Both (user-specified)	SPASQL-based Meta Schema Language	Quad Storage	Both	Both	Genetic	Enables	Potentially Multiple
10. D2RQ (Bizer et al, 2007)	Both (user-specified)	D2RQ language	D2RQ mapping file	Both	Both	Genetic	Enables	Potentially Multiple
11. R2O (Barrasa et al, 2006)	Both (user-specified)	R2O language	R2O mapping document	Both	Both	Genetic	Enables	Potentially Multiple
12. Dartgnd (Wu et al, 2006)	Automatic (Table to Class)	XML File	Visualized Mapping tool	Dynamic	SPARQL \rightarrow SQL \rightarrow RDB (Provide search and query interface)	Life Science (Traditional Chinese Medicine, TCM)	Yes	Test included databases for herb, compound formulas, disease, drug, TCM treatment.
13. RDBtoOnto (Cerbah, 2008)	Automatic (Table to Class, allows user intervention)	Constraint rules	Not explicitly stored	Static	Potentially SPARQL (on generated populated ontology)	Genetic	No	None
14. Asic Tools	Automatic (Table to Class)	OWL Full based language	File based	Both	SPARQL \rightarrow SQL \rightarrow RDB	Genetic	Enables	Potentially Multiple

Fig. 10: Comparison of approaches from [177].

Mapping Language. The used mapping language as an important factor for reusability and initial learning cost.

Values: Visual Tool, intern (internal self-designed language), FOL, n/a (no information available), R2O, XSLT, D2RQ, proprietary, SQL

Domain reliance. Domain reliance (general or domain-dependent): requiring a pre-defined ontology is a clear indicator of domain dependency.

Values: Dependent, General

Type. Although not used in the table the paper discusses four different classes:

Values: Alignment, Database Mining, Integration, Languages/Servers

R2RML - RDB to RDF Mapping Language The R2RML W3C recommendation⁷ specifies an RDF notation for mapping relational tables, views or queries into RDF. The primary area of applicability of this is extracting RDF from relational databases, but in special cases R2RML could lend itself to on-the-fly translation of SPARQL into SQL or to converting RDF data to a relational form. The latter application is not the primary intended use of R2RML but may be desirable for importing linked data into relational stores. This is possible if the constituent mappings and underlying SQL objects constitute updateable views in the SQL sense.

Data integration is often mentioned as a motivating use case for the adoption of RDF. This integration will very often be between relational databases which have logical entities in common, each with its local schema and identifiers. Thus, we expect to see relational to RDF mapping use cases involving the possibility of a triple coming from multiple sources. This does not present any problem if RDF is being extracted but does lead to complications if SPARQL queries are mapped into SQL. In specific, one will end up with potentially very long queries consisting of joins of unions. Most of the joins between terms of the unions will often be provably empty and can thus be optimized away. This capability however requires the mapping language to be able to express metadata about mappings, i.e. that IRIs coming from one place are always disjoint from IRIs coming from another place. Without such metadata optimizing SPARQL to SQL translation is not possible, which will significantly limit the possibility of querying collections of SQL databases through a SPARQL end point without ETL-ing the mapped RDF into an RDF store.

RDF is emerging as a format for interoperable data publishing. This does not entail that RDF were preferable as a data warehousing model. Besides, for large warehouses, RDF is not cost competitive with relational technology, even though projects such as LOD2 and LDBC expect to narrow this gap (see, e.g., [137,138] for recent SPARQL benchmarks). Thus it follows that on the fly mapping of SPARQL to SQL will be important. Regardless of the relative cost or performance of relational or RDF technology, it is not a feasible proposition to convert relational warehouses to RDF in general, rather existing investments must be protected and reused. Due to these reasons, R2RML will have to evolve in the direction of facilitating querying of federated relational resources.

⁷ <http://www.w3.org/TR/r2rml/>

Supervised Extraction Example: Sparqlify The challenges encountered with large scale relational data sources LinkedGeoData [12,188] indicate that ETL style approaches based on the conversion of all underlying data to RDF have severe deficiencies. For instance, the RDF conversion process is very time consuming for large-scale, crowdsourced data. Furthermore, changes in data modelling require many changes in the extracted RDF data or the creation of a completely new dump. In summary, the ETL approach is not sufficiently flexible for very large and frequently changing data. It seems preferable to establish virtual RDF views over the existing relational database. In contrast to other tools, such as *D2R* and *Virtuoso RDF views*, Sparqlify converts each SPARQL query to a single SQL query. This allows all optimisations of the underlying database to be applied and can lead to better scalability.

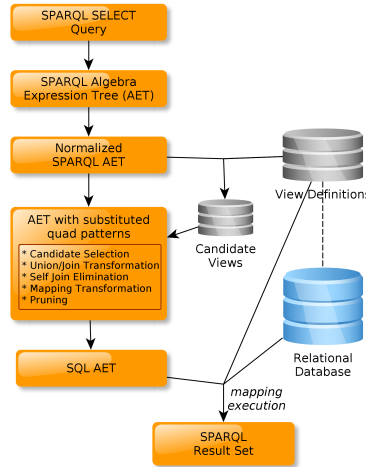


Fig. 11: The Sparqlify concepts and query rewriting workflow.

Figure 11 shows the query rewriting workflow in Sparqlify. The rationale of Sparqlify is to leave the schema of the underlying relational database schema unmodified and define RDF views over it. SPARQL queries can then be written against those views, which are expressed in the *Sparqlify-ML* (mapping language). Sparqlify-ML is easy to learn for users, who are experienced in SPARQL and SQL and more compact than other syntactic variants such as R2RML. The left part of Figure 11 shows all steps, which are performed to answer a query. First, the query is converted into an algebra expression. This expression is subsequently converted to a normal form. Given the query patterns, relevant Sparqlify-ML views need to be detected. After this is done, the algebra expression is rewritten to include those relevant views. In a next step, optimisations on the algebra expression are performed to improve efficiency. Finally, this algebra expression can be transformed to an SQL algebra expression. For accomplishing this, we define a general relational algebra for RDB-to-RDF mappings. The SQL query, which was obtained, is executed against the relational database. Using the defined mappings, the

SQL result set returned by the relational database can be converted to a SPARQL result set.

All of the above steps are explained in detail throughout the next sections. The main contribution of the Sparqlify project is a formalization, which goes beyond previous work by being capable to push the complete query execution using a single SQL query into the DBMS.

4 Authoring with Semantic Wikis

Semantic Wikis are an extension to conventional, text-based Wikis. While in conventional Wikis pages are stored as blocks of text using a special Wiki markup for structuring the display of the text and adding links to other pages, semantic Wikis aim at adding rich structure to the information itself. To this end, two initially orthogonal approaches have been used: a) extending the markup language to allow semantic annotations and links with meaning or b) building the Wiki software directly with structured information in mind. Nowadays, both approaches have somewhat converged, for instance Semantic MediaWiki [103] also provides forms for entering structured data (see Figure 12). Characteristics of both approaches are summarized in Table 4 for the two prototypical representatives of both approaches, i.e. Semantic MediaWiki and OntoWiki.

	Semantic MediaWiki	OntoWiki
<i>Managed entities</i>	Articles	Resources
<i>Editing</i>	Wiki markup	Forms
<i>Atomic element</i>	Text blob	Statement

Table 4: Conceptual differences between Semantic MediaWiki and OntoWiki.

Extending Wikis with Semantic Markup. The benefit of a Wiki system comes from the amount of interlinking between Wiki pages. Those links clearly state a relationship between the linked-to and the linking page. However, in conventional Wiki systems this relationship cannot be made explicit. Semantic Wiki systems therefore add a means to specify typed relations by extending the Wiki markup with semantic (i.e. typed) links. Once in place, those links form a knowledge base underlying the Wiki which can be used to improve search, browsing or automatically generated lists and category pages. Examples of approaches for extending Wikis with semantic markup can be found in [103,180,15,159,186]. They represent a straightforward combination of existing Wiki systems and the Semantic Web knowledge representation paradigms. Yet, we see the following obstacles:

Usability: The main advantage of Wiki systems is their unbeatable usability. Adding more and more syntactic possibilities counteracts ease of use for editors.

Redundancy: To allow the answering of real-time queries to the knowledge base, statements have to be additionally kept in a triple store. This introduces a redundancy, which complicates the implementation.

Evolution: As a result of storing information in both Wiki texts and triple store, supporting evolution of knowledge is difficult.

Wikis for Editing Structured Data. In contrast to text-based systems, Wikis for structured data – also called Data Wikis – are built on a structured model of the data being

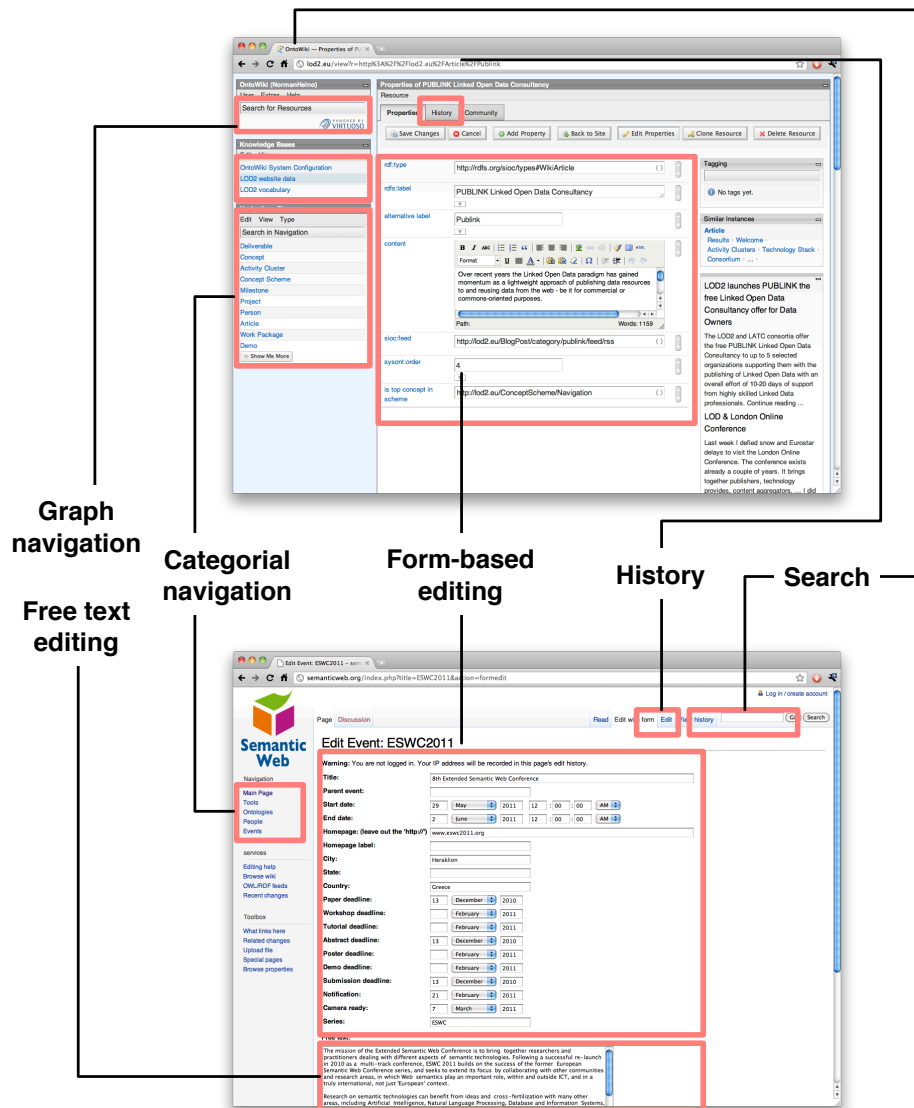


Fig. 12: Comparison of Semantic MediaWiki and OntoWiki GUI building blocks.

edited. The Wiki software can be used to add instances according to the schema or (in some systems) edit the schema itself. One of those systems is OntoWiki⁸ [9] which bases its data model on RDF. This way, both schema and instance data are represented using the same low-level model (i.e. statements) and can therefore be handled identically by the Wiki.

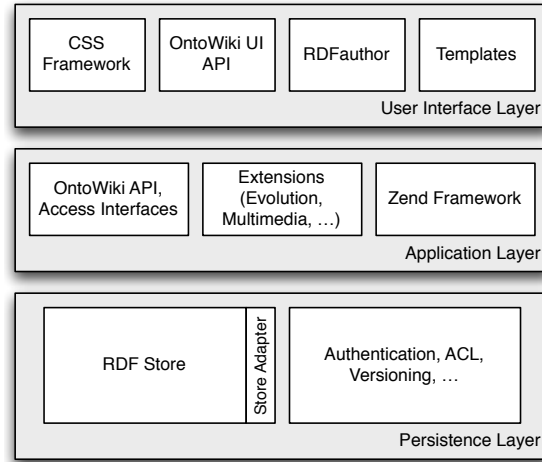


Fig. 13: Overview of OntoWiki’s architecture with extension API and Zend web framework (modified according to [80]).

4.1 OntoWiki - a Semantic Data Wiki

OntoWiki started as an RDF-based data wiki with emphasis on collaboration but has meanwhile evolved into a comprehensive framework for developing Semantic Web applications [80]. This involved not only the development of a sophisticated extension interface allowing for a wide range of customizations but also the addition of several access and consumption interfaces allowing OntoWiki installations to play both a provider and a consumer role in the emerging Web of Data.

OntoWiki is inspired by classical Wiki systems, its design, however, (as mentioned above) is independent and complementary to conventional Wiki technologies. In contrast to other semantic Wiki approaches, in OntoWiki text editing and knowledge engineering (i. e. working with structured knowledge bases) are not mixed. Instead, OntoWiki directly applies the Wiki paradigm of “making it easy to correct mistakes, rather than making it hard to make them” [121] to collaborative management of structured knowledge. This paradigm is achieved by interpreting knowledge bases as *information maps* where every node is represented visually and interlinked to related resources.

⁸ Available at: <http://ontowiki.net>

Furthermore, it is possible to enhance the knowledge schema gradually as well as the related instance data agreeing on it. As a result, the following requirements and corresponding features characterize OntoWiki:

Intuitive display and editing of instance data should be provided in generic ways, yet enabling domain-specific presentation of knowledge.

Semantic views allow the generation of different views and aggregations of the knowledge base.

Versioning and evolution provides the opportunity to track, review and roll-back changes selectively.

Semantic search facilitates easy-to-use full-text searches on all literal data, search results can be filtered and sorted (using semantic relations).

Community support enables discussions about small information chunks. Users are encouraged to vote about distinct facts or prospective changes.

Online statistics interactively measures the popularity of content and activity of users.

Semantic syndication supports the distribution of information and their integration into desktop applications.

OntoWiki enables the easy creation of highly structured content by distributed communities. The following points summarize some limitations and weaknesses of OntoWiki and thus characterize the application domain:

Environment: OntoWiki is a Web application and presumes all collaborators to work in a Web environment, possibly distributed.

Usage Scenario: OntoWiki focuses on knowledge engineering projects where a single, precise usage scenario is either initially (yet) unknown or not (easily) definable.

Reasoning: Application of reasoning services was (initially) not the primary focus.

4.2 Generic and Domain-specific Views

OntoWiki can be used as a tool for presenting, authoring and managing knowledge bases adhering to the RDF data model. As such, it provides generic methods and views, independent of the domain concerned. Two generic views included in OntoWiki are the resource view and the list view. While the former is generally used for displaying all known information about a resource, the latter can present a set of resources, typically instances of a certain concept. That concept must not necessarily be explicitly defined as `rdfs:Class` or `owl:Class` in the knowledge base. Via its faceted browsing, OntoWiki allows the construction of complex concept definitions, with a pre-defined class as a starting point by means of property value restrictions. These two views are sufficient for browsing and editing all information contained in a knowledge base in a generic way. For domain-specific use cases, OntoWiki provides an easy-to-use extension interface that enables the integration of custom components. By providing such a custom view, it is even possible to hide completely the fact that an RDF knowledge base is worked on. This permits OntoWiki to be used as a data-entry frontend for users with a less profound knowledge of Semantic Web technologies.

4.3 Workflow

With the use of RDFS [34] and OWL [163] as ontology languages, resource definition is divisible into different layers: a terminology box for conceptual information (i. e. classes and properties) and an assertion box for entities using the concepts defined (i. e. instances). There are characteristics of RDF which, for end users, are not easy to comprehend (e. g. *classes* can be defined as *instances* of `owl:Class`). OntoWiki's user interface, therefore, provides elements for these two layers, simultaneously increasing usability and improving a user's comprehension for the structure of the data. After starting and logging in into OntoWiki with registered user credentials, it is possible to select one of the existing ontologies. The user is then presented with general information about the ontology (i. e. all statements expressed about the knowledge base as a resource) and a list of defined classes, as part of the conceptual layer.

After starting and logging in into OntoWiki with registered user credentials, it is possible to select one of the existing knowledge bases. The user is then presented with general information about the ontology (i. e. all statements expressed about the knowledge base as a resource) and a list of defined classes, as part of the conceptual layer. By selecting one of these classes, the user obtains a list of the class' instances. OntoWiki applies basic `rdfs:subClassOf` reasoning automatically. After selecting an instance from the list – or alternatively creating a new one – it is possible to manage (i. e. insert, edit and update) information in the details view. OntoWiki focuses primarily on the assertion layer, but also provides ways to manage resources on the conceptual layer. By enabling the visualization of schema elements, called *System Classes* in the OntoWiki nomenclature, conceptual resources can be managed in a similar fashion as instance data.

4.4 Authoring

Semantic content in OntoWiki is represented as resource descriptions. Following the RDF data model representing one of the foundations of the Semantic Web vision, resource descriptions are represented (at the lowest level) in the form of *statements*. Each of these statements (or triples) consist of a *subject* which identifies a resource as well as a *predicate* and an *object* which together represent data about said resource in a fashion reminiscent of key-value pairs. By means of RDFa [3], these statements are retained in the HTML view (i.e. user interface) part and are thus accessible to client-side techniques like JavaScript.

Authoring of such content is based on said client-side representation by employing the RDFauthor approach [191]: views are declared in terms of the model language (RDF) which allows the underlying model be restored. Based on this model, a user interface can be generated with the model being providing all the domain knowledge required to do so. The RDFauthor system provides an extensible set of authoring widgets specialized for certain editing tasks. RDFauthor was also extended by adding capabilities for automatically translating literal object values between different languages. Since the semantic context is known to the system, these translation functionality can be bound to arbitrary characteristics of the data (e. g. to a certain property or a missing language).

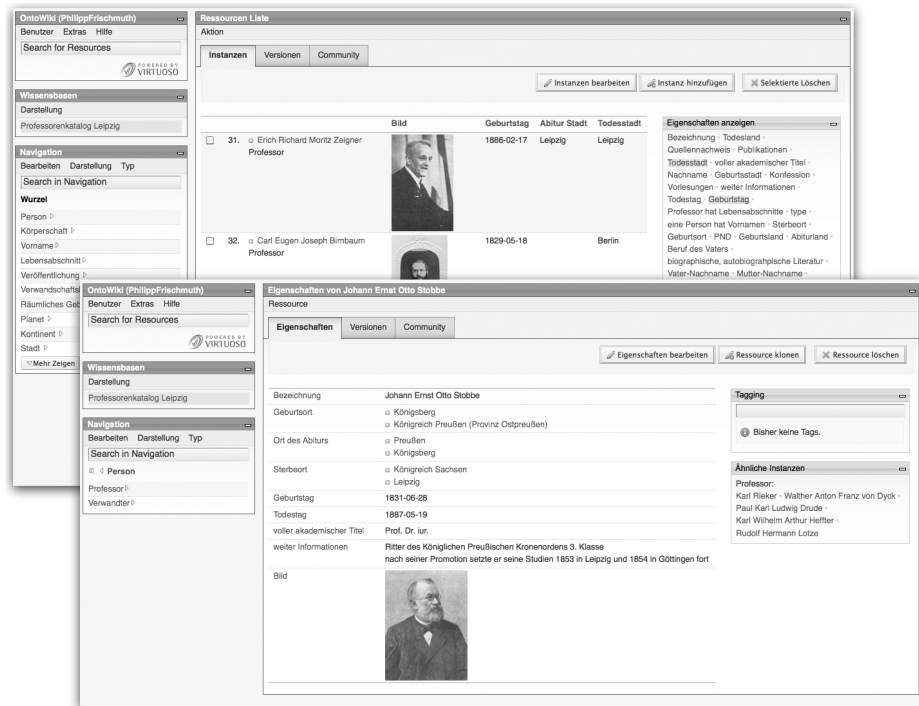


Fig. 14: OntoWiki views: (background) A tabular list view, which contains a filtered list of resources highlighting some specific properties of those resources and (foreground) a resource view which allows to tag and comment a specific resource as well as editing all property values.

Versioning & Evolution. As outlined in the wiki principles, keeping track of all changes is an important task in order to encourage user participation. OntoWiki applies this concept to RDF-based knowledge engineering in that all changes are tracked on the statement level [10]. These low-level changes can be grouped to reflect application- and domain-specific tasks involving modifications to several statements as a single versioned item. Provenance information as well as other metadata (such as time, user or context) of a particular changeset can be attached to each individual changeset. All changes on the knowledge base can be easily reviewed and rolled-back if needed. The loosely typed data model of RDF encourages continuous evolution and refinement of knowledge bases. With *EvoPat*, OntoWiki supports this in a declarative, pattern-based manner (see section on evolution).

4.5 Access Interfaces

In addition to human-targeted graphical user interfaces, OntoWiki supports a number of machine-accessible data interfaces. These are based on established Semantic Web

standards like SPARQL or accepted best practices like publication and consumption of Linked Data.

SPARQL Endpoint. The SPARQL recommendation not only defines a query language for RDF but also a protocol for sending queries to and receiving results from remote endpoints⁹. OntoWiki implements this specification, allowing all resources managed in an OntoWiki be queried over the Web. In fact, the aforementioned RDFauthor authoring interface makes use of SPARQL to query for additional schema-related information, treating OntoWiki as a remote endpoint in that case.

Linked Data. Each OntoWiki installation can be part of the emerging Linked Data Web. According to the Linked Data publication principles (cf. section 2), OntoWiki makes all resources accessible by its IRI (provided, the resource's IRI is in the same namespace as the OntoWiki instance). Furthermore, for each resource used in OntoWiki additional triples can be fetched if the resource is de-referenceable.

Semantic Pingback. Pingback is an established notification system that gained wide popularity in the blogosphere. With Semantic Pingback [190], OntoWiki adapts this idea to Linked Data providing a *notification mechanism* for resource usage. If a Pingback-enabled resource is mentioned (i. e. linked to) by another party, its pingback server is notified of the usage. Provided, the Semantic Pingback extension is enabled all resources used in OntoWiki are pinged automatically and all resources defined in OntoWiki are Pingback-enabled.

4.6 Exploration Interfaces

For exploring semantic content, OntoWiki provides several exploration interfaces that range from generic views over search interfaces to sophisticated querying capabilities for more RDF-knowledgeable users. The subsequent paragraphs give an overview of each of them.

Knowledge base as an information map. The compromise between, on the one hand, providing a generic user interface for arbitrary RDF knowledge bases and, on the other hand, aiming at being as intuitive as possible is tackled by regarding knowledge bases as *information maps*. Each node at the information map, i. e. RDF resource, is represented as a Web accessible page and interlinked to related digital resources. These Web pages representing nodes in the information map are divided into three parts: a left sidebar, a main content section and a right sidebar. The left sidebar offers the selection of content to display in the main content section. Selection opportunities include the set of available knowledge bases, a hierarchical browser and a full-text search.

Full-text search. The full-text search makes use of special indexes (mapped to proprietary extensions to the SPARQL syntax) if the underlying knowledge store provides this feature, else, plain SPARQL string matching is used. In both cases, the resulting

⁹ <http://www.w3.org/TR/rdf-sparql-protocol/>

SPARQL query is stored as an object which can later be modified (e. g. have its filter clauses refined). Thus, full-text search is seamlessly integrated with faceted browsing (see below).

Content specific browsing interfaces. For domain-specific use cases, OntoWiki provides an easy-to-use extension interface that enables the integration of custom components. By providing such a custom view, it is even possible to hide completely the fact that an RDF knowledge base is worked on. This permits OntoWiki to be used as a data-entry frontend for users with a less profound knowledge of Semantic Web technologies.

Faceted-browsing. Via its faceted browsing, OntoWiki allows the construction of complex concept definitions, with a pre-defined class as a starting point by means of property value restrictions. These two views are sufficient for browsing and editing all information contained in a knowledge base in a generic way.

Query-builder. OntoWiki serves as a SPARQL endpoint, however, it quickly turned out that formulating SPARQL queries is too tedious for end users. In order to simplify the creation of queries, we developed the *Visual Query Builder*¹⁰ (VQB) as an OntoWiki extension, which is implemented in JavaScript and communicates with the triple store using the SPARQL language and protocol. VQB allows to visually create queries to the stored knowledge base and supports domain experts with an intuitive visual representation of query and data. Developed queries can be stored and added via drag-and-drop to the current query. This enables the reuse of existing queries as building blocks for more complex ones.

4.7 Applications

Catalogous Professorum. The World Wide Web, as an ubiquitous medium for publication and exchange, already significantly influenced the way historians work: the online availability of catalogs and bibliographies allows to efficiently search for content relevant for a certain investigation; the increasing digitization of works from historical archives and libraries, in addition, enables historians to directly access historical sources remotely. The capabilities of the Web as a medium for collaboration, however, are only starting to be explored. Many, historical questions can only be answered by combining information from different sources, from different researchers and organizations. Also, after original sources are analyzed, the derived information is often much richer, than can be captured by simple keyword indexing. These factors pave the way for the successful application of knowledge engineering techniques in historical research communities.

In [171] we report about the application of an adaptive, semantics-based knowledge engineering approach using OntoWiki for the development of a prosopographical knowledge base. In prosopographical research, historians analyze common characteristics of historical groups by studying statistically relevant quantities of individual biographies. Untraceable periods of biographies can be determined on the basis of such

¹⁰ <http://aksw.org/Projects/OntoWiki/Extension/VQB>

accomplished analyses in combination with statistically examinations as well as patterns of relationships between individuals and their activities.

In our case, researchers from the historical seminar at Universität Leipzig aimed at creating a prosopographical knowledge base about the life and work of professors in the 600 years history of Universität Leipzig ranging from the year 1409 till 2009 - the *Catalogus Professorum Lipsiensis* (CPL). In order to enable historians to collect, structure and publish this prosopographical knowledge an ontological knowledge model was developed and incrementally refined over a period of three years. The community of historians working on the project was enabled to add information to the knowledge base using an adapted version of OntoWiki. For the general public, a simplified user interface¹¹ is dynamically generated based on the content of the knowledge base. For access and exploration of the knowledge base by other historians a number of access interfaces was developed and deployed, such as a graphical SPARQL query builder, a relationship finder and plain RDF and Linked Data interfaces. As a result, a group of 10 historians supported by a much larger group of volunteers and external contributors collected information about 1,300 professors, 10,000 associated periods of life, 400 institutions and many more related entities.

The benefits of the developed knowledge engineering platform for historians are twofold: Firstly, the collaboration between the participating historians has significantly improved: The ontological structuring helped to quickly establish a common understanding of the domain. Collaborators within the project, peers in the historic community as well as the general public were enabled to directly observe the progress, thus facilitating peer-review, feedback and giving direct benefits to the contributors. Secondly, the ontological representation of the knowledge facilitated original historical investigations, such as historical social network analysis, professor appointment analysis (e.g. with regard to the influence of cousin-hood or political influence) or the relation between religion and university. The use of the developed model and knowledge engineering techniques is easily transferable to other prosopographical research projects and with adaptations to the ontology model to other historical research in general. In the long term, the use of collaborative knowledge engineering in historian research communities can facilitate the transition from largely individual-driven research (where one historian investigates a certain research question solitarily) to more community-oriented research (where many participants contribute pieces of information in order to enlighten a larger research question). Also, this will improve the reusability of the results of historic research, since knowledge represented in structured ways can be used for previously not anticipated research questions.

OntoWiki Mobile. As comparatively powerful mobile computing devices are becoming more common, mobile web applications have started gaining in popularity. An important feature of these applications is their ability to provide *offline functionality* with local updates for later synchronization with a web server. The key problem here is the reconciliation, i.e. the problem of potentially *conflicting updates* from *disconnected clients*. Another problem current mobile application developers face is the plethora of mobile application development platforms as well as the incompatibilities between

¹¹ Available at: <http://www.uni-leipzig.de/unigeschichte/professorenkatalog/>

them. *Android* (Google), *iOS* (Apple), *Blackberry OS* (RIM), *WebOS* (HP/Palm), *Symbian* (Nokia) are popular and currently widely deployed platforms, with many more proprietary ones being available as well. As a consequence of this fragmentation, realizing a special purpose application, which works with many or all of these platforms is extremely time consuming and inefficient due to the large amount of duplicate work required.

The W3C addressed this problem, by enriching HTML in its 5th revision with access interfaces to local storage (beyond simple cookies) as well as a number of devices and sensors commonly found on mobile devices (e. g. GPS, camera, compass etc.). We argue, that in combination with semantic technologies these features can be used to realize a *general purpose*, mobile collaboration platform, which can support the long tail of mobile special interest applications, for which the development of individual tools would not be (economically) feasible.

In [51] we present the *OntoWiki Mobile* approach realizing a mobile semantic collaboration platform based on the OntoWiki. It comprises specifically adopted user interfaces for browsing, faceted navigation as well as authoring of knowledge bases. It allows users to collect instance data and refine the structured knowledge bases on-the-go. OntoWiki Mobile is implemented as an *HTML5 web application*, thus being completely mobile device platform independent. In order to allow offline use in cases with restricted network coverage (or in order to avoid roaming charges) it uses the novel HTML5 local storage feature for replicating parts of the knowledge base on the mobile device. Hence, a crucial part of OntoWiki Mobile is the advanced conflict resolution for RDF stores. The approach is based on a combination of the EvoPat [172] method for data evolution and ontology refactoring along with a versioning system inspired by distributed version control systems like Git. OntoWiki Mobile is a generic, application domain agnostic tool, which can be utilized in a wide range of very different usage scenarios ranging from instance acquisition to browsing of semantic data on the go. Typical OntoWiki Mobile usage scenarios are settings where users need to author and access semantically structured information on the go or in settings where users are away from regular power supply and restricted to light-weight equipment (e. g. scientific expeditions).

Semantics-based Requirements Engineering. Semantic interoperability, linked data, and a shared conceptual foundation become increasingly important prerequisites in software development projects that are characterized by spatial dispersion, large numbers of stakeholders, and heterogeneous development tools. The SoftWiki OntoWiki extension [124] focuses specifically on semantic collaboration with respect to requirements engineering. Potentially very large and spatially distributed groups of stakeholders, including developers, experts, managers, and average users, shall be enabled to collect, semantically enrich, classify, and aggregate software requirements. OntoWiki is used to support collaboration as well as interlinking and exchange of requirements data. To ensure a shared conceptual foundation and semantic interoperability, we developed the SoftWiki Ontology for Requirements Engineering (SWORE) that defines core concepts of requirement engineering and the way they are interrelated. For instance, the ontology defines frequent relation types to describe requirements interdependencies such as details, conflicts, related to, depends on, etc. The flexible SWORE design allows for

easy extension. Moreover, the requirements can be linked to external resources, such as publicly available domain knowledge or company-specific policies. The whole process is called semantification of requirements. It is envisioned as an evolutionary process: The requirements are successively linked to each other and to further concepts in a collaborative way, jointly by all stakeholders. Whenever a requirement is formulated, reformulated, analyzed, or exchanged, it might be semantically enriched by the respective participant.

5 Linking

The fourth Linked Data Principle, i.e., “Include links to other URIs, so that they can discover more things” (cf. section 2) is the most important Linked Data principle as it enables the paradigm change from data silos to interoperable data distributed across the Web. Furthermore, it plays a key role in important tasks such as cross-ontology question answering [23,125], large-scale inferences [194,131] and data integration [126,21]. Yet, while the number of triples in Linked Data sources increases steadily and has surpassed 31 billions¹², links between knowledge bases still constitute less than 5% of these triples. The goal of linking is to tackle this sparseness so as to transform the Web into a platform for data and information integration as well as for search and querying.

5.1 Link Discovery

Linking can be generally defined as *connecting things that are somehow related*. In the context of Linked Data, the idea of linking is especially concerned with establishing typed links between entities (i.e., classes, properties or instances) contained in knowledge bases. Over the last years, several frameworks have been developed to address the lack of typed links between the different knowledge bases on the Linked Data web. Overall, two main categories of frameworks that aim to achieve this goal can be differentiated. The first category implements *ontology matching* techniques and aims to establish links between the ontologies underlying two data sources. The second and more prominent category of approaches, dubbed *instance matching* approaches (also called linking or link discovery approaches), aims to discover links between instances contained in two data sources. It is important to notice that while ontology and instance matching are similar to schema matching [167,166] and record linkage [205,50,27] respectively (as known in the research area of databases), linking on the Web of Data is a more generic and thus more complex task, as it is not limited to finding equivalent entities in two knowledge bases. Rather, it aims at finding semantically related entities and establishing typed links between them, most of these links being imbued with formal properties (e.g., transitivity, symmetry, etc.) that can be used by reasoners and other application to infer novel knowledge. In this section, we will focus on the discovery of links between instances and use the term link discovery as name for this process. An overview of ontology matching techniques is given in [54].

Formally, link discovery can be defined as follows:

Definition 3 (Link Discovery). *Given two sets S (source) and T (target) of instances and a relation R , find the set $M \subseteq S \times T$ that is such that $\forall (s, t) \in M : R(s, t)$.*

Solving this problem is obviously not trivial. Declarative frameworks thus reduce this problem to a similarity computation problem and try to approximate M by a set \tilde{M} for which the following holds:

Definition 4 (Link Discovery as Similarity Computation). *Given two sets S (source) and T (target) of instances, a (complex) semantic similarity measure $\sigma : S \times T \rightarrow [0, 1]$*

¹² <http://lod-cloud.net/>

and a threshold $\theta \in [0, 1]$, the goal of link discovery task is to compute the set $\tilde{M} = \{(s, t), \sigma(s, t) \geq \theta\}$.

In general, the similarity function used to carry out a link discovery task is described by using a *link specification* (sometimes called *linkage decision rule* [91]).

5.2 Challenges

Two key challenges arise when trying to discover links between two sets of instances: the computational complexity of the matching task *per se* and the selection of an appropriate link specification. The first challenge is intrinsically related to the link discovery process. The time complexity of a matching task can be measured by the number of comparisons necessary to complete this task. When comparing a source knowledge base S with a target knowledge base T , the completion of a matching task requires a-priori $O(|S||T|)$ comparisons, an impractical proposition as soon as the source and target knowledge bases become large. For example, discovering duplicate cities in DBpedia [7] alone would necessitate approximately 0.15×10^9 similarity computations. Hence, the provision of time-efficient approaches for the reduction of the time complexity of link discovery is a key requirement to instance linking frameworks for Linked Data.

The second challenge of the link discovery process lies in the selection of an appropriate link specification. The configuration of link discovery frameworks is usually carried out manually, in most cases simply by guessing. Yet, the choice of a suitable link specification measure is central for the discovery of satisfactory links. The large number of properties of instances and the large spectrum of measures available in literature underline the complexity of choosing the right specification manually¹³. Supporting the user during the process of finding the appropriate similarity measure and the right properties for each mapping task is a problem that still needs to be addressed by the Linked Data community. Methods such as supervised and active learning can be used to guide the user in need of mapping to a suitable linking configuration for his matching task. In the following, we give a short overview of existing frameworks for Link Discovery on the Web of Data. Subsequently, we present a time-efficient framework for link discovery in more detail and show how it can detect link specifications using active learning.

5.3 Approaches to Link Discovery

Current frameworks for link discovery can be subdivided into two main categories: *domain-specific* and *universal* frameworks. Domain-specific link discovery frameworks aim at discovering links between knowledge bases from a particular domain. One of the first domain-specific approaches to carry out instance linking for Linked Data was implemented in the *RKBExplorer*¹⁴ [69] with the aim of discovering links between entities from the domain of academics. Due to the lack of data available as Linked Data,

¹³ The SimMetrics project (<http://simmetrics.sf.net>) provides an overview of strings similarity measures.

¹⁴ <http://www.rkbexplorer.com>

the RKBExplorer had to extract RDF from heterogeneous data source so as to populate its knowledge bases with instances according to the AKT ontology¹⁵. Especially, instances of persons, publications and institutions were retrieved from several major metadata websites such as ACM and DBLP. The linking was implemented by the so-called Consistent Reference Service (CRS) which linked equivalent entities by comparing properties including their type and label. So far, the CRS is limited to linking objects in the knowledge bases underlying the RKBExplorer and cannot be used for other tasks without further implementation.

Another domain-specific tool is GNAT [168], which was developed for the music domain. It implements several instance matching algorithms of which the most sophisticated, the online graph matching algorithm (OGMA), applies a similarity propagation approach to discover equivalent resources. The basic approach implemented by OGMA starts with a single resource $s \in S$. Then, it retrieves candidate matching resources $t \in T$ by comparing properties such as `foaf:name` for artists and `dc:title` for albums. If $\sigma(s, t) \geq \theta$, then the algorithm terminates. In case a disambiguation is needed, the resources related to s and t in their respective knowledge bases are compared and their similarity value is cumulated to recompute $\sigma(s, t)$. This process is iterated until a mapping resource for s is found in T or no resource matches.

Universal link discovery frameworks are designed to carry out mapping tasks independently from the domain of the source and target knowledge bases. For example, RDF-AI [181], a framework for the integration of RDF data sets, implements a five-step approach that comprises the preprocessing, matching, fusion, interlinking and post-processing of data sets. RDF-AI contains a series of modules that allow for computing instances matches by comparing their properties. Especially, it contains translation modules that allow to process the information contained in data sources before mapping. By these means, it can boost the precision of the mapping process. These modules can be configured by means of XML-files. RDF-AI does not comprise means for querying distributed data sets via SPARQL¹⁶. In addition, it suffers from not being time-optimized. Thus, mapping by using this tool can be very time-consuming.

A time-optimized approach to link discovery is implemented by the LIMES framework [148,147,153] (Link Discovery Framework for metric spaces).¹⁷ The idea behind the LIMES framework is to use the mathematical characteristics of similarity and distance measures to reduce the number of computations that have to be carried out by the system without losing any link. For example, LIMES can make use of the fact that the edit distance is a distance metric to approximate distances without having to compute them [148]. Moreover, it implements the reductio-ratio-optimal space tiling algorithm \mathcal{HR}^3 to compute similarities in affine spaces with Minkowski measures [146]. In contrast to other frameworks (of which most rely on blocking), LIMES relies on time-efficient set operators to combine the results of these algorithms efficiently and has been shown to outperform the state of the art by these means [147]. Moreover, LIMES

¹⁵ <http://www.aktors.org/publications/ontology/>

¹⁶ <http://www.w3.org/TR/rdf-sparql-query/>

¹⁷ <http://limes.sf.net>. A graphical user interface can be found at <http://saim.aksw.org>.

implements unsupervised and supervised machine learning approaches for detecting high-quality link specifications [153,154].

Another link discovery framework is SILK [198]. SILK implements several approaches to minimize the time necessary for mapping instances from knowledge bases. In addition to implementing rough index pre-matching to reach a quasi-linear time-complexity, SILK also implements a lossless blocking algorithm called MultiBlock [92] to reduce its overall runtime. The approach relies on generating overlapping blocks of instances and only comparing pairs of instances that are located in the same block. Moreover, SILK provides supervised machine learning approaches for link discovery [93].

It is important to notice that the task of discovering links between knowledge bases is related with record linkage [205,50] and de-duplication [27]. The database community has produced a vast amount of literature on efficient algorithms for solving these problems. Different blocking techniques such as standard blocking, sorted-neighborhood, bigram indexing, canopy clustering and adaptive blocking [20,24,102] have been developed to address the problem of the quadratic time complexity of brute force comparison methods. The idea is to filter out obvious non-matches efficiently before executing the more detailed and time-consuming comparisons. In the following, we present a state-of-the-art framework that implements lossless instance matching based on a similar idea in detail.

5.4 The LIMES Algorithm

The original LIMES algorithm described in [148] addresses the scalability problem of link discovery by utilizing the *triangle inequality* in metric spaces to compute pessimistic estimates of instance similarities. Based on these approximations, LIMES can filter out a large number of instance pairs that cannot suffice the matching condition set by the user. The real similarities of the remaining instances pairs are then computed and the matching instances are returned.

Mathematical Framework In the remainder of this section, we use the following notations:

1. A is an affine space,
2. m, m_1, m_2, m_3 symbolize metrics on A ,
3. x, y and z represent points from A and
4. α, β, γ and δ are scalars, i.e., elements of \mathbb{R} .

Definition 5 (Metric space). A metric space is a pair (A, m) such that A is an affine space and $m : A \times A \rightarrow \mathbb{R}$ is a function such that for all x, y and $z \in A$

1. $m(x, y) \geq 0$ (M_1) (non-negativity),
2. $m(x, y) = 0 \Leftrightarrow x = y$ (M_2) (identity of indiscernibles),
3. $m(x, y) = m(y, x)$ (M_3) (symmetry) and
4. $m(x, z) \leq m(x, y) + m(y, z)$ (M_4) (triangle inequality).

Note that the definition of a matching based on a similarity function σ can be rewritten for metrics m as follows:

Definition 6 (Instance Matching in Metric Spaces). *Given two sets S (source) and T (target) of instances, a metric $m : S \times T \rightarrow [0, \infty[$ and a threshold $\theta \in [0, \infty[$, the goal of instance matching task is to compute the set $M = \{(s, t) | m(s, t) \leq \theta\}$.*

Example of metrics on strings include the Levenshtein distance and the block distance. However, some popular measures such as JaroWinkler [204] do not satisfy the triangle inequality and are consequently not metrics. The rationale behind the LIMES framework is to make use of the boundary conditions entailed by the triangle inequality (TI) to reduce the number of comparisons (and thus the time complexity) necessary to complete a matching task. Given a metric space (A, m) and three points x, y and z in A , the TI entails that

$$m(x, y) \leq m(x, z) + m(z, y). \quad (5)$$

Without restriction of generality, the TI also entails that

$$m(x, z) \leq m(x, y) + m(y, z), \quad (6)$$

thus leading to the following boundary conditions in metric spaces:

$$m(x, y) - m(y, z) \leq m(x, z) \leq m(x, y) + m(y, z). \quad (7)$$

Inequality 7 has two major implications. The first is that the distance from a point x to any point z in a metric space can be approximated given the distance from x to a reference point y and the distance from the reference point y to z . Such a reference point is called an *exemplar* following [62]. The role of an exemplar is to be used as a sample of a portion of the metric space A . Given an input point x , knowing the distance from x to an exemplar y allows to compute lower and upper bounds of the distance from x to any other point z at a known distance from y . An example of such an approximation is shown in Figure 20. In this figure, all the points on the circle are subject to the same distance approximation. The distance from x to z is close to the lower bound of inequality 7, while the distance from x to z' is close to the upper bound of the same inequality.

The second implication of inequality 7 is that the distance from x to z can only be smaller than θ if the lower bound of the approximation of the distance from x to z via any exemplar y is also smaller than θ . Thus, if the lower bound of the approximation of the distance $m(x, z)$ is larger than θ , then $m(x, z)$ itself must be larger than θ . Formally,

$$m(x, y) - m(y, z) > \theta \Rightarrow m(x, z) > \theta. \quad (8)$$

Supposing that all distances from instances $t \in T$ to exemplars are known, reducing the number of comparisons simply consists of using inequality 8 to compute an approximation of the distance from all $s \in S$ to all $t \in T$ and computing the real distance only for the (s, t) pairs for which the first term of inequality 8 does not hold. This is the core of the approach implemented by LIMES.

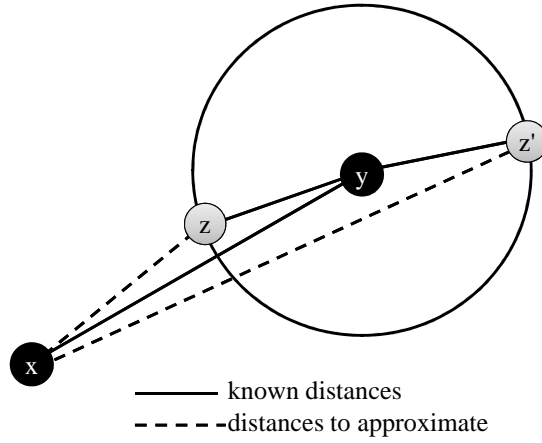


Fig. 15: Approximation of distances via exemplars. The lower bound of the distance from x to z can be approximated by $m(x, y) - m(y, z)$.

Computation of Exemplars The core idea underlying the computation of exemplars in LIMES is to select a set of exemplars in the metric space underlying the matching task in such a way that they are distributed uniformly in the metric space. One way to achieve this goal is by ensuring that the exemplars display a high dissimilarity. The approach used by LIMES to generate exemplars with this characteristic is shown in Algorithm 1.

Algorithm 1 Computation of Exemplars

Require: Number of exemplars n

Require: Target knowledge base T

1. Pick random point $e_1 \in T$

2. Set $E = E \cup \{e_1\}$;

3. Compute the distance from e_1 to all $t \in T$

while $|E| < n$ **do**

4. Get a random point e' such that $e' \in \operatorname{argmax}_t \sum_{t \in T} \sum_{e \in E} m(t, e)$

5. $E = E \cup \{e'\}$;

6. Compute the distance from e' to all $t \in T$

end while

7. Map each point in $t \in T$ to one of the exemplars $e \in E$ such that $m(t, e)$ is minimal

return E

Let n be the desired number of exemplars and E the set of all exemplars. In step 1 and 2, LIMES initializes E by picking a random point e_1 in the metric space (T, m) and setting $E = \{e_1\}$. Then, it computes the similarity from the exemplar e_1 to every other point in T (step 3). As long as the size of E has not reached n , LIMES repeats steps 4 to 6: In step 4, a point $e' \in T$ such that the sum of the distances from e' to the exemplars

$e \in E$ is maximal (there can be many of these points) is chosen randomly. This point is chosen as new exemplar and consequently added to E (step 5). Then, the distance from e' to all other points in T is computed (step 6). Once E has reached the size n , LIMES terminates the iteration. Finally, each point is mapped to the exemplar to which it is most similar (step 7) and the exemplar computation terminates (step 8). This algorithm has a constant time complexity of $O(|E||T|)$.

An example of the results of the exemplar computation algorithm ($|E| = 3$) is shown in Figure 16. The initial exemplar was the leftmost exemplar in the figure.

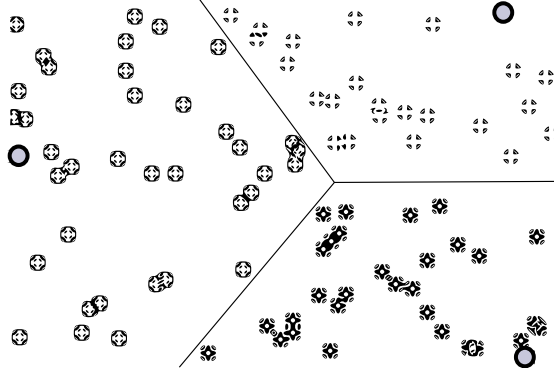


Fig. 16: Mapping of points to three exemplars in a metric space. The exemplars are displayed as gray disks.

Matching Based on Exemplars The instances associated with an exemplar $e \in E$ in step 7 of Algorithm 1 are stored in a list L_e sorted in descending order with respect to the distance to e . Let $\lambda_1^e \dots \lambda_m^e$ be the elements of the list L_e . The goal of matching an instance s from a source knowledge base to a target knowledge base w.r.t. a metric m is to find all instances t of the target knowledge source such that $m(s, t) \leq \theta$, where θ is a given threshold. LIMES achieves this goal by using the matching algorithm based on exemplars shown in Algorithm 2.

LIMES only carries out a comparison when the approximation of the distance is less than the threshold. Moreover, it terminates the similarity computation for an exemplar e as soon as the first λ^e is found such that the lower bound of the distance is larger than θ . This break can be carried out because the list L_e is sorted, i.e., if $m(s, e) - m(e, \lambda_i^e) > \theta$ holds for the i^{th} element of L_e , then the same inequality holds for all $\lambda_j^e \in L_e$ with $j > i$. In the worst case, LIMES' matching algorithm has the time complexity $O(|S||T|)$, leading to a total worst time complexity of $O((|E| + |S|)|T|)$, which is larger than that of brute force approaches. However, as the results displayed in Figure 17 show, a correct parameterization of LIMES leads to significantly smaller numbers of comparisons and runtimes.

Algorithm 2 LIMES' Matching algorithm

Require: Set of exemplars E

Require: Instance $s \in S$

Require: Metric m

Require: threshold θ

```
1.  $M = \emptyset$ 
  for  $e \in |E|$  do
    if  $m(s, e) \leq \theta$  then
      2.  $M = M \cup \{e\}$ 
      for  $i = 1 \dots |L_e|$  do
        if  $(m(s, e) - m(e, \lambda_i^e)) \leq \theta$  then
          if  $m(s, \lambda_i^e) \leq \theta$  then
            3.  $M = M \cup \{(s, \lambda_i^e)\}$ 
          end if
        else
          break
        end if
      end for
    end if
  end for
return  $M$ 
```

5.5 The \mathcal{HR}^3 algorithm

Let S resp. T be the source and target of a Link Discovery task. One of the key ideas behind time-efficient Link Discovery algorithms \mathcal{A} is to reduce the number of comparisons that are effectively carried out to a number $C(\mathcal{A}) < |S||T|$. The reduction ratio RR of an algorithm \mathcal{A} is given by

$$RR(\mathcal{A}) = 1 - \frac{C(\mathcal{A})}{|S||T|}. \quad (9)$$

$RR(\mathcal{A})$ captures how much of the Cartesian product $|S||T|$ was not explored before the output of \mathcal{A} was reached. It is obvious that even an optimal lossless solution which performs only the necessary comparisons cannot achieve a RR of 1. Let C_{min} be the minimal number of comparisons necessary to complete the Link Discovery task without losing recall, i.e., $C_{min} = |\mathcal{M}|$. The relative reduction ratio $RRR(\mathcal{A})$ is defined as the portion of the minimal number of comparisons that was carried out by the algorithm \mathcal{A} before it terminated. Formally

$$RRR(\mathcal{A}) = \frac{1 - \frac{C_{min}}{|S||T|}}{1 - \frac{C(\mathcal{A})}{|S||T|}} = \frac{|S||T| - C_{min}}{|S||T| - C(\mathcal{A})}. \quad (10)$$

$RRR(\mathcal{A})$ indicates how close \mathcal{A} is to the optimal solution with respect to the number of candidates it tests. Given that $C(\mathcal{A}) \geq C_{min}$, $RRR(\mathcal{A}) \geq 1$. Note that the larger the value of $RRR(\mathcal{A})$, the poorer the performance of \mathcal{A} with respect to the task at hand.

The main observation that led \mathcal{HR}^3 is that while most algorithms aim to optimize their RR (and consequently their RRR), most approaches do not provide any guarantee

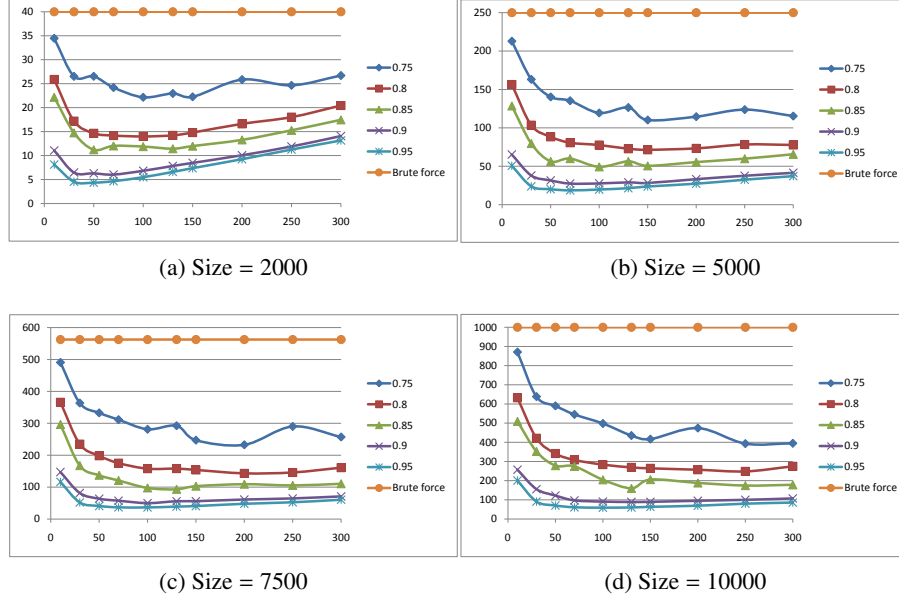


Fig. 17: Comparisons required by LIMES for different numbers of exemplars on knowledge bases of different sizes. The x-axis shows the number of exemplars, the y-axis the number of comparisons in multiples of 10^5 .

with respect to the RR (and consequently the RRR) that they can achieve. The approach is the first mathematically optimal algorithm w.r.t the reduction ratio that it can achieve, i.e., the first approach such that given any relative reduction ratio r , there is always a setting that leads \mathcal{HR}^3 achieving a relative reduction ratio r' with $r' \leq r$. To achieve this goal \mathcal{HR}^3 relies on space tiling as introduced by the HYPPO algorithm [145].

Space Tiling for Link Discovery HYPPO addresses the problem of efficiently mapping instance pairs $(s, t) \in S \times T$ described by using exclusively numeric values in a n -dimensional metric space and has been shown to outperform the state of the art in previous work [145]. The observation behind space tiling is that in spaces (\mathcal{Q}, δ) with orthogonal, (i.e., uncorrelated) dimensions¹⁸, common metrics for Link Discovery can be decomposed into the combination of functions $\phi_{i,i \in \{1 \dots n\}}$ which operate on exactly one dimension of \mathcal{Q} : $\delta = f(\phi_1, \dots, \phi_n)$. For Minkowski distances of order p , $\phi_i(x, \omega) = |x_i - \omega_i|$ for all values of i and $\delta(x, \omega) = \sqrt[p]{\sum_{i=1}^n \phi_i^p(x, \omega)}$. A direct consequence of this observation is the inequality $\phi_i(x, \omega) \leq \delta(x, \omega)$. The basic insight that results this observation is that the hypersphere $H(\omega, \theta) = \{x \in \mathcal{Q} : \delta(x, \omega) \leq \theta\}$ is a subset of the hypercube V

¹⁸ Note that in all cases, a space transformation exists that can map a space with correlated dimensions to a space with uncorrelated dimensions.

defined as $V(\omega, \theta) = \{x \in \Omega : \forall i \in \{1 \dots n\}, \phi_i(x_i, \omega_i) \leq \theta\}$. Consequently, one can reduce the number of comparisons necessary to detect all elements of $H(\omega, \theta)$ by discarding all elements which are not in $V(\omega, \theta)$ as non-matches. Let $\Delta = \theta/\alpha$, where $\alpha \in \mathbb{N}$ is the *granularity parameter* that controls how fine-grained the space tiling should be (see Figure 18 for an example). We first tile Ω into the adjacent hypercubes (short: cubes) C that contain all the points ω such that

$$\forall i \in \{1 \dots n\}, c_i \Delta \leq \omega_i < (c_i + 1) \Delta \text{ with } (c_1, \dots, c_n) \in \mathbb{N}^n. \quad (11)$$

We call the vector (c_1, \dots, c_n) the coordinates of the cube C . Each point $\omega \in \Omega$ lies in the cube $C(\omega)$ with coordinates $(\lfloor \omega_i / \Delta \rfloor)_{i=1 \dots n}$. Given such a space tiling, it is obvious that $V(\omega, \theta)$ consists of the union of the cubes such that $\forall i \in \{1 \dots n\} : |c_i - c(\omega)_i| \leq \alpha$.

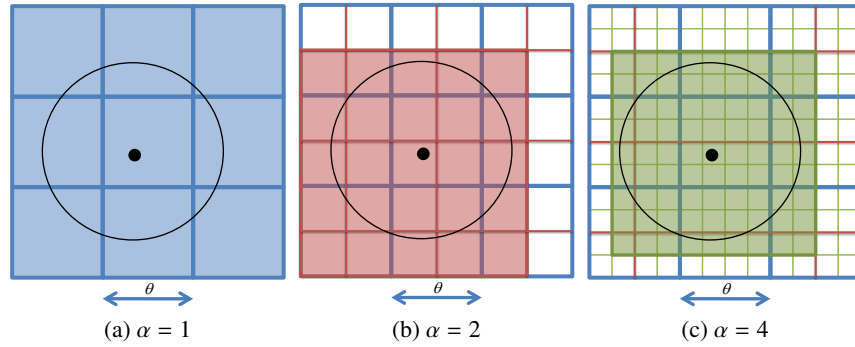


Fig. 18: Space tiling for different values of α . The colored squares show the set of elements that must be compared with the instance located at the black dot. The points within the circle lie within the distance θ of the black dot. Note that higher values of α lead to a better approximation of the hypersphere but also to more hypercubes.

\mathcal{HR}^3 's indexing scheme Let $\omega \in \Omega = S \cup T$ be an arbitrary reference point. Furthermore, let δ be the Minkowski distance of order p . The *index* function is defined as follows:

$$index(C, \omega) = \begin{cases} 0 & \text{if } \exists i : |c_i - c(\omega)_i| \leq 1 \text{ with } i \in \{1, \dots, n\}, \\ \sum_{i=1}^n (|c_i - c(\omega)_i| - 1)^p & \text{else,} \end{cases} \quad (12)$$

where C is a hypercube resulting from a space tiling and $\omega \in \Omega$. Figure 19 shows an example of such indexes for $p = 2$ with $\alpha = 2$ (Figure 19a) and $\alpha = 4$ (Figure 19b). Note that the blue square with index 0 contains the reference point ω . All elements of C must only be compared with the elements of cubes C' such that $index(C, C') \leq \alpha^p$. The authors of [146] prove formally that given this approach to space tiling, the following theorem holds:

Theorem 1. $\lim_{\alpha \rightarrow \infty} RRR(\mathcal{HR}^3, \alpha) = 1.$

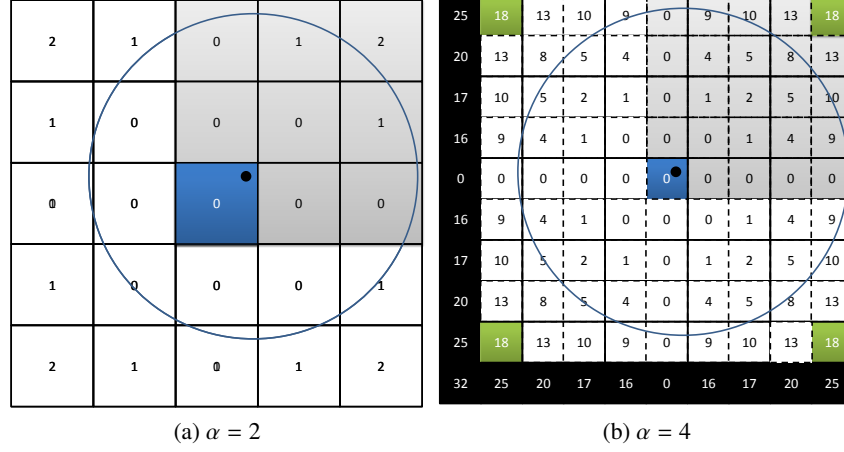


Fig. 19: Space tiling and resulting index for a two-dimensional example. Note that the index in both subfigures was generated for exactly the same portion of space. The black dot stands for the position of ω .

This conclusion is illustrated by Figure 20, which shows the space tiling computed by \mathcal{HR}^3 for different values of α with $p = 2$ and $n = 2$. The higher α , the closer the approximation is to a circle. Note that these results allow to conclude that for any RRR -value r larger than 1, there is a setting of \mathcal{HR}^3 that can compute links with a RRR smaller or equal to r .

Evaluation \mathcal{HR}^3 was evaluated against HYPPO w.r.t. to the number of comparisons that it has to carry out in several settings. In the first and second experiments, the goal was to deduplicate DBpedia places by comparing their names (`rdfs:label`), minimum elevation, elevation and maximum elevation. 2988 entities possessed all four properties. The Euclidean metric was applied to the last three values with the thresholds 49 meters resp. 99 meters for the first resp. second experiment. The third and fourth experiments aimed to discover links between Geonames and LinkedGeoData. This experiment was of considerably larger scale than the first one, as we compared 74458 entities in Geonames with 50031 entities from LinkedGeoData. Again, the number of comparisons necessary to complete the task by using the Euclidean metric was measured. The distance thresholds were set to 1 resp. 9° in experiment 3 resp. 4. We ran all experiments on the same Windows 7 Enterprise 64-bit computer with a 2.8GHz i7 processor with 8GB RAM. The JVM was allocated 7GB RAM to ensure that the runtimes were not influenced by swapping. Only one of the kernels of the processors was used.

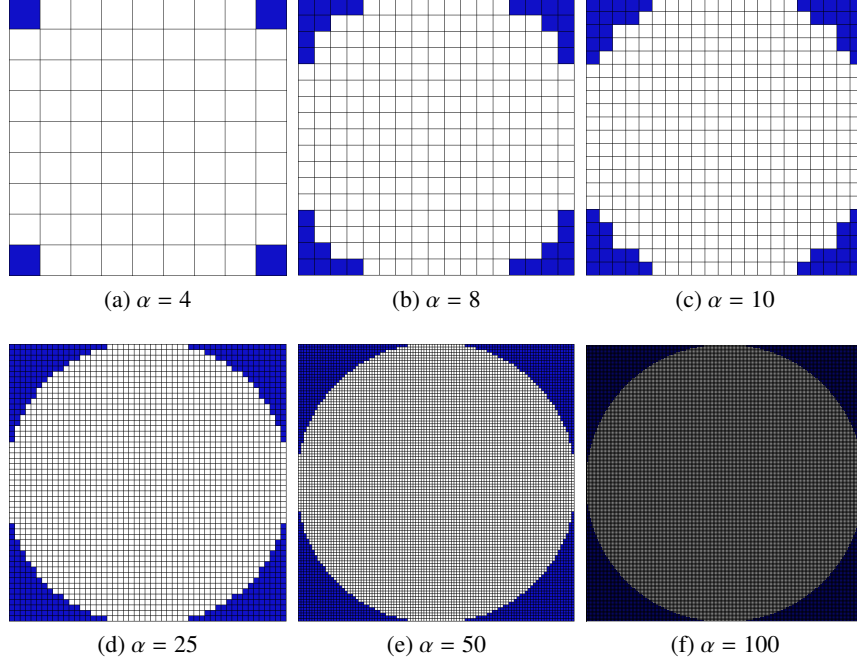


Fig. 20: Space tilings generated by \mathcal{HR}^3 for different values of α . The white squares are selected for comparisons with the elements of the square in the middle of the figure whilst the colored ones are discarded.

The results (see Figure 21) show that \mathcal{HR}^3 can reduce the overhead in comparisons (i.e., the number of unnecessary comparisons divided by the number of necessary comparisons) from approximately 24% for HYPPO to approximately 6% (granularity = 32). In experiment 2, the overhead is reduced from 4.1% to 2%. This difference in overhead reduction is mainly due to the data clustering around certain values and the clusters having a radius between 49 meters and 99 meters. Thus, running the algorithms with a threshold of 99 meters led to only a small a-priori overhead and HYPPO performing remarkably well. Still, even on such data distributions, \mathcal{HR}^3 was able to discard even more data and to reduce the number of unnecessary computations by more than 50% relative. In the best case (Exp. 4, $\alpha = 32$, see Figure 21d), \mathcal{HR}^3 required approximately 4.13×10^6 less comparisons than HYPPO for $\alpha = 32$. Even for the smallest setting (Exp. 1, see Figure 21a), \mathcal{HR}^3 still required 0.64×10^6 less comparisons.

5.6 Active Learning of Link Specifications

The second challenge of Link Discovery is the time-efficient discovery of link specifications for a particular linking task. Several approaches have been proposed to achieve this goal, of which most rely on genetic programming [93,154,152]. The COALA

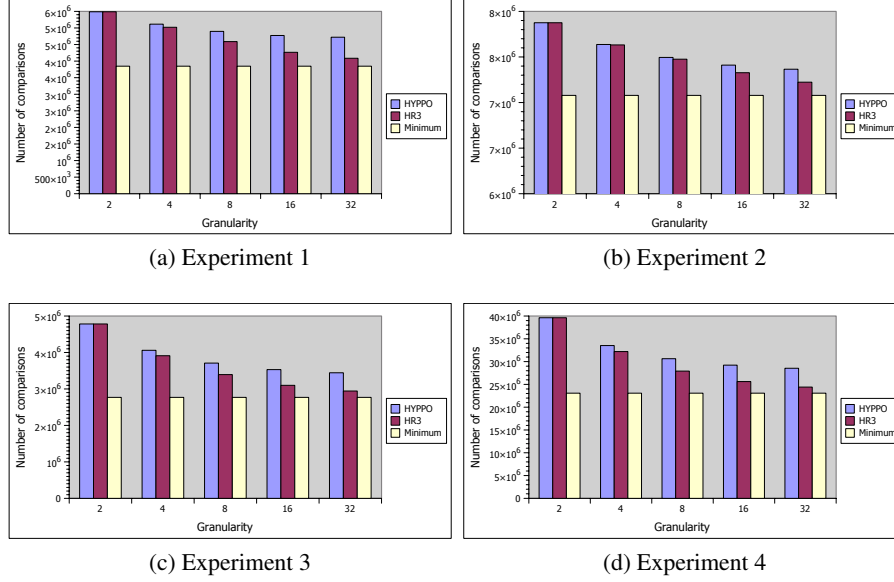


Fig. 21: Number of comparisons for \mathcal{HR}^3 and HYPPO.

(Correlation-Aware Active Learning) approach was implemented on top of the genetic programming approach EAGLE [153] with the aim of improving the selection of positive and negative examples during active learning. In the following, we give an overview of COALA.

Intuition Let \mathcal{N} be the set of most informative negative and \mathcal{P} the set of most informative negative examples w.r.t. an informativeness function ifm (e.g., the distance from the decision boundary). used by a curious classifier [183] The basic insight behind COALA is that the correlation between the features of the elements of \mathcal{N} and \mathcal{P} should play a role when computing the sets \mathcal{I}^+ and \mathcal{I}^- of positive resp. negative queries for the oracle. In particular, two main factors affect the information content of a link candidate: its similarity to elements of its presumed class and to elements of the other class. For the sake of simplicity, we will assume that the presumed class of the link candidate of interest is +1, i.e., that the link candidate was classified as positive by the current curious classifier. Our insights yet hold symmetrically for link candidates whose presumed class is -1.

Let $A = (s_A, t_A), B = (s_B, t_B) \in \mathcal{P}$ to be two link candidates which are equidistant from C 's boundary. Consider Figure 22a, where $\mathcal{P} = \{A, B, C\}$ and $\mathcal{N} = \{D\}$. The link candidate B is on on average most distant from any other elements of \mathcal{P} . Thus, it is more likely to be a statistical outlier than A . Hence, making a classification error on B should not have the same impact as an erroneous classification of link candidate A , which is close to another presumably positive link candidate, C . Consequently, B should

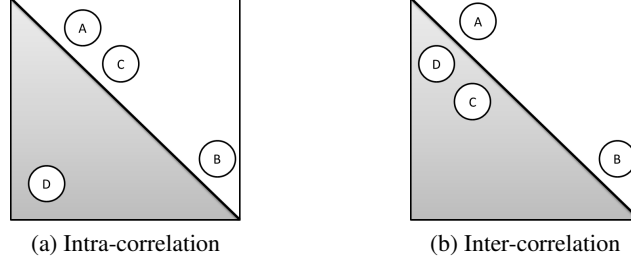


Fig. 22: Examples of correlations within classes and between classes. In each subfigure, the gray surface represent \mathcal{N} while the white surface stands for \mathcal{P} . The oblique line is C 's boundary.

be considered less informative than A . Approaches that make use of this information are said to exploit the *intra-class correlation*. Now, consider Figure 22b, where $\mathcal{P} = \{A, B\}$ and $\mathcal{N} = \{C, D\}$. While the probability of A being an outlier is the same as B 's, A is still to be considered more informative than B as it is located closer to elements of \mathcal{N} and can thus provide more information on where to set the classifier boundary. This information is dubbed *inter-class correlation*. Several approaches that make use of these two types of correlations can be envisaged. In the following, we present two approaches for these purposes. The first makes use of intra-class correlations and relies on graph clustering. The second approach relies on the spreading activation principle in combination with weight decay. We assume that the complex similarity function σ underlying C is computed by combining n atomic similarity functions $\sigma_1, \dots, \sigma_n$. This combination is most commonly carried out by using metric operators such as min, max or linear combinations.¹⁹ Consequently, each link candidate (s, t) can be described by a vector $(\sigma_1(s, t), \dots, \sigma_n(s, t)) \in [0, 1]^n$. We define the *similarity of link candidates* $sim : (S \times T)^2 \rightarrow [0, 1]$ to be the inverse of the Euclidean distance in the space spawned by the similarities σ_1 to σ_n . Hence, the similarity of two link candidates (s, t) and (s', t') is given by:

$$sim((s, t), (s', t')) = \frac{1}{1 + \sqrt{\sum_{i=1}^n (\sigma_i(s, t) - \sigma_i(s', t'))^2}}. \quad (13)$$

Note that we added 1 to the denominator to prevent divisions by 0.

Graph Clustering The basic intuition behind using clustering for COALA is that groups of very similar link candidates can be represented by a single link candidate. Consequently, once a representative of a group has been chosen, all other elements of the group become less informative. An example that illustrates this intuition is given in Figure 23. We implemented COALA based on clustering as follows: In each iteration, we begin by first selecting two sets $\mathcal{S}^+ \subseteq \mathcal{P}$ resp. $\mathcal{S}^- \subseteq \mathcal{N}$ that contain the positive resp.

¹⁹ See [147] for a more complete description of a grammar for link specifications.

negative link candidates that are most informative for the classifier at hand. Formally, \mathcal{S}^+ fulfills

$$\forall x \in \mathcal{S}^+ \forall y \in \mathcal{P}, y \notin \mathcal{S}^+ \rightarrow \text{ifm}(y) \leq \text{ifm}(x). \quad (14)$$

The analogous equation holds for \mathcal{S}^- . In the following, we will explain the further steps of the algorithm for \mathcal{S}^+ . The same steps are carried out for \mathcal{S}^- .

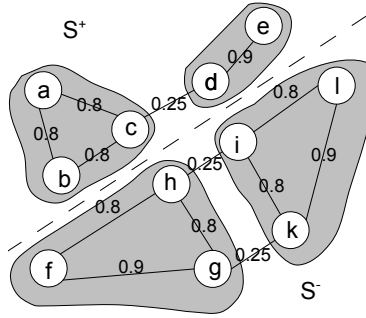


Fig. 23: Example of clustering. One of the most informative single link candidate is selected from each cluster. For example, d is selected from the cluster $\{d, e\}$.

First, we compute the similarity of all elements of \mathcal{S}^+ by using the similarity function shown in Equation 13. In the resulting similarity matrix, we set all elements of the diagonal to 0. Then, for each $x \in \mathcal{S}^+$, we only retain a fixed number ec of highest similarity values and set all others to 0. The resulting similarity matrix is regarded as the adjacency matrix of an undirected weighted graph $G = (V, E, sim)$. G 's set of nodes V is equal to \mathcal{S}^+ . The set of edges E is a set of 2-sets²⁰ of link candidates. Finally, the weighted function is the similarity function sim . Note that ec is the minimal degree of nodes in G .

In a second step, we use the graph G as input for a graph clustering approach. The resulting clustering is assumed to be a partition \mathcal{V} of the set V of vertices of G . The informativeness of partition $V_i \in \mathcal{V}$ is set to $\max_{x \in V_i} \text{ifm}(x)$. The final step of our approach consists of selecting the most informative node from each of the k most informative partitions. These are merged to generate I^+ , which is sent as query to the oracle. The computation of I^- is carried out analogously. Note that this approach is generic in the sense that it can be combined with any graph clustering algorithm that can process weighted graphs as well as with any informativeness function ifm . Here, we use Border-Flow [155] as clustering algorithm because (1) it has been used successfully in several other applications such as the creation of SPARQL benchmarks [137] and the analysis of protein-protein interactions [144]. and (2) it is parameter-free and does not require any tuning.

²⁰ A n -set is a set of magnitude n .

Spreading Activation with Weight Decay The idea behind spreading activation with weight decay (WD) is to combine the intra- and inter-class correlation to determine the informativeness of each link candidate. Here, we begin by computing the set $\mathcal{S} = \mathcal{S}^+ \cup \mathcal{S}^-$, where \mathcal{S}^+ and \mathcal{S}^- are described as above. Let s_i and s_j be the i^{th} and j^{th} elements of \mathcal{S} . We then compute the quadratic similarity matrix \mathcal{M} with entries $m_{ij} = \text{sim}(s_i, s_j)$ for $i \neq j$ and 0 else. Note that both negative and positive link candidates belong to \mathcal{S} . Thus, \mathcal{M} encodes both inter- and intra-class correlation. In addition to \mathcal{M} , we compute the activation vector \mathcal{A} by setting its entries to $a_i = \text{ifm}(s_i)$. In the following, \mathcal{A} is considered to be a column vector.

In a first step, we normalize the activation vector \mathcal{A} to ensure that the values contained therein do not grow indefinitely. Then, in a second step, we set $\mathcal{A} = \mathcal{A} + \mathcal{M} \times \mathcal{A}$. This has the effect of propagating the activation of each s to all its neighbors according to the weights of the edges between s and its neighbors. Note that elements of \mathcal{S}^+ that are close to elements of \mathcal{S}^- get a higher activation than elements of \mathcal{S}^+ that are further away from \mathcal{S}^- and vice-versa. Moreover, elements at the center of node clusters (i.e., elements that are probably no statistical outliers) also get a higher activation than elements that are probably outliers. The idea behind the weight decay step is to update the matrix by setting each m_{ij} to m_{ij}^r , where $r > 1$ is a fix exponent. This is the third step of the algorithm. Given that $\forall i \forall j m_{ij} \leq 1$, the entries in the matrix get smaller with time. By these means, the amount of activation transferred across long paths is reduced. We run this three-step procedure iteratively until all non-1 entries of the matrix are less or equal to a threshold $\epsilon = 10^{-2}$. The k elements of \mathcal{S}^+ resp. \mathcal{S}^- with maximal activation are returned as I^+ resp. I^- . In the example shown in Figure 24, while all nodes from \mathcal{S}^+ and \mathcal{S}^- start with the same activation, two nodes get the highest activation after only 3 iterations.

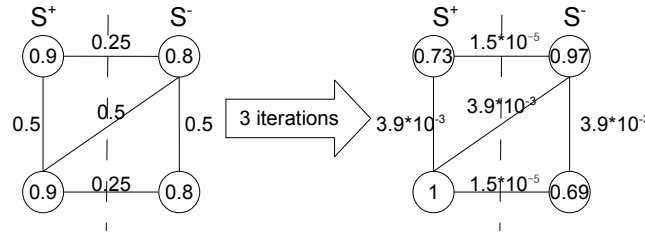


Fig. 24: Example of weight decay. Here r was set to 2. The left picture shows the initial activations and similarity scores while the right picture shows the results after 3 iterations. Note that for the sake of completeness the weights of the edges were not set to 0 when they reached ϵ .

Evaluation COALA was evaluated by running weight decay and clustering in combination with EAGLE, an active genetic programming approach for learning link specifi-

cations. Throughout the following experiments, EAGLE’s mutation and crossover rates were set to 0.6. Individuals were given a 70% chance to get selected for reproduction. The population sizes were set to 20 and 100. We set $k = 5$ and ran our experiments for 10 iterations. Between each iteration we evolved the populations for 50 generations. We ran our experiments on two real-world datasets and three synthetic datasets. The synthetic datasets consisted of the Persons1, Person2 and Restaurants datasets from the OAEI 2010 benchmark²¹. The real-world datasets consisted of the ACM-DBLP and Abt-Buy datasets, which were extracted from websites or databases [102]²². Given that genetic programming is non-deterministic, all results presented below are the means of 5 runs. Each experiment was ran on a single thread of a server running JDK1.7 on Ubuntu 10.0.4 and was allocated maximally 2GB of RAM. The processors were 2.0GHz Quad-core AMD Opterons. An excerpt of the results is shown in Figure 25. While the results show that COALA outperform EAGLE, it remains unclear whether WD or CL is the best approach to achieving a faster convergence towards the optimal solution.

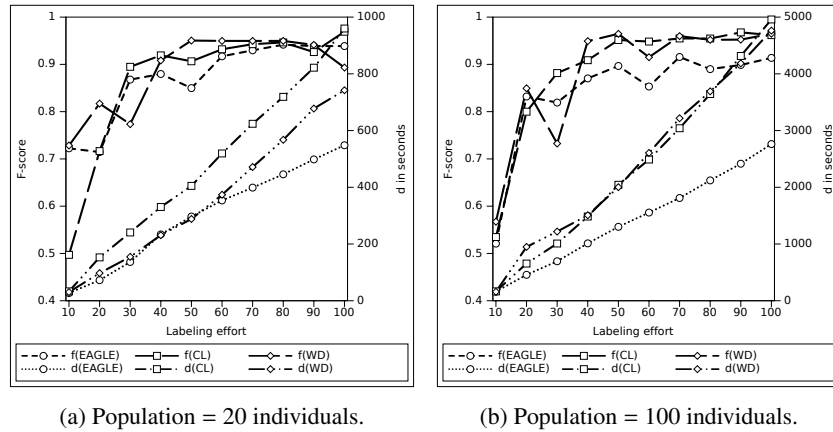


Fig. 25: F-score and runtime on the ACM-DBLP dataset. $f(X)$ stands for the F-score achieved by algorithm X , while $d(X)$ stands for the total duration required by the algorithm.

5.7 Conclusion

We presented and discussed linking approaches for Linked Data and the challenges they face. In addition, we gave an overview of several state-of-the-art approaches for instance

²¹ <http://oaei.ontologymatching.org/2010/>

²² http://dbs.uni-leipzig.de/en/research/projects/object_matching/fever/benchmark_datasets_for_entity_resolution

matching for Linked Data. We then presented time-efficient approaches for link discovery. Finally, we presented a state-of-the-art approach for the active learning of link specifications. This approach can be easily extended to learn specifications automatically.²³ Novel challenges that need to be addressed include the automatic management of resources for link specifications. First works on running link discovery in parallel have shown that using massively parallel hardware such as GPUs can lead to better results than using cloud implementations even on considerably large datasets [151]. Detecting the right resources for linking automatically given a hardware landscape is yet still a dream to achieve.

²³ Such an extension is the basis of the self-configuration algorithm of the SAIM framework.

6 Enrichment

The term *enrichment* in this chapter refers to the (semi-)automatic extension of a knowledge base schema. It describes the process of increasing the expressiveness and semantic richness of a knowledge base. Usually, this is achieved by adding or refining terminological axioms.

Enrichment methods can typically be applied in a *grass-roots* approach to knowledge base creation. In such an approach, the whole ontological structure is not created upfront, but evolves with the data in a knowledge base. Ideally, this enables a more agile development of knowledge bases. In particular, in the context of the Web of Linked Data such an approach appears to be an interesting alternative to more traditional ontology engineering methods. Amongst others, Tim Berners-Lee advocates to get “raw data now”²⁴ and worry about the more complex issues later.

Knowledge base enrichment can be seen as a sub-discipline of ontology learning. Ontology learning is more general in that it can rely on external sources, e.g. written text, to create an ontology. The term knowledge base enrichment is typically used when already existing data in the knowledge base is analysed to improve its schema.

Enrichment methods span several research areas like knowledge representation and reasoning, machine learning, statistics, natural language processing, formal concept analysis and game playing. Considering the variety of methods, we structure this section as follows: First, we give an overview of different types of enrichment and list some typical methods and give pointers to references, which allow the reader to obtain more information on a topic. In the second part, we describe a specific software – the ORE tool – in more detail.

6.1 State of the Art and Types of Enrichment

Ontology enrichment usually involves applying heuristics or machine learning techniques to find axioms, which can be added to an existing ontology. Naturally, different techniques have been applied depending on the specific type of axiom.

One of the most complex tasks in ontology enrichment is to find *definitions* of classes. This is strongly related to Inductive Logic Programming (ILP) [158] and more specifically supervised learning in description logics. Research in those fields has many applications apart from being applied to enrich ontologies. For instance, it is used in the life sciences to detect whether drugs are likely to be efficient for particular diseases. Work on learning in description logics goes back to e.g. [43,44], which used so-called *least common subsumers* to solve the learning problem (a modified variant of the problem defined in this article). Later, [19] invented a refinement operator for $\mathcal{AL}\mathcal{ER}$ and proposed to solve the problem by using a top-down approach. [52,89,90] combine both techniques and implement them in the YINYANG tool. However, those algorithms tend to produce very long and hard-to-understand class expressions. The algorithms implemented in DL-Learner [115,116,107,117,109] overcome this problem and investigate the learning problem and the use of top down refinement in detail. DL-FOIL [55] is a similar approach, which is based on a mixture of upward and downward refinement

²⁴ http://www.ted.com/talks/tim_bern timers_lee_on_the_next_web.html

of class expressions. They use alternative measures in their evaluation, which take the open world assumption into account, which was not done in ILP previously. Most recently, [111] implements appropriate heuristics and adaptations for learning definitions in ontologies. The focus in this work is efficiency and practical application of learning methods. The article presents plugins for two ontology editors (Protégé and OntoWiki) as well stochastic methods, which improve previous methods by an order of magnitude. For this reason, we will analyse it in more detail in the next subsection. The algorithms presented in the article can also learn *super class axioms*.

A different approach to learning the definition of a named class is to compute the so called *most specific concept* (msc) for all instances of the class. The most specific concept of an individual is the most specific class expression, such that the individual is instance of the expression. One can then compute the *least common subsumer* (lcs) [18] of those expressions to obtain a description of the named class. However, in expressive description logics, an msc does not need to exist and the lcs is simply the disjunction of all expressions. For light-weight logics, such as \mathcal{EL} , the approach appears to be promising.

Other approaches, e.g. [122] focus on learning in hybrid knowledge bases combining ontologies and *rules*. Ontology evolution [123] has been discussed in this context. Usually, hybrid approaches are a generalisation of concept learning methods, which enable powerful rules at the cost of efficiency (because of the larger search space). Similar as in knowledge representation, the tradeoff between expressiveness of the target language and efficiency of learning algorithms is a critical choice in symbolic machine learning.

Another enrichment task is *knowledge base completion*. The goal of such a task is to make the knowledge base complete in a particular well-defined sense. For instance, a goal could be to ensure that all subclass relationships between named classes can be inferred. The line of work starting in [173] and further pursued in e.g. [17] investigates the use of *formal concept analysis* for completing knowledge bases. It is promising, although it may not be able to handle noise as well as a machine learning technique. A Protégé plugin [182] is available. [196] proposes to improve knowledge bases through relational exploration and implemented it in the *RELExO framework*²⁵. It focuses on simple relationships and the knowledge engineer is asked a series of questions. The knowledge engineer either must positively answer the question or provide a counterexample.

[197] focuses on learning *disjointness* between classes in an ontology to allow for more powerful reasoning and consistency checking. To achieve this, it can use the ontology itself, but also texts, e.g. Wikipedia articles corresponding to a concept. The article includes an extensive study, which shows that proper modelling disjointness is actually a difficult task, which can be simplified via this ontology enrichment method.

Another type of ontology enrichment is schema mapping. This task has been widely studied and will not be discussed in depth within this chapter. Instead, we refer to [41] for a survey on ontology mapping.

There are further more light-weight ontology enrichment methods. For instance, *taxonomies* can be learned from simple tag structures via heuristics [36,195]. Similarly,

²⁵ <http://code.google.com/p/relexo/>

“properties of properties” can be derived via simple statistical analysis. This includes the detection whether a particular property might be symmetric, function, reflexive, inverse functional etc. Similarly, domains and ranges of properties can be determined from existing data. Enriching the schema with domain and range axioms allows to find cases, where properties are misused via OWL reasoning.

Type/Aim	References
Taxonomies	[207,36,195]
Definitions	often done via ILP approaches such as [115,116,117,111,55,52,89,90,19], genetic approaches [107] have also been used
Super Class Axioms	[111,195,36]
Rules in Ontologies	[122,123]
Disjointness	[197]
Properties of properties	[36,58]
Alignment	challenges: [185], recent survey: [41]
Completion	formal concept analysis and relational exploration [17,196,182]

Table 5: Work in ontology enrichment grouped by type or aim of learned structures.

In the following subsection, we describe an enrichment approach for learning definitions and super class axioms in more detail. The algorithm was recently developed by the first authors and is described in full detail in [111].

6.2 Class Expression Learning in DL-Learner

The Semantic Web has recently seen a rise in the availability and usage of knowledge bases, as can be observed within the Linking Open Data Initiative, the TONES and Protégé ontology repositories, or the Watson search engine. Despite this growth, there is still a lack of knowledge bases that consist of sophisticated schema information and instance data adhering to this schema. Several knowledge bases, e.g. in the life sciences, only consist of schema information, while others are, to a large extent, a collection of facts without a clear structure, e.g. information extracted from data bases or texts. The combination of sophisticated schema and instance data allows powerful reasoning, consistency checking, and improved querying possibilities. We argue that being able to learn OWL class expressions²⁶ is a step towards achieving this goal.

Example 1. As an example, consider a knowledge base containing a class **Capital** and instances of this class, e.g. London, Paris, Washington, Canberra etc. A machine learning algorithm could, then, suggest that the class **Capital** may be equivalent to one of the following OWL class expressions in Manchester OWL syntax²⁷:

City and isCapitalOf at least one GeopoliticalRegion
City and isCapitalOf at least one Country

²⁶ http://www.w3.org/TR/owl2-syntax/#Class_Expressions

²⁷ For details on Manchester OWL syntax (e.g. used in Protégé, OntoWiki) see [87].

Both suggestions could be plausible: The first one is more general and includes cities that are capitals of states, whereas the latter one is stricter and limits the instances to capitals of countries. A knowledge engineer can decide which one is more appropriate, i.e. a semi-automatic approach is used, and the machine learning algorithm should guide her by pointing out which one fits the existing instances better. Assuming the knowledge engineer decides for the latter, an algorithm can show her whether there are instances of the class `Capital` which are neither instances of `City` nor related via the property `isCapitalOf` to an instance of `Country`.²⁸ The knowledge engineer can then continue to look at those instances and assign them to a different class as well as provide more complete information; thus improving the quality of the knowledge base. After adding the definition of `Capital`, an OWL reasoner can compute further instances of the class which have not been explicitly assigned before.

Using machine learning for the generation of suggestions instead of entering them manually has the advantage that 1.) the given suggestions fit the instance data, i.e. schema and instances are developed in concordance, and 2.) the entrance barrier for knowledge engineers is significantly lower, since understanding an OWL class expression is easier than analysing the structure of the knowledge base and creating a class expression manually. Disadvantages of the approach are the dependency on the availability of instance data in the knowledge base and requirements on the quality of the ontology, i.e. modelling errors in the ontology can reduce the quality of results.

Overall, we describe the following in this chapter:

- extension of an existing learning algorithm for learning class expressions to the ontology engineering scenario,
- presentation and evaluation of different heuristics,
- showcase how the enhanced ontology engineering process can be supported with plugins for Protégé and OntoWiki,
- evaluation of the presented algorithm with several real ontologies from various domains.

The adapted algorithm for solving the learning problems, which occur in the ontology engineering process, is called *CELOE* (*Class Expression Learning for Ontology Engineering*). It was implemented within the open-source framework DL-Learner.²⁹ DL-Learner [108,109] leverages a modular architecture, which allows to define different types of components: knowledge sources (e.g. OWL files), reasoners (e.g. DIG³⁰ or OWL API based), learning problems, and learning algorithms. In this overview, we focus on the latter two component types, i.e. we define the class expression learning problem in ontology engineering and provide an algorithm for solving it.

Learning Problem The process of learning in logics, i.e. trying to find high-level explanations for given data, is also called *inductive reasoning* as opposed to *inference* or *deductive reasoning*. The main difference is that in deductive reasoning it is formally

²⁸ This is not an inconsistency under the standard OWL open world assumption, but rather a hint towards a potential modelling error.

²⁹ <http://dl-learner.org>

³⁰ <http://dl.kr.org/dig/>

shown whether a statement follows from a knowledge base, whereas in inductive learning new statements are invented. Learning problems, which are similar to the one we will analyse, have been investigated in *Inductive Logic Programming* [158] and, in fact, the method presented here can be used to solve a variety of machine learning tasks apart from ontology engineering.

In the ontology learning problem we consider, we want to learn a formal description of a class A , which has (inferred or asserted) instances in the considered ontology. In the case that A is already described by a class expression C via axioms of the form $A \sqsubseteq C$ or $A \equiv C$, those can be either refined, i.e. specialised/generalised, or relearned from scratch by the learning algorithm. To define the class learning problem, we need the notion of a *retrieval* reasoner operation $R_{\mathcal{K}}(C)$. $R_{\mathcal{K}}(C)$ returns the set of all instances of C in a knowledge base \mathcal{K} . If \mathcal{K} is clear from the context, the subscript can be omitted.

Definition 7 (class learning problem). *Let an existing named class A in a knowledge base \mathcal{K} be given. The class learning problem is to find an expression C such that $R_{\mathcal{K}}(C) = R_{\mathcal{K}}(A)$.*

Clearly, the learned expression C is a description of (the instances of) A . Such an expression is a candidate for adding an axiom of the form $A \equiv C$ or $A \sqsubseteq C$ to the knowledge base \mathcal{K} . If a solution of the learning problem exists, then the used base learning algorithm (as presented in the following subsection) is complete, i.e. guaranteed to find a correct solution if one exists in the target language and there are no time and memory constraints (see [116,117] for the proof). In most cases, we will not find a solution to the learning problem, but rather an approximation. This is natural, since a knowledge base may contain false class assignments or some objects in the knowledge base are described at different levels of detail. For instance, in Example 1, the city “Apia” might be typed as “Capital” in a knowledge base, but not related to the country “Samoa”. However, if most of the other cities are related to countries via a role `isCapitalOf`, then the learning algorithm may still suggest `City` and `isCapitalOf` at least one `Country` since this describes the majority of capitals in the knowledge base well. If the knowledge engineer agrees with such a definition, then a tool can assist him in completing missing information about some capitals.

According to Occam’s razor [29] simple solutions of the learning problem are to be preferred over more complex ones, because they are more readable. This is even more important in the ontology engineering context, where it is essential to suggest simple expressions to the knowledge engineer. We measure simplicity as the *length* of an expression, which is defined in a straightforward way, namely as the sum of the numbers of concept, role, quantifier, and connective symbols occurring in the expression. The algorithm is biased towards shorter expressions. Also note that, for simplicity the definition of the learning problem itself does enforce coverage, but not prediction, i.e. correct classification of objects which are added to the knowledge base in the future. Concepts with high coverage and poor prediction are said to *overfit* the data. However, due to the strong bias towards short expressions this problem occurs empirically rarely in description logics [117].

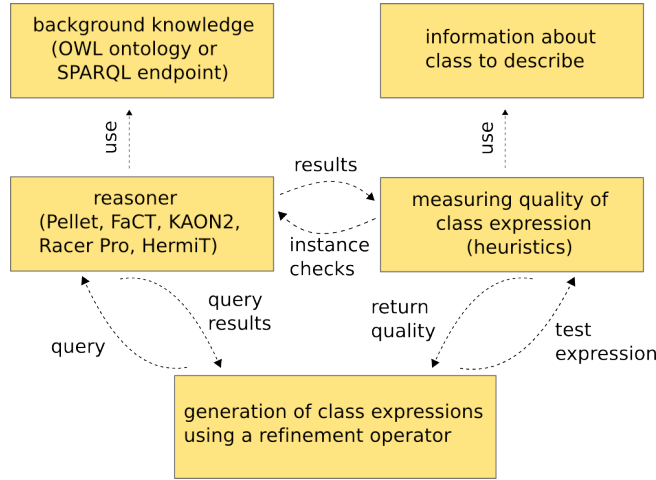


Fig. 26: Outline of the general learning approach in CELOE: One part of the algorithm is the generation of promising class expressions taking the available background knowledge into account. Another part is a heuristic measure of how close an expression is to being a solution of the learning problem. Figure adapted from [81].

Base Learning Algorithm Figure 26 gives a brief overview of the *CELOE* algorithm, which follows the common “generate and test” approach in ILP. This means that learning is seen as a search process and several class expressions are generated and tested against a background knowledge base. Each of those class expressions is evaluated using a heuristic, which is described in the next section. A challenging part of a learning algorithm is to decide which expressions to test. In particular, such a decision should take the computed heuristic values and the structure of the background knowledge into account. For *CELOE*, we use the approach described in [116,117] as base, where this problem has already been analysed, implemented, and evaluated in depth. It is based on the idea of *refinement operators*:

Definition 8 (refinement operator). A quasi-ordering is a reflexive and transitive relation. In a quasi-ordered space (S, \leq) a downward (upward) refinement operator ρ is a mapping from S to 2^S , such that for any $C \in S$ we have that $C' \in \rho(C)$ implies $C' \leq C$ ($C \leq C'$). C' is called a specialisation (generalisation) of C .

Refinement operators can be used for searching in the space of expressions. As ordering we can use subsumption. (Note that the subsumption relation \sqsubseteq is a quasi-ordering.) If an expression C subsumes an expression D ($D \sqsubseteq C$), then C will cover all examples which are covered by D . This makes subsumption a suitable order for searching in expressions as it allows to prune parts of the search space without losing possible solutions.

The approach we used is a top-down algorithm based on refinement operators as illustrated in Figure 27 (more detailed schemata can be found in the slides³¹ of the ontology learning lecture of Reasoning Web 2010 [110]). This means that the first class expression, which will be tested is the most general expression (\top), which is then mapped

³¹ <http://reasoningweb.org/2010/teaching-material/lehmann.pdf>

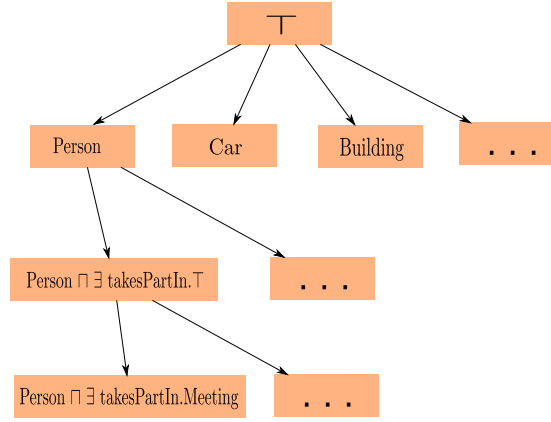


Fig. 27: Illustration of a search tree in a top down refinement approach.

to a set of more specific expressions by means of a downward refinement operator. Naturally, the refinement operator can be applied to the obtained expressions again, thereby spanning a *search tree*. The search tree can be pruned when an expression does not cover sufficiently many instances of the class A we want to describe. One example for a path in a search tree spanned up by a downward refinement operator is the following (\rightsquigarrow denotes a refinement step):

$$\begin{aligned} T &\rightsquigarrow \text{Person} \rightsquigarrow \text{Person} \sqcap \text{takesPartIn}.T \\ &\rightsquigarrow \text{Person} \sqcap \text{takesPartIn.Meeting} \end{aligned}$$

The heart of such a learning strategy is to define a suitable refinement operator and an appropriate search heuristics for deciding which nodes in the search tree should be expanded. The refinement operator in the considered algorithm is defined in [117]. It is based on earlier work in [116] which in turn is built on the theoretical foundations of [115]. It has been shown to be the best achievable operator with respect to a set of properties (not further described here), which are used to assess the performance of refinement operators. The learning algorithm supports conjunction, disjunction, negation, existential and universal quantifiers, cardinality restrictions, hasValue restrictions as well as boolean and double datatypes.

6.3 Finding a Suitable Heuristic

A heuristic measures how well a given class expression fits a learning problem and is used to guide the search in a learning process. To define a suitable heuristic, we first need to address the question of how to measure the accuracy of a class expression. We introduce several heuristics, which can be used for *CELOE* and later evaluate them.

We cannot simply use supervised learning from examples directly, since we do not have positive and negative examples available. We can try to tackle this problem by using the existing instances of the class as positive examples and the remaining instances

as negative examples. This is illustrated in Figure 28, where \mathcal{K} stands for the knowledge base and A for the class to describe. We can then measure accuracy as the number of correctly classified examples divided by the number of all examples. This can be computed as follows for a class expression C and is known as *predictive accuracy* in Machine Learning:

$$predacc(C) = 1 - \frac{|R(A) \setminus R(C)| + |R(C) \setminus R(A)|}{n} \quad n = |Ind(\mathcal{K})|$$

Here, $Ind(\mathcal{K})$ stands for the set of individuals occurring in the knowledge base. $R(A) \setminus R(C)$ are the false negatives whereas $R(C) \setminus R(A)$ are false positives. n is the number of all examples, which is equal to the number of individuals in the knowledge base in this case. Apart from learning definitions, we also want to be able to learn super class axioms ($A \sqsubseteq C$). Naturally, in this scenario $R(C)$ should be a superset of $R(A)$. However, we still do want $R(C)$ to be as small as possible, otherwise \top would always be a solution. To reflect this in our accuracy computation, we penalise false negatives more than false positives by a factor of t ($t > 1$) and map the result to the interval $[0, 1]$:

$$predacc(C, t) = 1 - 2 \cdot \frac{t \cdot |R(A) \setminus R(C)| + |R(C) \setminus R(A)|}{(t + 1) \cdot n} \quad n = |Ind(\mathcal{K})|$$

While being straightforward, the outlined approach of casting class learning into a standard learning problem with positive and negative examples has the disadvantage that the number of negative examples will usually be much higher than the number of positive examples. As shown in Table 6, this may lead to overly optimistic estimates. More importantly, this accuracy measure has the drawback of having a dependency on the number of instances in the knowledge base.

Therefore, we investigated further heuristics, which overcome this problem, in particular by transferring common heuristics from information retrieval to the class learning problem:

1. *F-Measure*: F_β -Measure is based on *precision* and *recall* weighted by β . They can be computed for the class learning problem without having negative examples. Instead, we perform a retrieval for the expression C , which we want to evaluate. We can then define precision as the percentage of instances of C , which are also instances of A and recall as percentage of instances of A , which are also instances of C . This is visualised in Figure 28. F-Measure is defined as harmonic mean of precision and recall. For learning super classes, we use F_3 measure by default, which gives recall a higher weight than precision.
2. *A-Measure*: We denote the arithmetic mean of precision and recall as A-Measure. Super class learning is achieved by assigning a higher weight to recall. Using the arithmetic mean of precision and recall is uncommon in Machine Learning, since it results in too optimistic estimates. However, we found that it is useful in super class learning, where F_n is often too pessimistic even for higher n .
3. *Generalised F-Measure*: Generalised F-Measure has been published in [48] and extends the idea of F-measure by taking the three valued nature of classification in OWL/DLs into account: An individual can either belong to a class, the negation of a class or none of both cases can be proven. This differs from common binary

classification tasks and, therefore, appropriate measures have been introduced (see [48] for details). Adaption for super class learning can be done in a similar fashion as for F-Measure itself.

4. *Jaccard Distance*: Since $R(A)$ and $R(C)$ are sets, we can use the well-known Jaccard coefficient to measure the similarity between both sets.

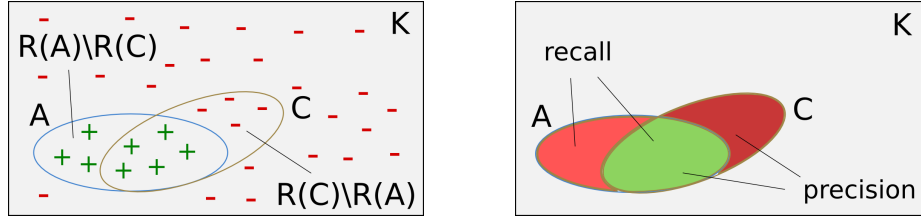


Fig. 28: Visualisation of different accuracy measurement approaches. \mathcal{K} is the knowledge base, A the class to describe and C a class expression to be tested. Left side: Standard supervised approach based on using positive (instances of A) and negative (remaining instances) examples. Here, the accuracy of C depends on the number of individuals in the knowledge base. Right side: Evaluation based on two criteria: recall (Which fraction of $R(A)$ is in $R(C)$?) and precision (Which fraction of $R(C)$ is in $R(A)$?).

We argue that those four measures are more appropriate than predictive accuracy when applying standard learning algorithms to the ontology engineering use case. Table 6 provides some example calculations, which allow the reader to compare the different heuristics.

Efficient Heuristic Computation Several optimisations for computing the heuristics are described in [111]. In particular, adapted approximate reasoning and stochastic approximations are discussed. Those improvements have shown to lead to order of magnitude gains in efficiency for many ontologies. We refrain from describing those methods in this chapter.

The Protégé Plugin After implementing and testing the described learning algorithm, we integrated it into *Protégé* and *OntoWiki*. Together with the Protégé developers, we extended the Protégé 4 plugin mechanism to be able to seamlessly integrate the DL-Learner plugin as an additional method to create class expressions. This means that the knowledge engineer can use the algorithm exactly where it is needed without any additional configuration steps. The plugin has also become part of the official Protégé 4 repository, i.e. it can be directly installed from within Protégé.

A screenshot of the plugin is shown in Figure 29. To use the plugin, the knowledge engineer is only required to press a button, which then starts a new thread in the background. This thread executes the learning algorithm. The used algorithm is an *anytime algorithm*, i.e. at each point in time we can always see the currently best suggestions.

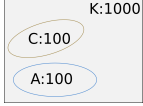
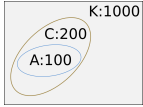
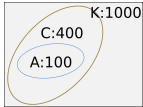
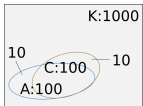
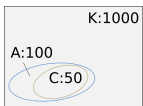
illustration	pred. acc.		F-Measure		A-Measure		Jaccard
	eq	sc	eq	sc	eq	sc	
	80%	67%	0%	0%	0%	0%	0%
	90%	92%	67%	73%	75%	88%	50%
	70%	75%	40%	48%	63%	82%	25%
	98%	97%	90%	90%	90%	90%	82%
	95%	88%	67%	61%	75%	63%	50%

Table 6: Example accuracies for selected cases (eq = equivalence class axiom, sc = super class axiom). The images on the left represent an imaginary knowledge base \mathcal{K} with 1000 individuals, where we want to describe the class A by using expression C . It is apparent that using predictive accuracy leads to impractical accuracies, e.g. in the first row C cannot possibly be a good description of A , but we still get 80% accuracy, since all the negative examples outside of A and C are correctly classified.

The GUI updates the suggestion list each second until the maximum runtime – 10 seconds by default – is reached. This means that the perceived runtime, i.e. the time after which only minor updates occur in the suggestion list, is often only one or two seconds for small ontologies. For each suggestion, the plugin displays its accuracy.

When clicking on a suggestion, it is visualized by displaying two circles: One stands for the instances of the class to describe and another circle for the instances of the suggested class expression. Ideally, both circles overlap completely, but in practice this will often not be the case. Clicking on the plus symbol in each circle shows its list of individuals. Those individuals are also presented as points in the circles and moving the mouse over such a point shows information about the respective individual. Red points show potential problems detected by the plugin. Please note that we use closed world reasoning to detect those problems. For instance, in our initial example, a capital which is not related via the property `isCapitalOf` to an instance of `Country` is marked red. If there is not only a potential problem, but adding the expression would render the ontology inconsistent, the suggestion is marked red and a warning message is displayed.

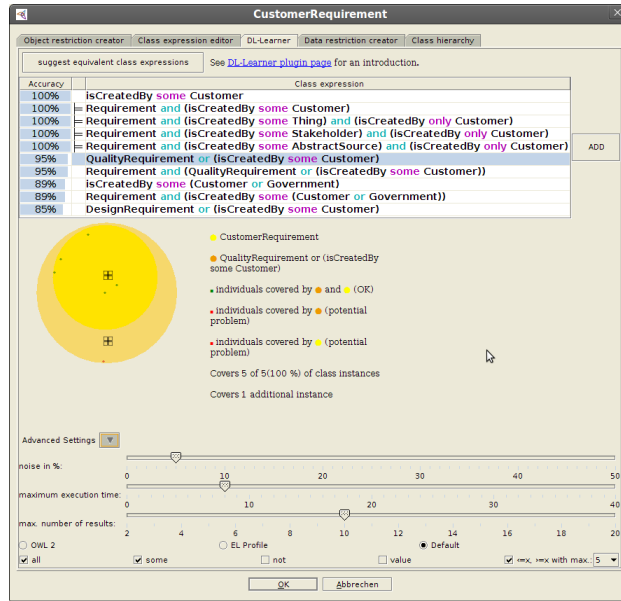


Fig. 29: A screenshot of the DL-Learner Protégé plugin. It is integrated as additional tab to create class expressions in Protégé. The user is only required to press the “suggest equivalent class expressions” button and within a few seconds they will be displayed ordered by accuracy. If desired, the knowledge engineer can visualize the instances of the expression to detect potential problems. At the bottom, optional expert configuration settings can be adopted.

Accepting such a suggestion can still be a good choice, because the problem often lies elsewhere in the knowledge base, but was not obvious before, since the ontology was not sufficiently expressive for reasoners to detect it. This is illustrated by a screencast available from the plugin homepage,³² where the ontology becomes inconsistent after adding the axiom, and the real source of the problem is fixed afterwards. Being able to make such suggestions can be seen as a strength of the plugin.

The plugin allows the knowledge engineer to change expert settings. Those settings include the maximum suggestion search time, the number of results returned and settings related to the desired target language, e.g. the knowledge engineer can choose to stay within the OWL 2 EL profile or enable/disable certain class expression constructors. The learning algorithm is designed to be able to handle noisy data and the visualisation of the suggestions will reveal false class assignments so that they can be fixed afterwards.

The OntoWiki Plugin Analogous to Protégé, we created a similar plugin for OntoWiki (cf. section 4). OntoWiki is a lightweight ontology editor, which allows distributed and collaborative editing of knowledge bases. It focuses on wiki-like, simple and intuitive authoring of semantic content, e.g. through inline editing of RDF content, and provides different views on instance data.

Recently, a fine-grained plugin mechanism and extensions architecture was added to OntoWiki. The DL-Learner plugin is technically realised by implementing an OntoWiki component, which contains the core functionality, and a module, which implements the

³² <http://dl-learner.org/wiki/ProtegePlugin>

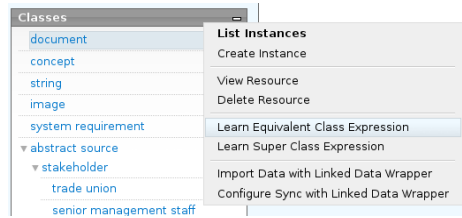


Fig. 30: The DL-Learner plugin can be invoked from the context menu of a class in OntoWiki.

UI embedding. The DL-Learner plugin can be invoked from several places in OntoWiki, for instance through the context menu of classes as shown in Figure 30.

The plugin accesses DL-Learner functionality through its WSDL-based web service interface. Jar files containing all necessary libraries are provided by the plugin. If a user invokes the plugin, it scans whether the web service is online at its default address. If not, it is started automatically.

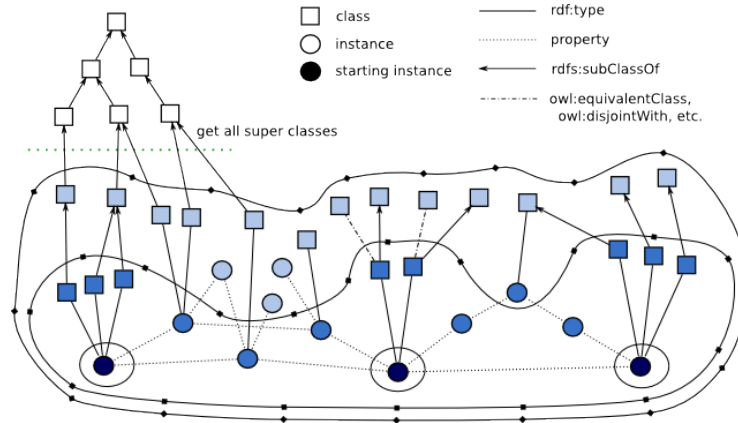


Fig. 31: Extraction with three starting instances. The circles represent different recursion depths. The circles around the starting instances signify recursion depth 0. The larger inner circle represents the fragment with recursion depth 1 and the largest outer circle with recursion depth 2. Figure taken from [81].

A major technical difference compared to the Protégé plugin is that the knowledge base is accessed via SPARQL, since OntoWiki is a SPARQL-based web application. In Protégé, the current state of the knowledge base is stored in memory in a Java object. As a result, we cannot easily apply a reasoner on an OntoWiki knowledge base. To overcome this problem, we use the DL-Learner fragment selection mechanism described in [81,82,40]. Starting from a set of instances, the mechanism extracts a relevant fragment from the underlying knowledge base up to some specified recursion depth. Figure 31

provides an overview of the fragment selection process. The fragment has the property that learning results on it are similar to those on the complete knowledge base. For a detailed description we refer the reader to the full article.

The fragment selection is only performed for medium to large-sized knowledge bases. Small knowledge bases are retrieved completely and loaded into the reasoner. While the fragment selection can cause a delay of several seconds before the learning algorithm starts, it also offers flexibility and scalability. For instance, we can learn class expressions in large knowledge bases such as DBpedia in OntoWiki.³³

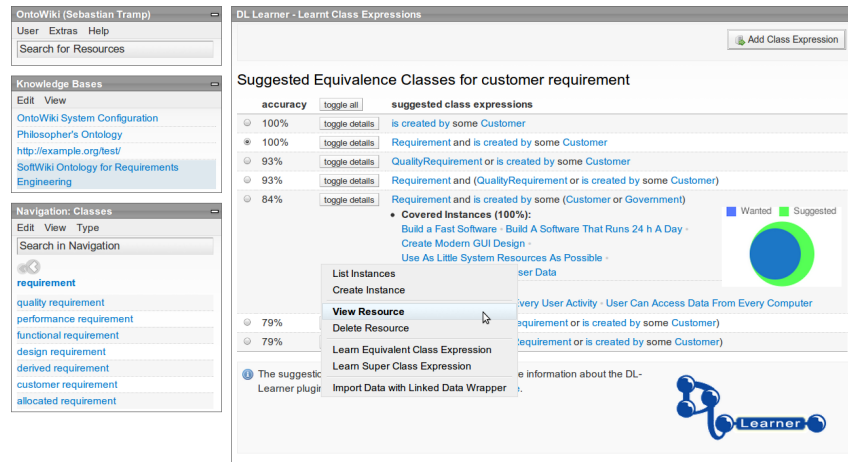


Fig. 32: Screenshot of the result table of the DL-Learner plugin in OntoWiki.

Figure 32 shows a screenshot of the OntoWiki plugin applied to the SWORE [170] ontology. Suggestions for learning the class “customer requirement” are shown in Manchester OWL Syntax. Similar to the Protégé plugin, the user is presented a table of suggestions along with their accuracy value. Additional details about the instances of “customer requirement”, covered by a suggested class expressions and additionally contained instances can be viewed via a toggle button. The modular design of OntoWiki allows rich user interaction: Each resource, e.g. a class, property, or individual, can be viewed and subsequently modified directly from the result table as shown for “design requirement” in the screenshot. For instance, a knowledge engineer could decide to import additional information available as Linked Data and run the CELOE algorithm again to see whether different suggestions are provided with additional background knowledge.

Evaluation To evaluate the suggestions made by our learning algorithm, we tested it on a variety of real-world ontologies of different sizes and domains. Please note that we

³³ OntoWiki is undergoing an extensive development, aiming to support handling such large knowledge bases. A release supporting this is expected for the first half of 2012.

intentionally do not perform an evaluation of the machine learning technique as such on existing benchmarks, since we build on the base algorithm already evaluated in detail in [117]. It was shown that this algorithm is superior to other supervised learning algorithms for OWL and at least competitive with the state of the art in ILP. Instead, we focus on its use within the ontology engineering scenario. The goals of the evaluation are to 1. determine the influence of reasoning and heuristics on suggestions, 2. to evaluate whether the method is sufficiently efficient to work on large real-world ontologies.

To perform the evaluation, we wrote a dedicated plugin for the Protégé ontology editor. This allows the evaluators to browse the ontology while deciding whether the suggestions made are reasonable. The plugin works as follows: First, all classes with at least 5 inferred instances are determined. For each such class, we run CELOE with different settings to generate suggestions for definitions. Specifically, we tested two reasoners and five different heuristics. The two reasoners are standard Pellet and Pellet combined with approximate reasoning (not described in detail here). The five heuristics are those described in Section 6.3. For each configuration of CELOE, we generate at most 10 suggestions exceeding a heuristic threshold of 90%. Overall, this means that there can be at most $2 * 5 * 10 = 100$ suggestions per class – usually less, because different settings of CELOE will still result in similar suggestions. This list is shuffled and presented to the evaluators. For each suggestion, the evaluators can choose between 6 options (see Table 8):

- 1 The suggestion improves the ontology (improvement),
- 2 The suggestion is no improvement and should not be included (not acceptable) and
- 3 Adding the suggestion would be a modelling error (error).

In the case of existing definitions for class *A*, we removed them prior to learning. In this case, the evaluator could choose between three further options:

- 4 The learned definition is equal to the previous one and both are good (equal +),
- 5 The learned definition is equal to the previous one and both are bad (equal -) and
- 6 The learned definition is inferior to the previous one (inferior).

We used the default settings of CELOE, e.g. a maximum execution time of 10 seconds for the algorithm. The knowledge engineers were five experienced members of our research group, who made themselves familiar with the domain of the test ontologies. Each researcher worked independently and had to make 998 decisions for 92 classes between one of the options. The time required to make those decisions was approximately 40 working hours per researcher. The raw agreement value of all evaluators is 0.535 (see e.g. [5] for details) with 4 out of 5 evaluators in strong pairwise agreement (90%). The evaluation machine was a notebook with a 2 GHz CPU and 3 GB RAM.

³⁴ <http://www.mindswap.org/ontologies/SC.owl>

³⁵ <http://www.sbcny.org/datasets/adhesome.owl>

³⁶ <http://i2geo.net/ontologies/current/GeoSkills.owl>

³⁷ <http://www.co-ode.org/ontologies/eukariotic/2005/06/01/eukariotic.owl>

³⁸ <http://acl.icnet.uk/%7Emw/MDM0.73.owl>

³⁹ <http://reliant.teknowledge.com/DAML/Economy.owl>

⁴⁰ <http://www.ecs.soton.ac.uk/~aoj04r/resist.owl>

⁴¹ <http://www.fadyart.com/Finance.owl>

⁴² <http://sweet.jpl.nasa.gov/1.1/earthrealm.owl>

Ontology	#logical axioms	#classes	#object properties	#data properties	#individuals	DL expressivity
SC Ontology ³⁴	20081	28	8	5	3542	$\mathcal{AL}(\mathcal{D})$
Adhesome ³⁵	12043	40	33	37	2032	$\mathcal{ALCHN}(\mathcal{D})$
GeoSkills ³⁶	14966	613	23	21	2620	$\mathcal{ALCHOIN}(\mathcal{D})$
Eukariotic ³⁷	38	11	1	0	11	\mathcal{ALCON}
Breast Cancer ³⁸	878	196	22	3	113	$\mathcal{ALCROF}(\mathcal{D})$
Economy ³⁹	1625	339	45	8	482	$\mathcal{ALCH}(\mathcal{D})$
Resist ⁴⁰	239	349	134	38	75	$\mathcal{ALUF}(\mathcal{D})$
Finance ⁴¹	16014	323	247	74	2466	$\mathcal{ALCROIQ}(\mathcal{D})$
Earthrealm ⁴²	931	2364	215	36	171	$\mathcal{ALCHO}(\mathcal{D})$

Table 7: Statistics about test ontologies

reasoner/heuristic	improvement	equal quality (+)	equal quality (-)	inferior	not acceptable	error	missed improvements in %	selected position on suggestion list (incl. std. deviation)	avg. accuracy of selected suggestion in %
Pellet/F-Measure	16.70	0.44	0.66	0.00	64.66	17.54	14.95	2.82 ± 2.93	96.91
Pellet/Gen. F-Measure	15.24	0.44	0.66	0.11	66.60	16.95	16.30	2.78 ± 3.01	92.76
Pellet/A-Measure	16.70	0.44	0.66	0.00	64.66	17.54	14.95	2.84 ± 2.93	98.59
Pellet/pred. acc.	16.59	0.44	0.66	0.00	64.83	17.48	15.22	2.69 ± 2.82	98.05
Pellet/Jaccard	16.81	0.44	0.66	0.00	64.66	17.43	14.67	2.80 ± 2.91	95.26
Pellet FIC/F-Measure	36.30	0.55	0.55	0.11	52.62	9.87	1.90	2.25 ± 2.74	95.01
Pellet FIC/Gen. F-M.	33.41	0.44	0.66	0.00	53.41	12.09	7.07	1.77 ± 2.69	89.42
Pellet FIC/A-Measure	36.19	0.55	0.55	0.00	52.84	9.87	1.63	2.21 ± 2.71	98.65
Pellet FIC/pred. acc.	32.99	0.55	0.55	0.11	55.58	10.22	4.35	2.17 ± 2.55	98.92
Pellet FIC/Jaccard	36.30	0.55	0.55	0.11	52.62	9.87	1.90	2.25 ± 2.74	94.07

Table 8: Options chosen by evaluators aggregated by class. FIC stands for the fast instance checker, which is an approximate reasoning procedure.

Table 8 shows the evaluation results. All ontologies were taken from the Protégé OWL⁴³ and TONES⁴⁴ repositories. We randomly selected 5 ontologies comprising instance data from these two repositories, specifically the Earthrealm, Finance, Resist, Economy and Breast Cancer ontologies (see Table 7).

The results in Table 8 show which options were selected by the evaluators. It clearly indicates that the usage of approximate reasoning is sensible. The results are, however, more difficult to interpret with regard to the different employed heuristics. Using predictive accuracy did not yield good results and, surprisingly, generalised F-Measure also had a lower percentage of cases where option 1 was selected. The other three heuristics generated very similar results. One reason is that those heuristics are all based on precision and recall, but in addition the low quality of some of the randomly selected test ontologies posed a problem. In cases of too many very severe modelling errors, e.g. conjunctions and disjunctions mixed up in an ontology or inappropriate domain and range restrictions, the quality of suggestions decreases for each of the heuristics. This is the main reason why the results for the different heuristics are very close. Particularly, generalised F-Measure can show its strengths mainly for properly designed ontologies. For instance, column 2 of Table 8 shows that it missed 7% of possible improvements. This means that for 7% of all classes, one of the other four heuristics was able to find an appropriate definition, which was not suggested when employing generalised F-Measure. The last column in this table shows that the average value of generalised F-Measure is quite low. As explained previously, it distinguishes between cases when an individual is instance of the observed class expression, its negation, or none of both. In many cases, the reasoner could not detect that an individual is instance of the negation of a class expression, because of the absence of disjointness axioms and negation in the knowledge base, which explains the low average values of generalised F-Measure. Column 4 of Table 8 shows that many selected expressions are amongst the top 5 (out of 10) in the suggestion list, i.e. providing 10 suggestions appears to be a reasonable choice.

In general, the improvement rate is only at about 35% according to Table 8 whereas it usually exceeded 50% in preliminary experiments with other real-world ontologies with fewer or less severe modelling errors. Since CELOE is based on OWL reasoning, it is clear that schema modelling errors will have an impact on the quality of suggestions. As a consequence, we believe that the CELOE algorithm should be combined with ontology debugging techniques. We have obtained first positive results in this direction and plan to pursue it in future work. However, the evaluation also showed that CELOE does still work in ontologies, which probably were never verified by an OWL reasoner.

Summary We presented the CELOE learning method specifically designed for extending OWL ontologies. Five heuristics were implemented and analysed in conjunction with CELOE along with several performance improvements. A method for approximating heuristic values has been introduced, which is useful beyond the ontology engineering scenario to solve the challenge of dealing with a large number of examples in ILP [203]. Furthermore, we biased the algorithm towards short solutions and implemented optimisations to increase readability of the suggestions made. The resulting

⁴³ http://protegewiki.stanford.edu/index.php/Protege_Ontology_Library

⁴⁴ <http://owl.cs.manchester.ac.uk/repository/>

algorithm was implemented in the open source DL-Learner framework. We argue that CELOE is the first ILP based algorithm, which turns the idea of learning class expressions for extending ontologies into practice. CELOE is integrated into two plugins for the ontology editors Protégé and OntoWiki and can be invoked using just a few mouse clicks.

7 Linked Data Quality

Linked Open Data (LOD) has provided, over the past several years, an unprecedented volume of structured data currently amount to 50 billion facts, represented as RDF triples. Although publishing large amounts of data on the Web is certainly a step in the right direction, the published data is only as useful as its quality. On the Data Web we have very varying quality of information covering various domains since data is merged together from different autonomous evolving data sources on the Web. For example, data extracted from semi-structured or even unstructured sources, such as DBpedia, often contains inconsistencies as well as mis-represented and incomplete information. Despite data quality in LOD being an essential concept, the autonomy and openness of the information providers makes the web vulnerable to missing, inaccurate, incomplete, inconsistent or outdated information.

Data quality is commonly conceived as *fitness for use* [97,202] for a certain application or use case. However, even datasets with quality problems might be useful for certain applications, as long as the quality is in the required range. For example, in the case of DBpedia the data quality is perfectly sufficient for enriching Web search with facts or suggestions about common sense information, such as entertainment topics. In such a scenario, DBpedia can be used to show related movies and personal information, when a user searches for an actor. In this case, it is rather neglectable, when in relatively few cases, a related movie or some personal fact is missing. For developing a medical application, on the other hand, the quality of DBpedia is probably completely insufficient. It should be noted that even the traditional, document-oriented Web has content of varying quality and is still perceived to be extremely useful by most people.

Consequently, one of the key challenges is to determine the quality of datasets published on the Web and making this quality information explicitly available. Assuring data quality is particularly a challenge in LOD as it involves a set of autonomously evolving data sources. Other than on the document Web, where information quality can be only indirectly (e.g. via page rank) or vaguely defined, there are much more concrete and measurable data quality metrics available for structured information such as accuracy of facts, completeness, adequacy of semantic representation or degree of understandability.

In this chapter, we first define the basic concepts of data quality (subsection 7.1), then report the formal definitions of a set of 18 different dimensions along with their respective metrics identified in [212] (subsection 7.2). Thereafter, we compare a set of currently available tools specially designed to assess the quality of Linked Data (subsection 7.7).

7.1 Data Quality Concepts

In this section, we introduce the basic concepts of data quality to help the readers understand these terminologies in their consequent usage.

Data Quality. The term *data quality* is commonly conceived as a multi-dimensional construct with a popular definition as the "fitness for use" [97]. In case of the Semantic Web, there are varying concepts of data quality such as the semantic metadata on the

one hand and the notion of link quality on the other. There are several characteristics of data quality that should be considered i.e. the completeness, accuracy, consistency and validity on the one hand and the representational consistency, conciseness as well as the timeliness, understandability, availability and verifiability on the other hand.

Data Quality Problems. A set of issues that can affect the potentiality of the applications that use the data are termed as data quality problems. The problems may vary from the incompleteness of data, inconsistency in representation, invalid syntax or inaccuracy.

Data Quality Dimensions and Metrics. Data quality assessment involves the measurement of quality *dimensions* (or *criteria*) that are relevant to the user. A data quality assessment *metric* (or *measure*) is a procedure for measuring an information quality dimension [26]. The metrics are basically heuristics designed to fit a specific assessment situation [120]. Since the dimensions are rather abstract concepts, the assessment metrics rely on quality *indicators* that can be used for the assessment of the quality of a data source w.r.t the criteria [59].

Data Quality Assessment Method. A data quality assessment methodology is the process of evaluating if a piece of data meets the information consumers need for a specific use case [26]. The process involves measuring the quality dimensions that are relevant to the user and comparing the assessment results with the users quality requirements.

7.2 Linked Data quality dimensions

In [212], a core set of 18 different data quality dimensions were reported that can be applied to assess the quality of Linked Data. These dimensions are divided into the following groups:

- Accessibility dimensions
- Intrinsic dimensions
- Contextual dimensions
- Representational dimensions

Figure 33 shows the classification of the dimensions into these four different groups as well as the relations between them.

Use case scenario. Since data quality is conceived as “fitness for use”, we introduce a specific use case that will allow us to illustrate the importance of each dimension with the help of an example. The use case is about an intelligent flight search engine, which relies on aggregating data from several datasets. The search engine obtains information about airports and airlines from an airline dataset (e.g. *OurAirports*⁴⁵, *OpenFlights*⁴⁶). Information about the location of countries, cities and particular addresses is obtained from a spatial dataset (e.g. *LinkedGeoData*⁴⁷). Additionally, aggregators pull all the

⁴⁵ <http://thedatahub.org/dataset/ourairports>

⁴⁶ <http://thedatahub.org/dataset/open-flights>

⁴⁷ linkedgeo.org

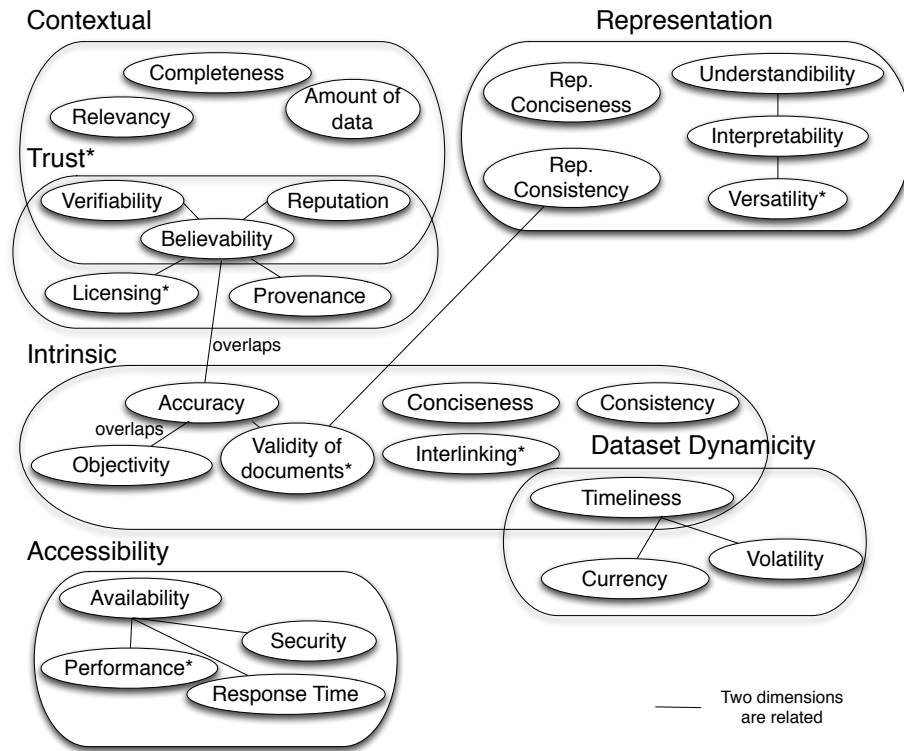


Fig. 33: Linked data quality dimensions and the relations between them [Source: [212].]

information related to flights from different booking services (e.g., *Expedia*⁴⁸) and represent this information as RDF. This allows a user to query the integrated dataset for a flight between any start and end destination for any time period. We will use this scenario throughout as an example to explain each quality dimension through a quality issue.

7.3 Accessibility dimensions

The dimensions belonging to this category involve aspects related to the access, authenticity and retrieval of data to obtain either the entire or some portion of the data (or from another linked dataset) for a particular use case. There are five dimensions that are part of this group, which are *availability*, *licensing*, *interlinking*, *security* and *performance*.

Availability. Bizer [25] adopted the definition of availability from Pipino et al. [164] as “the extent to which information is available, or easily and quickly retrievable”. Flemming [59] referred to availability as the proper functioning of all access methods. However, the definition by Pipino et al. is more related to the measurement of available

⁴⁸ <http://www.expedia.com/>

information rather than to the method of accessing the information as implied in the latter explanation by Flemming.

Definition 9 (Availability). *Availability of a dataset is the extent to which data (or some portion of it) is present, obtainable and ready for use.*

Metrics. The metrics identified for availability are:

- A1: checking whether the server responds to a SPARQL query [59]
- A2: checking whether an RDF dump is provided and can be downloaded [59]
- A3: detection of dereferencability of URIs by checking:
 - for dead or broken links [85], i.e. that when an HTTP-GET request is sent, the status code `404 Not Found` is not returned [59]
 - that useful data (particularly RDF) is returned upon lookup of a URI [85]
 - for changes in the URI, i.e. compliance with the recommended way of implementing redirections using the status code `303 See Other` [59]
- A4: detect whether the HTTP response contains the header field stating the appropriate content type of the returned file, e.g. `application/rdf+xml` [85]
- A5: dereferncability of all forward links: all available triples where the local URI is mentioned in the subject (i.e. the description of the resource) [86]

Example. Let us consider the case in which a user looks up a flight in our flight search engine. She requires additional information such as car rental and hotel booking at the destination, which is present in another dataset and interlinked with the flight dataset. However, instead of retrieving the results, she receives an error response code `404 Not Found`. This is an indication that the requested resource cannot be dereferenced and is therefore unavailable. Thus, with this error code, she may assume that either there is no information present at that specified URI or the information is unavailable.

Licensing. Licensing is a new quality dimensions not considered for relational databases but mandatory in the data world such as LD. Flemming [59] and Hogan et al. [86] both stated that each RDF document should contain a license under which the content can be (re-)used, in order to enable information consumers to use the data under clear legal terms. Additionally, the existence of a machine-readable indication (by including the specifications in a VoID⁴⁹ description) as well as a human-readable indication of a license are important not only for the permissions a licence grants but as an indication of which requirements the consumer has to meet [59]. Although both these studies do not provide a formal definition, they agree on the use and importance of licensing in terms of data quality.

Definition 10 (Licensing). *Licensing is defined as the granting of permission for a consumer to re-use a dataset under defined conditions.*

Metrics. The metrics identified for licensing are:

⁴⁹ <http://vocab.deri.ie/void>

- L1: machine-readable indication of a license in the VoID description or in the dataset itself [59,86]
- L2: human-readable indication of a license in the documentation of the dataset [59,86]
- L3: detection of whether the dataset is attributed under the same license as the original [59]

Example. Since our flight search engine aggregates data from several existing data sources, a clear indication of the license allows the search engine to re-use the data from the airlines websites. For example, the LinkedGeoData dataset is licensed under the Open Database License⁵⁰, which allows others to copy, distribute and use the data and produce work from the data allowing modifications and transformations. Due to the presence of this specific license, the flight search engine is able to re-use this dataset to pull geo-spatial information and feed it to the search engine.

Interlinking. Interlinking is a relevant dimension in LD since it supports data integration. Interlinking is provided by RDF triples that establish a link between the entity identified by the subject with the entity identified by the object. Through the typed RDF links, data items are effectively interlinked. The importance of interlinking, also known as “mapping coherence” can be classified in one of the four scenarios: (i) Frameworks; (ii) Terminological Reasoning; (iii) Data Transformation; (iv) Query Processing, as identified in [133]. Even though the core articles in this survey do not contain a formal definition for interlinking, they provide metrics on how to measure this dimension.

Definition 11 (Interlinking). *Interlinking refers to the degree to which entities that represent the same concept are linked to each other, be it within or between two or more data sources.*

Metrics. The metrics identified for interlinking are:

- I1: detection of:
 - interlinking degree: how many hubs there are in a network⁵¹ [73]
 - clustering coefficient: how dense is the network [73]
 - centrality: indicates the likelihood of a node being on the shortest path between two other nodes [73]
 - whether there are open sameAs chains in the network [73]
 - how much value is added to the description of a resource through the use of sameAs edges [73]
- I2: detection of the existence and usage of external URIs (e.g. using `owl:sameAs` links) [85,86]
- I3: detection of all local in-links or back-links: all triples from a dataset that have the resource’s URI as the object [86]

⁵⁰ <http://opendatacommons.org/licenses/odbl/>

⁵¹ In [73], a network is a set of facts provided by the graph of the Web of Data, excluding the blank nodes.

Example. In our flight search engine, the instance of the country “United States” in the airline dataset should be interlinked with the instance “America” in the spatial dataset. This interlinking can help when a user queries for a flight, as the search engine can display the correct route from the start destination to the end destination by correctly combining information for the same country from both the datasets. Since names of various entities can have different URIs in different datasets, their interlinking can help in disambiguation.

Security. Flemming [59] referred to security as “the possibility to restrict access to the data and to guarantee the confidentiality of the communication between a source and its consumers”. Additionally, Flemming referred to the verifiability dimension as the mean a consumer is provided with to examine the data for correctness. Bizer [25] adopted the definition of verifiability from Naumann et al. [143] as the “degree and ease with which the information can be checked for correctness”. Without such means, the assurance of the correctness of the data would come from the consumer’s trust in that source. It can be observed here that on the one hand Naumann et al. provided a formal definition whereas Flemming described the dimension by providing its advantages and metrics. Thus, security and verifiability point towards the same quality dimension i.e. to avoid alterations of the dataset and verify its correctness.

Definition 12 (Security). *Security is the extent to which data is protected against alteration and misuse.*

Metrics. The metrics identified for security are:

- S1: degree of using digital signatures to sign documents containing an RDF serialization, a SPARQL result set or signing an RDF graph [37,59]
- S2: verifying authenticity of the dataset based on provenance information such as the author and his contributors, the publisher of the data and its sources, if present in the dataset [59]

Example: In our use case, if we assume that the flight search engine obtains flight information from arbitrary airline websites, there is a risk for receiving incorrect information from malicious websites. For instance, an airline or sales agency website can pose as its competitor and display incorrect expensive flight fares. Thus, by this spoofing attack, this airline can prevent users to book with the competitor. In that case, the use of digital signatures for published RDF data allows to verify the identity of the publisher.

Performance. Performance is a dimension that has an influence on the quality of the information system or search engine, however not on the data set itself. Flemming [59] states that “the performance criterion comprises aspects of enhancing the performance of a source as well as measuring of the actual values”. Flemming [59] gave a general description of performance without explaining the meaning while Hogan et al. [86] described the issues related to performance. Moreover, Bizer [25], defined response-time as “the delay between submission of a request by the user and reception of the

response from the system”. Thus, response-time and performance point towards the same quality dimension.

Definition 13 (Performance). *Performance refers to the efficiency of a system that binds to a large dataset, that is, the more performant a data source is the more efficiently a system can process data.*

Metrics. The metrics identified for performance are:

- P1: checking for usage of slash-URIs where large amounts of data is provided⁵² [59]
- P2: low latency⁵³: (minimum) delay between submission of a request by the user and reception of the response from the system [25,59]
- P3: high throughput: (maximum) number of answered HTTP-requests per second [59]
- P4: scalability - detection of whether the time to answer an amount of ten requests divided by ten is not longer than the time it takes to answer one request [59]

Example. In our use case, the performance may depend on the type and complexity of the query by a large number of users. Our flight search engine can perform well by considering response-time when deciding which sources to use to answer a query.

Intra-relations The dimensions in this group are related with each other as follows: performance (response-time) of a system is related to the availability dimension. Only if a dataset is available and has low response time, it can perform well. Also, interlinking is related to availability because only if a dataset is available, it can be interlinked and the interlinks can be traversed. Additionally, the dimensions security and licensing are related since providing a license and specifying conditions for re-use helps secure the dataset against alterations and misuse.

7.4 Intrinsic dimensions

Intrinsic dimensions are those that are independent of the user’s context. There are five dimensions that are part of this group, which are *syntactic validity*, *semantic accuracy*, *consistency*, *conciseness* and *completeness*. These dimensions focus on whether information correctly (syntactically and semantically), compactly and completely represents the real world data and whether information is logically consistent in itself.

Syntactic validity. Flemming [59] defined the term validity of documents as “the valid usage of the underlying vocabularies and the valid syntax of the documents”. Fürber et al. [64] classified accuracy into syntactic and semantic accuracy. He explained that a “value is syntactically accurate, when it is part of a legal value set for the represented domain or it does not violate syntactical rules defined for the domain”. We associate

⁵² <http://www.w3.org/wiki/HashVsSlash>

⁵³ Latency is the amount of time from issuing the query until the first information reaches the user [143].

the validity of documents defined by Flemming to syntactic validity. We distinguish between the two types of accuracy defined by Fürber et al. and form two dimensions: *Syntactic validity* (syntactic accuracy) and *Semantic accuracy*. Additionally, Hogan et al. [85] identify syntax errors such as RDF/XML syntax errors, malformed datatype literals and literals incompatible with datatype range, which we associate with syntactic validity.

Definition 14 (Syntactic validity). *Syntactic validity is defined as the degree to which an RDF document conforms to the specification of the serialization format.*

Metrics. The metrics identified for syntactic validity are:

- SV1: detecting syntax errors using validators [59,85]
- SV2: detecting use of:
 - explicit definition of the allowed values for a certain datatype [64]
 - syntactic rules (type of characters allowed and/or the pattern of literal values) [64]
- SV3: detection of ill-typed literals which do not abide by the lexical syntax for their respective datatype that can occur if a value is (i) malformed or (ii) is a member of an incompatible datatype [85]

Example. In our use case, let us assume that the ID of the flight between Paris and New York is A123 while in our search engine the same flight instance is represented as A231. Since this ID is included in one of the datasets, it is considered to be syntactically accurate since it is a valid ID (even though it is incorrect).

Semantic accuracy. Bizer [25] adopted the definition of accuracy from Wang et al. [200] as the “degree of correctness and precision with which information in an information system represents states of the real world”. Furthermore, Fürber et al. [64] classified accuracy into syntactic and semantic accuracy. He explained that values are semantically accurate when they represent the correct state of an object. Based on this definition, we also considered the problems of *spurious annotation* and *inaccurate annotation* (inaccurate labeling and inaccurate classification) identified in Lei et al. [119] related to the semantic accuracy dimension.

Definition 15 (Semantic accuracy). *Semantic accuracy is defined as the degree to which data values correctly represent the real world facts.*

Metrics. The metrics identified for semantic accuracy are:

- SA1: detection of outliers by using distance-based, deviations-based and distribution-based methods [26]
- SA2: detection of inaccurate values by using functional dependency rules⁵⁴ [64,209]

⁵⁴ Functional dependencies are dependencies between the values of two or more different properties.

- SA3: detection of inaccurate annotations⁵⁵, labellings⁵⁶ or classifications⁵⁷ using the formula:

$$1 - \frac{\text{inaccurate instances}}{\text{total no. of instances}} * \frac{\text{balanced distance metric}}{\text{total no. of instances}} \text{ }^{58} \text{ [119]}$$
- SA4: verifying correctness of the dataset with the help of unbiased trusted third party (humans) [25]
- SA5: detection of misuse of properties⁵⁹ by using profiling statistics, which support the detection of discordant values or misused properties and facilitate to find valid values for specific properties [30]
- SA6: ratio of the number of semantically valid rules⁶⁰ to the number of nontrivial rules⁶¹ [39]

Example. In our use case, let us assume that the ID of the flight between Paris and New York is A123 while in our search engine the same flight instance is represented as A231. In this case, the instance is semantically inaccurate since the flight ID does not represent its real-world state i.e. A123.

Consistency. Bizer [25] adopted the definition of consistency from Mecella et al., [132] as when “two or more values do not conflict with each other”. Similarly, Hogan et al. [85] defined consistency as “no contradictions in the data”. Another definition was given by Mendes et al. [134] where “a dataset is consistent if it is free of conflicting information”. Additionally, Böhm et al. [30] and Mostafavi et al. [139] present metrics to assess consistency. However, it should be noted that for some languages such as OWL DL, there are clearly defined semantics, including clear definitions what inconsistency means. In description logics, model based semantics are used: A knowledge base is a set of axioms. A model is an interpretation, which satisfies all axioms in the knowledge base. A knowledge base is consistent if and only if it has a model [16].

Definition 16 (Consistency). *Consistency means that a knowledge base is free of (logical/formal) contradictions with respect to particular knowledge representation and inference mechanisms.*

Metrics. A straightforward way to check for consistency is to load the knowledge base into a reasoner and check whether it is consistent. However, for certain knowledge bases (e.g. very large or inherently inconsistent ones) this approach is not feasible. For such cases, specific aspects of consistency can be checked individually. Some important metrics identified in the literature are:

⁵⁵ where an instance of the semantic metadata set can be mapped back to more than one real world object or in other cases, where there is no object to be mapped back to an instance.

⁵⁶ where mapping from the instance to the object is correct but not properly labeled.

⁵⁷ in which the knowledge of the source object has been correctly identified by not accurately classified.

⁵⁸ Balanced distance metric is an algorithm that calculates the distance between the extracted (or learned) concept and the target concept [129]

⁵⁹ Properties are often misused when no applicable property exists.

⁶⁰ valid rules are generated from the real data and validated against a set of principles specified in the semantic network

⁶¹ The intuition is that the larger a dataset is, the more closely it should reflect the basic domain principles and the less semantically incorrect rules will be generated.

- CS1: detection of use of entities as members of disjoint classes using the formula: $\frac{\text{no. of entities described as members of disjoint classes}}{\text{total no. of entities described in the dataset}}$ [85]
- CS2: detection of misplaced classes or properties⁶² using entailment rules that indicate the position of a term in a triple [85]
- CS3: detection of misuse of `owl:DatatypeProperty` or `owl:ObjectProperty` through the ontology maintainer⁶³ [85]
- CS4: detection of use of members of `owl:DeprecatedClass` or `owl:DeprecatedProperty` through the ontology maintainer or by specifying manual mappings from deprecated terms to compatible terms [85]
- CS5: detection of bogus `owl:InverseFunctionalProperty` values by checking the uniqueness and validity of the inverse-functional values [85]
- CS6: detection of the re-definition by third parties of external classes/ properties (ontology hijacking) such that reasoning over data using those external terms is affected [85]
- CS7: detection of negative dependencies/correlation among properties using association rules [30]
- CS8: detection of inconsistencies in spatial data through semantic and geometric constraints [139]

Example. Let us assume a user looking for flights between Paris and New York on the 21st of December, 2013. Her query returns the following results:

Flight	From	To	Arrival	Departure
A123	Paris	NewYork	14:50	22:35
A123	Paris	London	14:50	22:35

The results show that the flight number A123 has two different destinations⁶⁴ at the same date and same time of arrival and departure, which is inconsistent with the ontology definition that one flight can only have one destination at a specific time and date. This contradiction arises due to inconsistency in data representation, which is detected by using inference and reasoning.

Conciseness. Mendes et al. [134] classified conciseness into schema and instance level conciseness. On the schema level (intensional), “a dataset is concise if it does not contain redundant attributes (two equivalent attributes with different names)”. Thus, intensional conciseness measures the number of unique schema elements (i.e. properties and classes) of a dataset in relation to the overall number of schema elements in a schema. On the data (instance) level (extensional), “a dataset is concise if it does not contain redundant objects (two equivalent objects with different identifiers)”. Thus, extensional conciseness measures the number of unique objects in relation to the overall number of objects in the dataset. This definition of conciseness is very similar to the definition of ‘uniqueness’ defined by Fürber et al. [64] as the “degree to which data is free of redundancies, in breadth, depth and scope”. This comparison shows that uniqueness and

⁶² For example, a URI defined as a class is used as a property or vice-a-versa.

⁶³ For example, attribute properties used between two resources and relation properties used with literal values.

⁶⁴ Under the assumption that we can infer that `NewYork` and `London` are different entities or, alternatively, make the unique names assumption.

conciseness point to the same dimension. Redundancy occurs when there are *equivalent* schema elements with different names/identifiers (in case of intensional conciseness) and when there are *equivalent* objects (instances) with different identifiers (in case of extensional conciseness) in a dataset.

Definition 17 (Conciseness). *Conciseness refers to the minimization of redundancy of entities at the schema and the data level. Conciseness is classified into (i) intensional conciseness (schema level) which refers to the case when the data does not contain redundant schema elements (properties and classes) and (ii) extensional conciseness (data level) which refers to the case when the data does not contain redundant objects (instances).*

Metrics. The metrics identified for conciseness are:

- CN1: intensional conciseness measured by $\frac{\text{no. of unique properties/classes of a dataset}}{\text{total no. of properties/classes in a target schema}}$ [134]
- CN2: extensional conciseness measured by:
 - $\frac{\text{no. of unique instances of a dataset}}{\text{total number of instances representations in the dataset}}$ [134] or
 - $1 - \frac{\text{total no. of instances that violate the uniqueness rule}}{\text{total no. of relevant instances}}$ [64,119]
- CN3: detection of unambiguous annotations using the formula:

$$1 - \frac{\text{no. of ambiguous instances}}{\text{no. of the instances contained in the semantic metadata set}}$$
⁶⁵ [119]

Example. In our flight search engine, an example of intensional conciseness would be a particular flight, say A123, being represented by two different properties in the same dataset, such as <http://flights.org/airlineID> and <http://flights.org/name>. This redundancy (‘airlineID’ and ‘name’ in this case) can ideally be solved by fusing the two properties and keeping only one unique identifier. On the other hand, an example of extensional conciseness is when both these identifiers of the same flight have the same information associated with them in both the datasets, thus duplicating the information.

Completeness. Bizer [25] adopted the definition of completeness from Pipino et al. [164] as “the degree to which information is not missing”. Fürber et al. [64] further classified completeness into: (i) Schema completeness, which is the degree to which classes and properties are not missing in a schema; (ii) Column completeness, which is a function of the missing property values for a specific property/column; and (iii) Population completeness, which refers to the ratio between classes represented in an information system and the complete population. Mendes et al. [134] distinguish completeness on the schema and the data level. On the schema level, a dataset is complete if it contains all of the attributes needed for a given task. On the data (i.e. instance) level, a dataset is complete if it contains all of the necessary objects for a given task. As can be observed, Pipino et al. provided a general definition whereas Fürber et al. provided a set of sub-categories for completeness. On the other hand, the two types of completeness defined in Mendes et al. can be mapped to the two categories (i) Schema completeness and (iii) Population completeness provided by Fürber et al.

⁶⁵ detection of an instance mapped back to more than one real world object leading to more than one interpretation.

Definition 18 (Completeness). *Completeness refers to the degree to which all required information is present in a particular dataset. In terms of LD, completeness comprises of the following aspects: (i) Schema completeness, the degree to which the classes and properties of an ontology are represented, thus can be called “ontology completeness”, (ii) Property completeness, measure of the missing values for a specific property, (iii) Population completeness is the percentage of all real-world objects of a particular type that are represented in the datasets and (iv) Interlinking completeness, which has to be considered especially in LD, refers to the degree to which instances in the dataset are interlinked.*

Metrics. The metrics identified for completeness are:

- CM1: schema completeness - no. of classes and properties represented / total no. of classes and properties [25,64,134]
- CM2: property completeness - no. of values represented for a specific property / total no. of values for a specific property [25,64]
- CM3: population completeness - no. of real-world objects are represented / total no. of real-world objects [25,64,134]
- CM4: interlinking completeness - no. of instances in the dataset that are interlinked / total no. of instances in a dataset [73]

It should be noted, that in this case, users should assume a closed-world-assumption where a gold standard dataset is available and can be used to compare against the converted dataset.

Example. In our use case, the flight search engine contains complete information to include all the airports and airport codes such that it allows a user to find an optimal route from the start to the end destination (even in cases when there is no direct flight). For example, the user wants to travel from Santa Barbara to San Francisco. Since our flight search engine contains interlinks between these close airports, the user is able to locate a direct flight easily.

Intra-relations The dimensions in this group are related to each other as follows: Data can be semantically accurate by representing the real world state but still can be inconsistent. However, if we merge accurate datasets, we will most likely get fewer inconsistencies than merging inaccurate datasets. On the other hand, being syntactically valid does not necessarily mean that the value is semantically accurate. Moreover, if a dataset is complete, syntactic validity, semantic accuracy and consistency checks need to be performed to determine if the values have been completed correctly. Additionally, conciseness is related to completeness since both point towards the dataset having all, however unique (non-redundant) information. However, if data integration leads to duplication of instances, it may lead to contradictory values thus leading to inconsistency [28].

7.5 Contextual dimensions

Contextual dimensions are those that highly depend on the context of the task at hand. There are four dimensions that are part of this group, namely *relevancy*, *trustworthiness*, *understandability* and *timeliness*.

Relevancy. Bizer [25] adopted the definition of relevancy from Pipino et al. [164] as “the extent to which information is applicable and helpful for the task at hand”. Additionally, Bizer [25] adopted the definition for the amount-of-data dimension from Pipino et al. [164] as “the extent to which the volume of data is appropriate for the task at hand”. Thus, since the amount-of-data dimension is similar to the relevancy dimension, we merge both dimensions. Flemming [59] defined amount-of-data as the “criterion influencing the usability of a data source”. While Pipino et al. provided a formal definition, Flemming and Chen et al. explained the dimension by mentioning its advantages.

Definition 19 (Relevancy). *Relevancy refers to the provision of information which is in accordance with the task at hand and important to the users’ query.*

Metrics. The metrics identified for relevancy are:

- R1: obtaining relevant data by:
 - counting the occurrence of relevant terms (keywords) within meta-data attributes (e.g. title, description, subject) [25]
 - using a combination of hyperlink analysis and information retrieval methods such as the vector space model that assigns higher weight to terms (keywords) that appear within the meta-data attributes [25]
 - ranking (a numerical value similar to PageRank), which determines the centrality of RDF documents and statements [31])
- R2: measuring the coverage (i.e. number of entities described in a dataset) and level of detail (i.e. number of properties) in a dataset to ensure that there exists an appropriate volume of relevant data for a particular task [59]

Example. When a user is looking for flights between any two cities, only relevant information i.e. departure and arrival airports and starting and ending time, duration and cost per person should be provided. Some datasets, in addition to relevant information, also contain much irrelevant data such as car rental, hotel booking, travel insurance etc. and as a consequence a lot of irrelevant extra information is provided. Providing irrelevant data distracts service developers and potentially users and wastes network resources. Instead, restricting the dataset to only flight related information, simplifies application development and increases the likelihood to return only relevant results to users.

Trustworthiness. Trustworthiness is a crucial topic due to the availability and the high volume of data from varying sources on the Web of Data. Bizer [25] adopted the definition of trust from Pipino et al. [164] as “the extent to which information is regarded as true and credible”. Jacobi et al. [94], similar to Pipino et al., referred to trustworthiness as a subjective measure of a user’s belief that the data is “true”. Gil et al. [67] used reputation of an entity or a dataset either as a result from direct experience or recommendations from others to establish trust. Additionally, Bizer [25] adopted the definition of objectivity from Pipino et al. [164] as “the extent to which information is unbiased, unprejudiced and impartial.” Thus, reputation as well as objectivity are part of the trustworthiness dimension. Other articles [31,66,68,70,71,76,134,184] provide metrics for assessing trustworthiness.

Definition 20 (Trustworthiness). *Trustworthiness is defined as the degree to which the information is accepted to be correct, true, real and credible.*

Metrics. The metrics identified for trustworthiness are:

- T1: computing statement trust values based on:
 - provenance information which can be either unknown or a value in the interval $[-1,1]$ where 1: absolute belief, -1: absolute disbelief and 0: lack of belief/disbelief [76]
 - opinion-based method, which use trust annotations made by several individuals [68,76]
 - provenance information and trust annotations in Semantic Web-based social-networks [70]
- T2: using annotations for data to encode two facets of information:
 - blacklists (indicates that the referent data is known to be harmful) [31] and
 - authority (a boolean value which uses the Linked Data principles to conservatively determine whether or not information can be trusted) [31]
- T3: using trust ontologies that assigns trust values that can be transferred from known to unknown data [94] using:
 - content-based methods (from content or rules) and
 - metadata-based methods (based on reputation assignments, user ratings, and provenance, rather than the content itself)
- T4: computing trust values between two entities through a path by using:
 - a propagation algorithm based on statistical techniques [184]
 - in case there are several paths, trust values from all paths are aggregated based on a weighting mechanism [184]
- T5: computing trustworthiness of the information provider by:
 - construction of decision networks informed by provenance graphs [66]
 - checking whether the provider/contributor is contained in a list of trusted providers [25]
 - indicating the level of trust for the publisher on a scale of 1 – 9 [67,71]
 - no bias or opinion expressed when a data provider interprets or analyses facts [25]
- T6: checking content trust⁶⁶ based on associations (e.g. anything having a relationship to a resource such as author of the dataset) that transfer trust from content to resources [67]
- T7: assignment of explicit trust ratings to the dataset by humans or analyzing external links or page ranks [134]

Example. In our flight search engine use case, if the flight information is provided by trusted and well-known airlines then a user is more likely to trust the information than when it is provided by an unknown travel agency. Generally information about a product or service (e.g. a flight) can be more trusted, when it is directly published by the producer or service provider (e.g. the airline). On the other hand, if a user retrieves information from a previously unknown source, she can decide whether to believe this information by checking whether the source is well-known or if it is contained in a list of trusted providers.

⁶⁶ Context trust is a trust judgment on a particular piece of information in a given context [67].

Understandability. Bizer [25] adopted the definition of understandability from Pipino et al. [164] stating that “understandability is the extent to which data is easily comprehended by the information consumer”. Flemming [59] related understandability also to the comprehensibility of data i.e. the ease with which human consumers can understand and utilize the data. Thus, comprehensibility can be interchangeably used with understandability.

Definition 21 (Understandability). *Understandability refers to the ease with which data can be comprehended without ambiguity and be used by a human information consumer.*

Metrics. The metrics identified for understandability are:

- U1: detection of human-readable labelling of classes, properties and entities as well as indication of metadata (e.g. name, description, website) of a dataset [59,86]
- U2: detect whether the pattern of the URIs is provided [59]
- U3: detect whether a regular expression that matches the URIs is present [59]
- U4: detect whether examples of SPARQL queries are provided [59]
- U5: checking whether a list of vocabularies used in the dataset is provided [59]
- U6: checking the effectiveness and the efficiency of the usage of the mailing list and/or the message boards [59]

Example. Let us assume that a user wants to search for flights between Boston and San Francisco using our flight search engine. Data related to Boston in the integrated dataset, for the required flight is represented as follows:

- `http://rdf.freebase.com/ns/m.049jnng`
- `http://rdf.freebase.com/ns/m.043j22x`
- “Boston Logan Airport”@en

For the first two items no human-readable label is available, therefore the machine is only able to display the URI as a result of the users query. This does not represent anything meaningful to the user besides perhaps that the information is from Freebase. The third entity, however, contains a human-readable label, which the user can easily understand.

Timeliness. Gamble et al. [66] defined timeliness as “a comparison of the date the annotation was updated with the consumer’s requirement”. The timeliness dimension is motivated by the fact that it is possible to have current data that is actually useless because it reflects a too old state of the real world for a specific usage. According to the timeliness dimension, data should ideally be recorded and reported as frequently as the source values change and thus never become outdated.

Definition 22. *Timeliness measures how up-to-date data is relative to a specific task.*

Metrics. The metrics identified for timeliness are:

- TI1: detecting freshness of datasets based on currency and volatility using the formula:
 $\max\{0, 1 - \text{currency}/\text{volatility}\}$ [77], which gives a value in a continuous scale from 0 to 1, where data with 1 is timely and 0 is unacceptable. In the formula, volatility is the length of time the data remains valid [64] and currency is the age of the data when delivered to the user [134,175].
- TI2: detecting freshness of datasets based on their data source by measuring the distance between last modified time of the data source and last modified time of the dataset [64,135]

Measuring currency of arbitrary documents or statements in LD presents several challenges: (i) it is unlikely that temporal metadata (e.g., the last modification date) are associated with statements; (ii) temporal metadata are not always available and (iii) LD is characterized by *autonomous providers* which use different vocabularies and different patterns for representing temporal metadata. In a recent study [174], the authors described the approaches used for representing temporal metadata associated with statements or documents and also showed the scarce availability of temporal metadata which impact the assessment of currency.

Example. Consider a user checking the flight timetable for her flight from a city A to a city B. Suppose that the result is a list of triples comprising of the description of the resource A such as the connecting airports, the time of arrival, the terminal, the gate, etc. This flight timetable is updated every 10 minutes (volatility). Assume there is a change of the flight departure time, specifically a delay of one hour. However, this information is communicated to the control room with a slight delay. They update this information in the system after 30 minutes. Thus, the timeliness constraint of updating the timetable within 10 minutes is not satisfied which renders the information out-of-date.

Intra-relations Data is of high relevance if data is current for the user needs. The timeliness of information thus influences its relevancy. On the other hand, if a dataset has current information, it is considered to be trustworthy. Moreover, to allow users to properly understand information in a dataset, a system should be able to provide sufficient relevant information.

7.6 Representational dimensions

Representational dimensions capture aspects related to the design of the data such as the *representational-conciseness*, *interoperability*, *interpretability* as well as the *versatility*.

Representational-conciseness. Bizer [25], adopted the definition of representational-conciseness from Pipino et al. [164] as “the extent to which information is compactly represented”. This is the only article that describes this dimension (from the core set of articles included in this survey).

Definition 23 (Representational-conciseness). *Representational-conciseness refers to the representation of the data which is compact and well formatted on the one hand and clear and complete on the other hand.*

Metrics. The metrics identified for representational-conciseness are:

- RC1: detection of long URIs or those that contain query parameters [86]
- RC2: detection of RDF primitives i.e. RDF reification, RDF containers and RDF collections [86]

Example. Our flight search engine represents the URIs for the destination compactly with the use of the airport codes. For example, LEJ is the airport code for Leipzig, therefore the URI is `http://airlines.org/LEJ`. This short representation of URIs helps users share and memorize them easily.

Interoperability. Bizer [25] adopted the definition of representational-consistency from Pipino et al. [164] as “the extent to which information is represented in the same format”. We use the term “interoperability” for this dimension. In addition, the definition of “uniformity”, which refers to the re-use of established formats to represent data as described by Flemming [59], can be associated to the interoperability of the dataset. Additionally, as stated in Hogan et al. [86], the re-use of well-known terms to describe resources in a uniform manner increases the interoperability of data published in this manner and contributes towards the interoperability of the entire dataset.

Definition 24 (Interoperability). *Interoperability is the degree to which the format and structure of the information conforms to previously returned information as well as data from other sources.*

Metrics. The metrics identified for interoperability are:

- IO1: detection of whether existing terms from all relevant vocabularies for that particular domain have been reused [86]
- IO2: usage of relevant vocabularies for that particular domain [59]

Example. Let us consider different airline datasets using different notations for representing temporal data, e.g. one dataset uses the time ontology while another dataset uses XSD notations. This makes querying the integrated dataset difficult as it requires users to understand the various heterogeneous schema. Additionally, with the difference in the vocabularies used to represent the same concept (in this case time), the consumers are faced with problem of how the data can be interpreted and displayed. In order to avoid these interoperability issues, we provide data based on the Linked Data principles, which are designed to support heterogeneous description models necessary to handle different formats of data.

Interpretability. Bizer [25] adopted the definition of interpretability from Pipino et al. [164] as the “extent to which information is in appropriate languages, symbols and units and the definitions are clear”. This is the only article that describes this dimension (from the core set of articles included in this survey).

Definition 25 (Interpretability). *Interpretability refers to technical aspects of the data, that is, whether information is represented using an appropriate notation and whether the machine is able to process the data.*

Metrics. The metrics identified for interpretability are:

- IN1: identifying objects and terms used to define these objects with globally unique identifiers⁶⁷ [25]
- IN2: detecting the use of appropriate language, symbols, units, datatypes and clear definitions [25,164,59]
- IN3: detection of invalid usage of undefined classes and properties (i.e. those without any formal definition) [85]
- IN4: detecting the use of blank nodes⁶⁸ [86]

Example. Consider our flight search engine and a user that is looking for a flight from Mumbai to Boston with a two day stop-over in Berlin. The user specifies the dates correctly. However, since the flights are operated by different airlines, thus different datasets, they have a different way of representing the date. In the first leg of the trip, the date is represented in the format dd/mm/yyyy whereas in the other case, date is represented as mm/dd/yy. Thus, the machine is unable to correctly interpret the data and cannot provide an optimal result for this query. This lack of consensus in the format of the date hinders the ability of the machine to interpret the data and thus provide the appropriate flights.

Versatility. Flemming [59] defined versatility as the “alternative representations of the data and its handling.” This is the only article that describes this dimension (from the core set of articles included in this survey).

Definition 26 (Versatility). *Versatility refers to the availability of the data in an internationalized way and alternative representations of data.*

Metrics. The metrics identified for versatility are:

- V1: checking whether data is available in different serialization formats [59]
- V2: checking whether data is available in different languages [14,59,104]

Example. Consider a user who does not understand English but only Spanish and wants to use our flight search engine. In order to cater to the needs of such a user, the dataset should provide labels and other language-dependent information in Spanish so that any user has the capability to understand it.

Intra-relations The dimensions in this group are related as follows: Interpretability is related to the interoperability of data since the consistent representation (e.g. re-use of established vocabularies) ensures that a system will be able to interpret the data correctly [49]. Versatility is also related to the interpretability of a dataset as the more versatile forms a dataset is represented in (e.g. in different languages), the more interpretable a dataset is. Additionally, concise representation of the data allows the data to be interpreted correctly.

⁶⁷ www.w3.org/TR/webarch/

⁶⁸ Blank nodes are not recommended since they cannot be externally referenced.

7.7 Data quality assessment frameworks

In Table 9, we compare the tools proposed by eight of the 21 core articles, identified in the survey, based on eight different attributes. These tools implement the methodologies and metrics defined in the respective approaches.

Table 9: Comparison of quality assessment tools according to several attributes.

	Trellis, Gil et al., 2002	TrustBot, Golbeck et al., 2003	tSPARQL, Hartig, 2008	WIQA, Bizer et al., 2009	ProLOD, Böhm et al., 2010	Flemming, et al., 2010	LinkQA, Gueret et al., 2012	Sieve, Mendes et al., 2012
Accessibility/Availability	–	–	✓	–	–	✓	✓	✓
Licensing	Open-source	–	GPL v3	Apache v2	–	–	Open-source	Apache
Automation	Semi-automated	Semi-automated	Semi-automated	Semi-automated	Semi-automated	Semi-automated	Automated	Semi-automated
Collaboration	Allows users to add observations and conclusions	No	No	No	No	No	No	No
Customizability	✓	✓	✓	✓	✓	✓	No	✓
Scalability	–	No	Yes	–	–	No	Yes	Yes
Usability/Documentation	2	4	4	2	2	3	2	4
Maintenance (Last updated)	2005	2003	2012	2006	2010	2010	2011	2012

Accessibility/Availability. In Table 9, only the tools marked with a tick are available to be used for quality assessment. The other tools are either available only as a demo or screencast (Trellis, ProLOD) or not available at all (TrustBot, WIQA).

Licensing. Each of the tools is available using a particular software license, which specifies the restrictions with which it can be redistributed. The Trellis and LinkQA tools are open-source and as such by default they are protected by copyright which is *All Rights Reserved*. Also, WIQA and Sieve are both available with open-source license: the Apache Version 2.0⁶⁹ and Apache licenses respectively. tSPARQL is distributed under the GPL v3 license⁷⁰. However, no licensing information is available for TrustBot, ProLOD and Flemming’s tool.

Automation. The automation of a system is the ability to automatically perform its intended tasks thereby reducing the need for human intervention. In this context, we classify the eight tools into semi-automated and automated approaches. As seen in Table 9, all the tools are semi-automated except for LinkQA, which is completely automated as there is no user involvement. LinkQA automatically selects a set of resources, information from the Web of Data (i.e. SPARQL endpoints and/or dereferencable resources) and a set of new triples as input and generates the respective quality assessment reports.

On the other hand, the WIQA and Sieve tools require a high degree of user involvement. Specifically in Sieve, the definition of metrics has to be done by creating an XML

⁶⁹ <http://www.apache.org/licenses/LICENSE-2.0>

⁷⁰ <http://www.gnu.org/licenses/gpl-3.0.html>

file which contains specific configurations for a quality assessment task. Although it gives the users the flexibility of tweaking the tool to match their needs, it requires much time for understanding the required XML file structure and specification. The other semi-automated tools, Trellis, TrustBot, tSPARQL, ProLOD and Flemming's tool require a minimum amount of user involvement. For example, Flemming's Data Quality Assessment Tool requires the user to answer a few questions regarding the dataset (e.g. existence of a human-readable license) or they have to assign weights to each of the pre-defined data quality metrics.

Collaboration. Collaboration is the ability of a system to support co-operation between different users of the system. None of the tools, except Trellis, support collaboration between different users of the tool. The Trellis user interface allows several users to express their trust value for a data source. The tool allows several users to add and store their observations and conclusions. Decisions made by users on a particular source are stored as annotations, which can be used to analyze conflicting information or handle incomplete information.

Customizability. Customizability is the ability of a system to be configured according to the users' needs and preferences. In this case, we measure the customizability of a tool based on whether the tool can be used with any dataset that the user is interested in. Only LinkQA cannot be customized since the user cannot add any dataset of her choice. The other seven tools can be customized according to the use case. For example, in TrustBot, an IRC bot that makes trust recommendations to users (based on the trust network it builds), the users have the flexibility to submit their own URIs to the bot at any time while incorporating the data into a graph. Similarly, Trellis, tSPARQL, WIQA, ProLOD, Flemming's tool and Sieve can be used with any dataset.

Scalability. Scalability is the ability of a system, network, or process to handle a growing amount of work or its ability to be enlarged to accommodate that growth. Out of the eight tools only three, the tSPARQL, LinkQA and Sieve, tools are scalable, that is, they can be used with large datasets. Flemming's tool and TrustBot are reportedly not scalable for large datasets [71,59]. Flemming's tool, on the one hand, performs analysis based on a sample of three entities whereas TrustBot takes as input two email addresses to calculate the weighted average trust value. Trellis, WIQA and ProLOD do not provide any information on the scalability.

Usability/Documentation. Usability is the ease of use and learnability of a human-made object, in this case the quality assessment tool. We assess the usability of the tools based on the ease of use as well as the complete and precise documentation available for each of them. We score them based on a scale from 1 (low usability) – 5 (high usability). TrustBot, tSPARQL and Sieve score high in terms of usability and documentation followed by Flemming's data quality assessment tool. Trellis, WIQA, ProLOD and LinkQA rank lower in terms of ease of use since they do not contain useful documentation of how to use the tool.

Maintenance/Last updated. While TrustBot, Trellis and WIQA have not been updated since they were first introduced in 2003, 2005 and 2006 respectively, ProLOD and Flemming's tool have been updated in 2010. The recently updated tools are LinkQA (2011), tSPARQL (2012) and Sieve (2012) and are currently being maintained.

8 Outlook and Future Challenges

Although the different approaches for aspects of the Linked Data life-cycle as presented in this chapter are already working together, more effort must be done to further integrate them in ways that they mutually fertilize themselves. The discovery of new links or the authoring of new resource descriptions, for example, should automatically trigger the enrichment of the linked knowledge bases. The enrichment in turn can trigger the application of inconsistency detection and repair techniques. This leads to recognizing data quality problems and their consequent assessment and improvement. The browsing and exploration paths followed by end-users can be taken into account for machine learning techniques to refine the knowledge bases etc. Ultimately, when the different aspects of Linked Data management are fully integrated we envision the Web of Data becoming a washing machine for knowledge. A progress in one particular aspect will automatically trigger improvements in many other ones as well. In the following we outline some research challenges and promising research directions regarding some of the Linked Data management aspects.

Extraction. One promising research direction with regard to the extraction from unstructured sources is the development of standardized, LOD enabled integration interfaces between existing NLP tools. An open question is whether and how efficient bi-directional synchronization between extraction source and target knowledge base can be established. With regard to the extraction from structured sources (e.g. relational, XML) we need a declarative syntax and semantics for data model transformations. Some orthogonal challenges include the use of LOD as background knowledge and the representation and tracing of provenance information.

Authoring. Current Semantic Wikis still suffer from a lack of scalability. Hence, an important research and development target are large-scale Semantic Wikis, which include functionality for access control and provenance. In order to further flexibilize and simplify the authoring an adaptive choreography of editing widgets based on underlying data structures is needed. Also, the joint authoring of unstructured and structured sources (i.e. HTML/RDFa authoring) and better support for the integrated semantic annotation of other modalities such as images, audio, video is of paramount importance.

Natural Language Queries. One of the future challenges for Linked Data is to create user interfaces, which are able to hide the complexity of the underlying systems. A possible path towards this goal is question answering, e.g. converting natural language queries to SPARQL [193,114]. In order to allow users to interact with such systems, there is ongoing work on converting the created SPARQL queries back to natural language [149] and employ feedback mechanisms [113,83]. Ultimately, a goal is to provide users enhanced functionality without the need to adapt to different kinds of interface.

Automatic Management of Resources for Linking. With the growth of the Cloud and of the datasets that need to be interlinked, the use of parallel hardware has been studied over the last few years [84,151]. The comparative study of parallel hardware for link discovery yet shows surprising results and suggests that the use of massively parallel

yet local hardware can lead to tremendous runtime improvements. Still, when result sets go beyond sizes of 10^{10} , the higher amount of resources available on remote devices in the Cloud is still to be used. Devising automatic solutions for selecting the right hardware to run a linking task is one of the most interesting research areas pertaining to the efficient execution of link specifications. In addition, developing reduction-ratio optimal algorithms for spaces other than Minkowski spaces promises to ensure the best possible use of available hardware. Finally, devising more efficient means to combine single algorithms is the third open area of research in this domain. The challenges faces with regard to learning link specifications are also manifold and include devising approaches that can efficiently detected most informative positive and negative examples as well even running in a fully unsupervised manner on properties that are not one-to-one relations.

Linked Data Visualization The potential of the vast amount of Linked Data on the Web is enormous but in most cases it is very difficult *and* cumbersome for users to visualize, explore and use this data, especially for lay users without experience with Semantic Web technologies. Visualizations are useful for obtaining an overview of the datasets, their main types, properties and the relationships between them. Compared to prior information visualization strategies, we have a unique opportunity on the Data Web. The unified RDF data model being prevalent on the Data Web enables us to bind data to visualizations in an *unforeseen* and *dynamic* way. An information visualization technique requires certain data structures to be present. When we can derive and generate these data structures automatically from reused vocabularies or semantic representations, we are able to realize a largely automatic visualization workflow. Ultimately, various visualizations techniques can develop an ecosystem of data extractions and visualizations, which can be bound together in a dynamic and unforeseen way. This will enable users to explore datasets even if the publisher of the data does not provide any exploration or visualization means. Yet, most existing work related to visualizing RDF is focused on concrete domains and concrete datatypes.

Linked Data Quality With the amount of Linked Data on the Web growing at an exponential rate, assessing the quality of the datasets is of utmost importance in order to ensure reliability of the applications built using that data. However, currently there are few methodologies in place to perform this assessment. Also, the tools currently available are either not scalable for large datasets or do not cover all the data quality aspects for assessing the quality of Linked Data, in particular. Moreover, these tools provide results that are difficult to interpret, in certain cases do not allow a user to choose the input dataset or require considerable amount of user involvement. Currently, crowd-sourcing based approaches for quality assessment are being explored, which prove to be cost-effective and accurate when used in combination with semi-automated approaches [210,2]. Thus, one of the main challenges is to devise such a methodology supported by a tool to perform quality assessment for large datasets. Additionally, the tools should provide features that allow dataset owners to amend the quality aspects that are detected during the quality assessment.

Acknowledgments

We would like to thank our colleagues from the AKSW research group in Leipzig as well as the LOD2, GeoKnow and BioASQ project consortia, without whom writing this chapter would not have been possible. In particular, we would like to thank Christian Bizer and Tom Heath, whose Chapter 2 of the book ‘Linked Data – Evolving the Web into a Global Data Space’ [79] served as a blueprint for section 2; Sebastian Hellmann, Claus Stadler, René Speck, Jörg Unbehauen and Ricardo Usbeck for their contributions to section 3, Sebastian Tramp, Michael Martin, Norman Heino, Phillip Frischmuth and Thomas Riechert for their contributions to the development of OntoWiki as described in section 4. This work was supported by a grant from the European Union’s 7th Framework Programme provided for the projects LOD2 (GA no. 257943), GeoKnow (GA no. 318159) and BioASQ (GA no. 318652).

References

1. Resource description framework (RDF): Concepts and abstract syntax. Technical report, W3C, 2 2004.
2. Maribel Acosta, Amrapali Zaveri, Elena Simperl, Dimitris Kontokostas, Sören Auer, and Jens Lehmann. Crowdsourcing linked data quality assessment. In *12th International Semantic Web Conference, 21-25 October 2013, Sydney, Australia*, pages 260–276, 2013.
3. Ben Adida, Mark Birbeck, Shane McCarron, and Steven Pemberton. RDFa in XHTML: Syntax and processing – a collection of attributes and processing rules for extending XHTML to support RDF. W3C Recommendation, Oct 2008. <http://www.w3.org/TR/rdfa-syntax/>.
4. Eugene Agichtein and Luis Gravano. Snowball: Extracting relations from large plain-text collections. In *In ACM DL*, pages 85–94, 2000.
5. Alan Agresti. *An Introduction to Categorical Data Analysis 2nd edition*. Wiley-Interscience, 1997.
6. R. Amsler. Research towards the development of a lexical knowledge base for natural language processing. *SIGIR Forum*, 23:1–2, 1989.
7. Sören Auer, Chris Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary Ives. Dbpedia: A nucleus for a web of open data. In *ISWC 2007*, volume 4825 of *LNCS*, pages 722–735. Springer.
8. Sören Auer, Sebastian Dietzold, Jens Lehmann, Sebastian Hellmann, and David Aumüller. Triplify: Light-weight linked data publication from relational databases. In Juan Quemada, Gonzalo León, Yoëlle S. Maarek, and Wolfgang Nejdl, editors, *Proceedings of the 18th International Conference on World Wide Web, WWW 2009, Madrid, Spain, April 20-24, 2009*, pages 621–630. ACM, 2009.
9. Sören Auer, Sebastian Dietzold, and Thomas Riechert. OntoWiki - A Tool for Social, Semantic Collaboration. In Isabel F. Cruz, Stefan Decker, Dean Allemang, Chris Preist, Daniel Schwabe, Peter Mika, Michael Uschold, and Lora Aroyo, editors, *The Semantic Web - ISWC 2006, 5th International Semantic Web Conference, ISWC 2006, Athens, GA, USA, November 5-9, 2006, Proceedings*, volume 4273 of *Lecture Notes in Computer Science*, pages 736–749, Berlin / Heidelberg, 2006. Springer.
10. Sören Auer and Heinrich Herre. A versioning and evolution framework for rdf knowledge bases. In *Proceedings of at Sixth International Andrei Ershov Memorial Conference - Perspectives of System Informatics (PSI’06), 27-30 June, Novosibirsk, Akademgorodok, Russia*, volume 4378, June 2006.

11. Sören Auer and Jens Lehmann. Making the web a data washing machine - creating knowledge out of interlinked data. *Semantic Web Journal*, 2010.
12. Sören Auer, Jens Lehmann, and Sebastian Hellmann. LinkedGeoData - adding a spatial dimension to the web of data. In *Proc. of 8th International Semantic Web Conference (ISWC)*, 2009.
13. Sören Auer, Jens Lehmann, Axel-Cyrille Ngonga Ngomo, and Amrapali Zaveri. Introduction to linked data and its lifecycle on the web. In Sebastian Rudolph, Georg Gottlob, Ian Horrocks, and Frank van Harmelen, editors, *Reasoning Web*, volume 8067 of *Lecture Notes in Computer Science*, pages 1–90. Springer, 2013.
14. Sören Auer, Matthias Weidl, Jens Lehmann, Amrapali Zaveri, and Key-Sun Choi. I18n of semantic web applications. In *ISWC 2010*, volume 6497 of *LNCIS*, pages 1 – 16, 2010.
15. David Aumüller. Semantic Authoring and Retrieval within a Wiki (WikSAR). In Demo Session at the Second European Semantic Web Conference (ESWC2005), May 2005. Available at <http://wiksar.sf.net>, 2005.
16. Franz Baader, Calvanese Diageo, Deborah McGuinness, Daniele Nardi, and Peter Patel-Schneider, editors. *The Description Logic Handbook*. Cambridge, 2003.
17. Franz Baader, Bernhard Ganter, Ulrike Sattler, and Baris Sertkaya. Completing description logic knowledge bases using formal concept analysis. In *IJCAI 2007*. AAAI Press, 2007.
18. Franz Baader, Baris Sertkaya, and Anni-Yasmin Turhan. Computing the least common subsumer w.r.t. a background terminology. *J. Applied Logic*, 5(3):392–420, 2007.
19. Liviu Badea and Shan-Hwei Nienhuys-Cheng. A refinement operator for description logics. In *ILP 2000*, volume 1866 of *LNAI*, pages 40–59. Springer, 2000.
20. Rohan Baxter, Peter Christen, and Tim Churches. A comparison of fast blocking methods for record linkage. In *KDD-2003 Workshop on Data Cleaning, Record Linkage, and Object Consolidation*, 2003.
21. David Ben-David, Tamar Domany, and Abigail Tarem. Enterprise data classification using semantic web technologies. In *ISWC2010*, 2010.
22. Tim Berners-Lee, Roy Thomas Fielding, and Larry Masinter. Uniform resource identifiers (URI): Generic syntax. Internet RFC 2396, August 1998.
23. Ravish Bhagdev, Sam Chapman, Fabio Ciravegna, Vitaveska Lanfranchi, and Daniela Petrelli. Hybrid search: Effectively combining keywords and semantic searches. In *ESWC 2008*, pages 554–568, 2008.
24. Mikhail Bilenko, Beena Kamath, and Raymond J. Mooney. Adaptive blocking: Learning to scale up record linkage. In *ICDM '06*, pages 87–96. IEEE, 2006.
25. Christian Bizer. *Quality-Driven Information Filtering in the Context of Web-Based Information Systems*. PhD thesis, Freie Universität Berlin, March 2007.
26. Christian Bizer and Richard Cyganiak. Quality-driven information filtering using the wiqua policy framework. *Web Semantics*, 7(1):1 – 10, Jan 2009.
27. Jens Bleiholder and Felix Naumann. Data fusion. *ACM Comput. Surv.*, 41(1):1–41, 2008.
28. Jens Bleiholder and Felix Naumann. Data fusion. *ACM Computing Surveys (CSUR)*, 41(1):1, 2008.
29. Anselm Blumer, Andrzej Ehrenfeucht, David Haussler, and Manfred K. Warmuth. Occam’s razor. In *Readings in Machine Learning*, pages 201–204. Morgan Kaufmann, 1990.
30. Christoph Böhm, Felix Naumann, Ziawasch Abedjan, Dandy Fenz, Toni Grütze, Daniel Hefenbrock, Matthias Pohl, and David Sonnabend. Profiling linked open data with PROLOD. In *ICDE Workshops*, pages 175–178. IEEE, 2010.
31. Piero A. Bonatti, Aidan Hogan, Axel Polleres, and Luigi Sauro. Robust and scalable linked data reasoning incorporating provenance and trust annotations. *Journal of Web Semantics*, 9(2):165 – 201, 2011.
32. Leo Breiman. Bagging predictors. *Machine Learning*, 24(2):123–140, 1996.

33. Leo Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001.
34. Dan Brickley and Ramanatgan V. Guha. RDF Vocabulary Description Language 1.0: RDF Schema. W3C recommendation, W3C, February 2004. <http://www.w3.org/TR/2004/REC-rdf-schema-20040210/>.
35. Sergey Brin. Extracting patterns and relations from the world wide web. In *WebDB*, pages 172–183, London, UK, 1999. Springer-Verlag.
36. Lorenz Bühmann and Jens Lehmann. Universal OWL axiom enrichment for large knowledge bases. In *Proceedings of EKAW 2012*, pages 57–71. Springer, 2012.
37. Jeremy J Carroll. Signing RDF graphs. In *ISWC*, pages 369–384. Springer, 2003.
38. Chih-Chung Chang and Chih-Jen Lin. Libsvm - a library for support vector machines, 2001. The Weka classifier works with version 2.82 of LIBSVM.
39. Ping Chen and Walter Garcia. Hypothesis generation and data quality assessment through association mining. In *IEEE ICCI*, pages 659–666. IEEE, 2010.
40. Didier Cherix, Sebastian Hellmann, and Jens Lehmann. Improving the performance of a sparql component for semantic web applications. In *JIST*, 2012.
41. Namyoun Choi, Il-Yeol Song, and Hyoil Han. A survey on ontology mapping. *SIGMOD Record*, 35(3):34–41, 2006.
42. Sam Coates-Stephens. The analysis and acquisition of proper names for the understanding of free text. *Computers and the Humanities*, 26:441–456, 1992. 10.1007/BF00136985.
43. William W. Cohen, Alexander Borgida, and Haym Hirsh. Computing least common subsumers in description logics. In *AAAI 1992*, pages 754–760, 1992.
44. William W. Cohen and Haym Hirsh. Learning the CLASSIC description logic: Theoretical and experimental results. In *KR 1994*, pages 121–133. Morgan Kaufmann, 1994.
45. Marco Cornolti, Paolo Ferragina, and Massimiliano Ciaramita. A framework for benchmarking entity-annotation systems. In *Proceedings of the 22nd international conference on World Wide Web*, pages 249–260. International World Wide Web Conferences Steering Committee, 2013.
46. Silviu Cucerzan. Large-scale named entity disambiguation based on wikipedia data. In *EMNLP-CoNLL*, pages 708–716, 2007.
47. James R. Curran and Stephen Clark. Language independent ner using a maximum entropy tagger. In *Proceedings of the seventh conference on Natural language learning at HLT-NAACL 2003 - Volume 4*, pages 164–167, Morristown, NJ, USA, 2003. Association for Computational Linguistics.
48. Claudia d’Amato, Nicola Fanizzi, and Floriana Esposito. A note on the evaluation of inductive concept classification procedures. In Aldo Gangemi, Johannes Keizer, Valentina Presutti, and Heiko Stoermer, editors, *SWAP 2008*, volume 426 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2008.
49. Li Ding and Tim Finin. Characterizing the semantic web on the web. In *5th ISWC*, 2006.
50. Ahmed K. Elmagarmid, Panagiotis G. Ipeirotis, and Vassilios S. Verykios. Duplicate record detection: A survey. *IEEE Transactions on Knowledge and Data Engineering*, 19:1–16, 2007.
51. Timofey Ermilov, Norman Heino, Sebastian Tramp, and Sören Auer. OntoWiki Mobile — Knowledge Management in your Pocket. In *Proceedings of the ESWC2011*, 2011.
52. Floriana Esposito, Nicola Fanizzi, Luigi Iannone, Ignazio Palmisano, and Giovanni Semeraro. Knowledge-intensive induction of terminologies from metadata. In *ISWC 2004*, pages 441–455. Springer, 2004.
53. Oren Etzioni, Michael Cafarella, Doug Downey, Ana-Maria Popescu, Tal Shaked, Stephen Soderland, Daniel S. Weld, and Alexander Yates. Unsupervised named-entity extraction from the web: an experimental study. *Artif. Intell.*, 165:91–134, June 2005.
54. Jérôme Euzenat and Pavel Shvaiko. *Ontology matching*. Springer-Verlag, Heidelberg (DE), 2007.

55. Nicola Fanizzi, Claudia d'Amato, and Floriana Esposito. DL-FOIL concept learning in description logics. In *ILP 2008*, volume 5194 of *LNCS*, pages 107–121. Springer, 2008.
56. R. Fielding, J. Gettys, J. Mogul, H. Frystyk, L. Masinter, P. Leach, and T. Berners-Lee. Hypertext transfer protocol – http/1.1 (rfc 2616). Request For Comments, 1999. available at <http://www.ietf.org/rfc/rfc2616.txt>, accessed 7 July 2006.
57. Jenny Rose Finkel, Trond Grenager, and Christopher Manning. Incorporating non-local information into information extraction systems by gibbs sampling. In *Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics*, ACL '05, pages 363–370, Morristown, NJ, USA, 2005. Association for Computational Linguistics.
58. Daniel Fleischhacker, Johanna Völker, and Heiner Stuckenschmidt. Mining rdf data for property axioms. In *OTM Conferences (2)*, pages 718–735, 2012.
59. Annika Flemming. Quality characteristics of linked data publishing datasources. Master's thesis, Humboldt-Universität zu Berlin, 2010.
60. Eibe Frank, Gordon W. Paynter, Ian H. Witten, Carl Gutwin, and Craig G. Nevill-Manning. Domain-specific keyphrase extraction. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*, IJCAI '99, pages 668–673, San Francisco, CA, USA, 1999. Morgan Kaufmann Publishers Inc.
61. Yoav Freund and Robert E. Schapire. Experiments with a New Boosting Algorithm. In *International Conference on Machine Learning*, pages 148–156, 1996.
62. Brendan J. Frey and Delbert Dueck. Clustering by passing messages between data points. *Science*, 315:972–976, 2007.
63. J. Friedman, T. Hastie, and R. Tibshirani. Additive logistic regression: a statistical view of boosting. Technical report, Stanford University, 1998.
64. Christian Fürber and Martin Hepp. SWIQA - a semantic web information quality assessment framework. In *ECIS*, 2011.
65. Joao Gama. Functional trees. 55(3):219–250, 2004.
66. Matthew Gamble and Carole Goble. Quality, trust, and utility of scientific data on the web: Towards a joint model. In *ACM WebSci*, pages 1–8, June 2011.
67. Yolanda Gil and Donovan Artz. Towards content trust of web resources. *Web Semantics*, 5(4):227 – 239, December 2007.
68. Yolanda Gil and Varun Ratnakar. Trusting information sources one citizen at a time. In *ISWC*, pages 162 – 176, 2002.
69. Hugh Glaser, Ian C. Millard, Won-Kyung Sung, Seungwoo Lee, Pyung Kim, and Beom-Jong You. Research on linked data and co-reference resolution. Technical report, University of Southampton, 2009.
70. Jennifer Golbeck. Using trust and provenance for content filtering on the semantic web. In *Workshop on Models of Trust on the Web at the 15th World Wide Web Conference*, 2006.
71. Jennifer Golbeck, Bijan Parsia, and James Hendler. Trust networks on the semantic web. In *CIA*, 2003.
72. R. Grishman and R. Yangarber. Nyu: Description of the Proteus/Pet system as used for MUC-7 ST. In *MUC-7*. Morgan Kaufmann, 1998.
73. Christophe Guéret, Paul Groth, Claus Stadler, and Jens Lehmann. Assessing linked data mappings using network measures. In *ESWC*, 2012.
74. Mark Hall, Eibe Frank, Geoffrey Holmes, Bernhard Pfahringer, Peter Reutemann, and Ian H. Witten. The weka data mining software: An update. *SIGKDD Explor. Newsl.*, 11(1):10–18, November 2009.
75. Sanda Harabagiu, Cosmin Adrian Bejan, and Paul Mărarescu. Shallow semantics for relation extraction. In *IJCAI*, pages 1061–1066, 2005.
76. Olaf Hartig. Trustworthiness of data on the web. In *STI Berlin and CSW PhD Workshop, Berlin, Germany*, 2008.

77. Olaf Hartig and Jun Zhao. Using web data provenance for quality assessment. In Juliana Freire, Paolo Missier, and Satya Sanket Sahoo, editors, *SWPM*, volume 526 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2009.
78. Trevor Hastie and Robert Tibshirani. Classification by pairwise coupling. In Michael I. Jordan, Michael J. Kearns, and Sara A. Solla, editors, *Advances in Neural Information Processing Systems*, volume 10. MIT Press, 1998.
79. Tom Heath and Christian Bizer. *Linked Data - Evolving the Web into a Global Data Space*, volume 1 of *Synthesis Lectures on the Semantic Web: Theory and Technology*. Morgan & Claypool, 2011.
80. Norman Heino, Sebastian Dietzold, Michael Martin, and Sören Auer. Developing semantic web applications with the ontowiki framework. In Tassilo Pellegrini, Sören Auer, Klaus Tochtermann, and Sebastian Schaffert, editors, *Networked Knowledge - Networked Media*, volume 221 of *Studies in Computational Intelligence*, pages 61–77. Springer, Berlin / Heidelberg, 2009.
81. Sebastian Hellmann, Jens Lehmann, and Sören Auer. Learning of OWL class descriptions on very large knowledge bases. *Int. J. Semantic Web Inf. Syst.*, 5(2):25–48, 2009.
82. Sebastian Hellmann, Jens Lehmann, and Sören Auer. Learning of owl class expressions on very large knowledge bases and its applications. In Interoperability Semantic Services and Web Applications: Emerging Concepts, editors, *Learning of OWL Class Expressions on Very Large Knowledge Bases and its Applications*, chapter 5, pages 104–130. IGI Global, 2011.
83. Sebastian Hellmann, Jens Lehmann, Jörg Unbehauen, Claus Stadler, Thanh Nghia Lam, and Markus Strohmaier. Navigation-induced knowledge engineering by example. In *JIST*, 2012.
84. Stanley Hillner and Axel-Cyrille Ngonga Ngomo. Parallelizing limes for large-scale link discovery. In *I'Semantics*, 2011.
85. Aidan Hogan, Andreas Harth, Alexandre Passant, Stefan Decker, and Axel Polleres. Weaving the pedantic web. In *LDOW*, 2010.
86. Aidan Hogan, Jürgen Umbrich, Andreas Harth, Richard Cyganiak, Axel Polleres, and Stefan Decker. An empirical survey of linked data conformance. *Journal of Web Semantics*, 2012.
87. Matthew Horridge and Peter F. Patel-Schneider. Manchester syntax for OWL 1.1. In *OWLED 2008*, 2008.
88. HTML 5: A vocabulary and associated APIs for HTML and XHTML. W3C Working Draft, August 2009. <http://www.w3.org/TR/2009/WD-html5-20090825/>.
89. Luigi Iannone and Ignazio Palmisano. An algorithm based on counterfactuals for concept learning in the semantic web. In *IEA/AIE 2005*, pages 370–379, June 2005.
90. Luigi Iannone, Ignazio Palmisano, and Nicola Fanizzi. An algorithm based on counterfactuals for concept learning in the semantic web. *Applied Intelligence*, 26(2):139–159, 2007.
91. Ali Inan, Murat Kantarcioglu, Elisa Bertino, and Monica Scannapieco. A hybrid approach to private record linkage. *ICDE*, 0:496–505, 2008.
92. Robert Isele, Anja Jentzsch, and Christian Bizer. Efficient multidimensional blocking for link discovery without losing recall. In *WebDB*, 2011.
93. Robert Isele, Anja Jentzsch, and Christian Bizer. Active learning of expressive linkage rules for the web of data. In *ICWE*, pages 411–418, 2012.
94. Ian Jacobi, Lalana Kagal, and Ankesh Khandelwal. Rule-based trust assessment on the semantic web. In *RuleML*, pages 227 – 241, 2011.
95. Ian Jacobs and Norman Walsh. Architecture of the world wide web, volume one. World Wide Web Consortium, Recommendation REC-webarch-20041215, December 2004.

96. George H. John and Pat Langley. Estimating continuous distributions in bayesian classifiers. In *Eleventh Conference on Uncertainty in Artificial Intelligence*, pages 338–345, San Mateo, 1995. Morgan Kaufmann.
97. Joseph Juran. *The Quality Control Handbook*. McGraw-Hill, New York, 1974.
98. Su Nam Kim and Min-Yen Kan. Re-examining automatic keyphrase extraction approaches in scientific articles. In *Proceedings of the Workshop on Multiword Expressions: Identification, Interpretation, Disambiguation and Applications*, MWE '09, pages 9–16, Stroudsburg, PA, USA, 2009. Association for Computational Linguistics.
99. Su Nam Kim, Olena Medelyan, Min-Yen Kan, and Timothy Baldwin. Semeval-2010 task 5: Automatic keyphrase extraction from scientific articles. In *Proceedings of the 5th International Workshop on Semantic Evaluation*, SemEval '10, pages 21–26, Stroudsburg, PA, USA, 2010. Association for Computational Linguistics.
100. J. Kittler, M. Hatef, R. P W Duin, and J. Matas. On combining classifiers. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 20(3):226–239, Mar 1998.
101. Ron Kohavi. The power of decision tables. In *8th European Conference on Machine Learning*, pages 174–189. Springer, 1995.
102. Hanna Köpcke, Andreas Thor, and Erhard Rahm. Comparative evaluation of entity resolution approaches with fever. *Proc. VLDB Endow.*, 2(2):1574–1577, 2009.
103. Markus Krötzsch, Denny Vrandečić, Max Völkel, Heiko Haller, and Rudi Studer. Semantic wikipedia. *Journal of Web Semantics*, 5:251–261, September 2007.
104. Jose Emilio Labra Gayo, Dimitris Kontokostas, and Sören Auer. Multilingual linked open data patterns. *Semantic Web Journal*, 2012.
105. Niels Landwehr, Mark Hall, and Eibe Frank. Logistic model trees. *Machine Learning*, 95(1-2):161–205, 2005.
106. S. le Cessie and J.C. van Houwelingen. Ridge estimators in logistic regression. *Applied Statistics*, 41(1):191–201, 1992.
107. Jens Lehmann. Hybrid learning of ontology classes. In *Proc. of the 5th Int. Conference on Machine Learning and Data Mining MLDM*, volume 4571 of *Lecture Notes in Computer Science*, pages 883–898. Springer, 2007.
108. Jens Lehmann. DL-Learner: learning concepts in description logics. *Journal of Machine Learning Research (JMLR)*, 10:2639–2642, 2009.
109. Jens Lehmann. *Learning OWL Class Expressions*. PhD thesis, University of Leipzig, 2010. PhD in Computer Science.
110. Jens Lehmann. Ontology learning. In *Proceedings of Reasoning Web Summer School*, 2010.
111. Jens Lehmann, Sören Auer, Lorenz Bühmann, and Sebastian Tramp. Class expression learning for ontology engineering. *Journal of Web Semantics*, 9:71 – 81, 2011.
112. Jens Lehmann, Chris Bizer, Georgi Kobilarov, Sören Auer, Christian Becker, Richard Cyganiak, and Sebastian Hellmann. DBpedia - a crystallization point for the web of data. *Journal of Web Semantics*, 7(3):154–165, 2009.
113. Jens Lehmann and Lorenz Bühmann. Autosparql: Let users query your knowledge base. In *Proceedings of ESWC 2011*, 2011.
114. Jens Lehmann, Tim Furche, Giovanni Grasso, Axel-Cyrille Ngonga Ngomo, Christian Schallhart, Andrew Sellers, Christina Unger, Lorenz Bühmann, Daniel Gerber, Konrad Höffner, David Liu, and Sören Auer. Deqa: Deep web extraction for question answering. In *Proceedings of ISWC*, 2012.
115. Jens Lehmann and Pascal Hitzler. Foundations of refinement operators for description logics. In Hendrik Blockeel, Jan Ramon, Jude W. Shavlik, and Prasad Tadepalli, editors, *Inductive Logic Programming, 17th International Conference, ILP 2007, Corvallis, OR, USA, June 19-21, 2007*, volume 4894 of *Lecture Notes in Computer Science*, pages 161–174. Springer, 2007. Best Student Paper Award.

116. Jens Lehmann and Pascal Hitzler. A refinement operator based learning algorithm for the \mathcal{ALC} description logic. In *Inductive Logic Programming, 17th International Conference, ILP 2007, Corvallis, OR, USA, June 19-21, 2007*, volume 4894 of *Lecture Notes in Computer Science*, pages 147–160. Springer, 2007. Best Student Paper Award.
117. Jens Lehmann and Pascal Hitzler. Concept learning in description logics using refinement operators. *Machine Learning journal*, 78(1-2):203–250, 2010.
118. 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. DBpedia - a large-scale, multilingual knowledge base extracted from wikipedia. *Semantic Web Journal*, 2014.
119. Yuanguai Lei, Victoria Uren, and Enrico Motta. A framework for evaluating semantic meta-data. In *4th International Conference on Knowledge Capture*, number 8 in K-CAP '07, pages 135 – 142. ACM, 2007.
120. David Kopcsó, Leo Pipino, Ricahrd Wang and William Rybold. *Developing Measurement Scales for Data-Quality Dimensions*, volume 1. M.E. Sharpe, New York, 2005.
121. Bo Leuf and Ward Cunningham. *The Wiki Way: Collaboration and Sharing on the Internet*. Addison-Wesley Professional, 2001.
122. Francesca A. Lisi. Building rules on top of ontologies for the semantic web with inductive logic programming. *Theory and Practice of Logic Programming*, 8(3):271–300, 2008.
123. Francesca A. Lisi and Floriana Esposito. Learning SHIQ+log rules for ontology evolution. In *SWAP 2008*, volume 426 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2008.
124. Steffen Lohmann, Philipp Heim, Sören Auer, Sebastian Dietzold, and Thomas Riechert. Semantifying requirements engineering – the softwiki approach. In *Proceedings of the 4th International Conference on Semantic Technologies (I-SEMANTICS '08)*, J.UCS, pages 182–185, 2008.
125. Vanessa Lopez, Victoria Uren, Marta Reka Sabou, and Enrico Motta. Cross ontology query answering on the semantic web: an initial evaluation. In *K-CAP '09*, pages 17–24, New York, NY, USA, 2009. ACM.
126. Li Ma, Xingzhi Sun, Feng Cao, Chen Wang, and Xiaoyuan Wang. Semantic enhancement for enterprise data management. In *ISWC2009*, 2009.
127. Michael Martin, Claus Stadler, Philipp Frischmuth, and Jens Lehmann. Increasing the financial transparency of european commission project funding. *Semantic Web Journal*, Special Call for Linked Dataset descriptions(2):157–164, 2013.
128. Y. Matsuo and M. Ishizuka. Keyword Extraction From A Single Document Using Word Co-Occurrence Statistical Information. *International Journal on Artificial Intelligence Tools*, 13(1):157–169, 2004.
129. Diana Maynard, Wim Peters, and Yaoyong Li. Metrics for evaluation of ontology-based information extraction. In *Workshop on “Evaluation of Ontologies for the Web” (EON) at WWW*, May 2006.
130. Brian McBride and Dave Beckett. Rdf/xml syntax specification. W3C Recommendation, February 2004.
131. James McCusker and Deborah McGuinness. Towards identity in linked data. In *Proceedings of OWL Experiences and Directions Seventh Annual Workshop*, 2010.
132. Massimo Mecella, Monica Scannapieco, Antonino Virgillito, Roberto Baldoni, Tiziana Catarci, and Carlo Batini. Managing data quality in cooperative information systems. In *Confederated International Conferences DOA, CoopIS and ODBASE*, pages 486–502, 2002.
133. Christian Meilicke and Heiner Stuckenschmidt. Incoherence as a basis for measuring the quality of ontology mappings. In *3rd International Workshop on Ontology Matching (OM) at the ISWC*, 2008.

134. P Mendes, H Mühleisen, and C Bizer. Sieve: Linked data quality assessment and fusion. In *LWDM*, March 2012.
135. Pablo Mendes, Christian Bizer, Zoltán Miklos, Jean-Paul Calbimonte, Alexandra Moraru, and Giorgos Flouris. D2.1: Conceptual model and best practices for high-quality metadata publishing. Technical report, PlanetData Deliverable, 2012.
136. Ryan Moats. Urn syntax. Internet RFC 2141, May 1997.
137. Mohamed Morsey, Jens Lehmann, Sören Auer, and Axel-Cyrille Ngonga Ngomo. DBpedia SPARQL Benchmark – Performance Assessment with Real Queries on Real Data. In *ISWC 2011*, 2011.
138. Mohamed Morsey, Jens Lehmann, Sören Auer, and Axel-Cyrille Ngonga Ngomo. Usage-Centric Benchmarking of RDF Triple Stores. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI 2012)*, 2012.
139. MA. Mostafavi, Edwards G., and R. Jeansoulin. Ontology-based method for quality assessment of spatial data bases. In *International Symposium on Spatial Data Quality*, volume 4, pages 49–66, 2004.
140. D. Nadeau. *Semi-Supervised Named Entity Recognition: Learning to Recognize 100 Entity Types with Little Supervision*. PhD thesis, University of Ottawa, 2007.
141. David Nadeau and Satoshi Sekine. A survey of named entity recognition and classification. *Linguisticae Investigationes*, 30(1):3–26, January 2007. Publisher: John Benjamins Publishing Company.
142. David Nadeau, Peter Turney, and Stan Matwin. Unsupervised named-entity recognition: Generating gazetteers and resolving ambiguity. pages 266–277, 2006.
143. Felix Naumann. *Quality-Driven Query Answering for Integrated Information Systems*, volume 2261 of *Lecture Notes in Computer Science*. Springer-Verlag, 2002.
144. Axel-Cyrille Ngonga Ngomo. Parameter-free clustering of protein-protein interaction graphs. In *Proceedings of Symposium on Machine Learning in Systems Biology 2010*, 2010.
145. Axel-Cyrille Ngonga Ngomo. A time-efficient hybrid approach to link discovery. In *Proceedings of OM@ISWC*, 2011.
146. Axel-Cyrille Ngonga Ngomo. Link discovery with guaranteed reduction ratio in affine spaces with minkowski measures. In *Proceedings of ISWC*, 2012.
147. Axel-Cyrille Ngonga Ngomo. On link discovery using a hybrid approach. *Journal on Data Semantics*, 1:203 – 217, December 2012.
148. Axel-Cyrille Ngonga Ngomo and Sören Auer. Limes - a time-efficient approach for large-scale link discovery on the web of data. In *Proceedings of IJCAI*, 2011.
149. Axel-Cyrille Ngonga Ngomo, Lorenz Bühmann, Christina Unger, Jens Lehmann, and Daniel Gerber. Sorry, i don’t speak sparql — translating sparql queries into natural language. In *Proceedings of WWW*, 2013.
150. Axel-Cyrille Ngonga Ngomo, Norman Heino, Klaus Lyko, René Speck, and Martin Kaltenböck. Scms - semantifying content management systems. In *ISWC 2011*, 2011.
151. Axel-Cyrille Ngonga Ngomo, Lars Kolb, Norman Heino, Michael Hartung, Sören Auer, and Erhard Rahm. When to reach for the cloud: Using parallel hardware for link discovery. In *Proceedings of ESCW*, 2013.
152. Axel-Cyrille Ngonga Ngomo, Jens Lehmann, Sören Auer, and Konrad Höffner. Raven: Active learning of link specifications. In *Proceedings of the Ontology Matching Workshop (co-located with ISWC)*, 2011.
153. Axel-Cyrille Ngonga Ngomo and Klaus Lyko. Eagle: Efficient active learning of link specifications using genetic programming. In *Proceedings of ESWC*, 2012.
154. Axel-Cyrille Ngonga Ngomo, Klaus Lyko, and Victor Christen. Coala – correlation-aware active learning of link specifications. In *Proceedings of ESWC*, 2013.

155. Axel-Cyrille Ngonga Ngomo and Frank Schumacher. Border flow – a local graph clustering algorithm for natural language processing. In *Proceedings of the 10th International Conference on Intelligent Text Processing and Computational Linguistics (CICLING 2009)*, pages 547–558, 2009. Best Presentation Award.
156. Dat P. T. Nguyen, Yutaka Matsuo, and Mitsuru Ishizuka. Relation extraction from wikipedia using subtree mining. In *AAAI*, pages 1414–1420, 2007.
157. Thuy Nguyen and Min-Yen Kan. Keyphrase Extraction in Scientific Publications. pages 317–326. 2007.
158. Shan-Hwei Nienhuys-Cheng and Ronald de Wolf, editors. *Foundations of Inductive Logic Programming*, volume 1228 of *LNCS*. Springer, 1997.
159. Eyal Oren. SemperWiki: A Semantic Personal Wiki. In Stefan Decker, Jack Park, Dennis Quan, and Leo Sauermann, editors, *Proc. of Semantic Desktop Workshop at the ISWC, Galway, Ireland, November 6*, volume 175, November 2005.
160. Patrick Pantel and Marco Pennacchiotti. Espresso: Leveraging generic patterns for automatically harvesting semantic relations. In *ACL*, pages 113–120. ACL Press, 2006.
161. Youngja Park, Roy J Byrd, and Branimir K Boguraev. Automatic glossary extraction: beyond terminology identification. In *Proceedings of the 19th international conference on Computational linguistics - Volume 1*, COLING '02, pages 1–7, Stroudsburg, PA, USA, 2002. Association for Computational Linguistics.
162. Marius Pasca, Dekang Lin, Jeffrey Bigham, Andrei Lifchits, and Alpa Jain. Organizing and searching the world wide web of facts - step one: the one-million fact extraction challenge. In *proceedings of the 21st national conference on Artificial intelligence - Volume 2*, pages 1400–1405. AAAI Press, 2006.
163. Peter F. Patel-Schneider, Patrick Hayes, and Ian Horrocks. OWL Web Ontology Language - Semantics and Abstract Syntax. W3c:rec, W3C, 10 feb 2004. <http://www.w3.org/TR/owl-semantics/>.
164. Leo L. Pipino, Yang W. Lee, and Richarg Y. Wang. Data quality assessment. *Communications of the ACM*, 45(4), 2002.
165. J. Ross Quinlan. *C4.5: Programs for Machine Learning*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1993.
166. Erhard Rahm. *Schema Matching and Mapping*. Springer-Verlag, Heidelberg (DE), 2011.
167. Erhard Rahm and Philip A. Bernstein. A survey of approaches to automatic schema matching. *The VLDB Journal*, 10:334–350, 2001.
168. Yves Raimond, Christopher Sutton, and Mark Sandler. Automatic interlinking of music datasets on the semantic web. In *1st Workshop about Linked Data on the Web*, 2008.
169. Lev Ratinov, Dan Roth, Doug Downey, and Mike Anderson. Local and global algorithms for disambiguation to wikipedia. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 1375–1384, Portland, Oregon, USA, June 2011. Association for Computational Linguistics.
170. Thomas Riechert, Kim Lauenroth, Jens Lehmann, and Sören Auer. Towards semantic based requirements engineering. In *Proceedings of the 7th International Conference on Knowledge Management (I-KNOW)*, 2007.
171. Thomas Riechert, Ulf Morgenstern, Sören Auer, Sebastian Tramp, and Michael Martin. Knowledge engineering for historians on the example of the catalogus professorum lipsiensis. In Peter F. Patel-Schneider, Yue Pan, Pascal Hitzler, Peter Mika, Lei Zhang, Jeff Z. Pan, Ian Horrocks, and Birte Glimm, editors, *Proceedings of the 9th International Semantic Web Conference (ISWC2010)*, volume 6497 of *Lecture Notes in Computer Science*, pages 225–240, Shanghai / China, 2010. Springer.
172. Christoph Rieß, Norman Heino, Sebastian Tramp, and Sören Auer. EvoPat – Pattern-Based Evolution and Refactoring of RDF Knowledge Bases. In *Proceedings of the 9th Interna-*

- tional Semantic Web Conference (ISWC2010)*, Lecture Notes in Computer Science, Berlin / Heidelberg, 2010. Springer.
173. Sebastian Rudolph. Exploring relational structures via FLE. In Karl Erich Wolff, Heather D. Pfeiffer, and Harry S. Delugach, editors, *ICCS 2004*, volume 3127 of *LNCS*, pages 196–212. Springer, 2004.
 174. Anisa Rula, Matteo Palmonari, Andreas Harth, Steffen Stadtmüller, and Andrea Maurino. On the Diversity and Availability of Temporal Information in Linked Open Data. In *ISWC*, 2012.
 175. Anisa Rula, Matteo Palmonari, and Andrea Maurino. Capturing the Age of Linked Open Data: Towards a Dataset-independent Framework. In *IEEE International Conference on Semantic Computing*, 2012.
 176. Michael Röder, Ricardo Usbeck, Sebastian Hellmann, Daniel Gerber, and Andreas Both. N3 - a collection of datasets for named entity recognition and disambiguation in the nlp interchange format. In *The 9th edition of the Language Resources and Evaluation Conference, 26-31 May, Reykjavik, Iceland*, 2014.
 177. Satya S. Sahoo, Wolfgang Halb, Sebastian Hellmann, Kingsley Idehen, Ted Thibodeau Jr, Sören Auer, Juan Sequeda, and Ahmed Ezzat. A survey of current approaches for mapping of relational databases to rdf, 01 2009.
 178. G. Sampson. How fully does a machine-usable dictionary cover english text. *Literary and Linguistic Computing*, 4(1), 1989.
 179. Leo Sauermann and Richard Cyganiak. Cool uris for the semantic web. W3C Interest Group Note, December 2008.
 180. Sebastian Schaffert. Ikewiki: A semantic wiki for collaborative knowledge management. In *Proceedings of the 1st International Workshop on Semantic Technologies in Collaborative Applications (STICA)*, 2006.
 181. Francois Scharffe, Yanbin Liu, and Chuguang Zhou. Rdf-ai: an architecture for rdf datasets matching, fusion and interlink. In *Proc. IJCAI 2009 IR-KR Workshop*.
 182. Baris Sertkaya. OntocomP system description. In Bernardo Cuenca Grau, Ian Horrocks, Boris Motik, and Ulrike Sattler, editors, *Proceedings of the 22nd International Workshop on Description Logics (DL 2009)*, Oxford, UK, July 27-30, 2009, volume 477 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2009.
 183. Burr Settles. *Active Learning*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers, 2012.
 184. Saeedeh Shekarpour and S.D. Katebi. Modeling and evaluation of trust with an extension in semantic web. *Web Semantics: Science, Services and Agents on the World Wide Web*, 8(1):26 – 36, March 2010.
 185. Pavel Shvaiko and Jerome Euzenat. Ten challenges for ontology matching. Technical report, August 01 2008.
 186. Adam Souzis. Building a Semantic Wiki. *IEEE Intelligent Systems*, 20(5):87–91, 2005.
 187. Dimitrios-Emmanuel Spanos, Periklis Stavrou, and Nikolas Mitrou. Bringing relational databases into the semantic web: A survey. *Semantic Web*, 3(2):169–209, 2012.
 188. Claus Stadler, Jens Lehmann, Konrad Höffner, and Sören Auer. Linkedgeodata: A core for a web of spatial open data. *Semantic Web Journal*, 3(4):333–354, 2012.
 189. Christine Thielen. An approach to proper name tagging for german. In *Proceedings of the EACL-95 SIGDAT Workshop*, 1995.
 190. Sebastian Tramp, Philipp Frischmuth, Timofey Ermilov, and Sören Auer. Weaving a Social Data Web with Semantic Pingback. In P. Cimiano and H.S. Pinto, editors, *Proceedings of the EKAW 2010 - Knowledge Engineering and Knowledge Management by the Masses; 11th October-15th October 2010 - Lisbon, Portugal*, volume 6317 of *Lecture Notes in Artificial Intelligence (LNAI)*, pages 135–149, Berlin / Heidelberg, October 2010. Springer.

191. Sebastian Tramp, Norman Heino, Sören Auer, and Philipp Frischmuth. RDFauthor: Employing RDFa for collaborative Knowledge Engineering. In P. Cimiano and H.S. Pinto, editors, *Proceedings of the EKAW 2010 - Knowledge Engineering and Knowledge Management by the Masses; 11th October-15th October 2010 - Lisbon, Portugal*, volume 6317 of *Lecture Notes in Artificial Intelligence (LNAI)*, pages 90–104, Berlin / Heidelberg, October 2010. Springer.
192. Peter D. Turney. Coherent keyphrase extraction via web mining. In *Proceedings of the 18th international joint conference on Artificial intelligence*, pages 434–439, San Francisco, CA, USA, 2003. Morgan Kaufmann Publishers Inc.
193. Christina Unger, Lorenz Bühmann, Jens Lehmann, Axel-Cyrille Ngonga Ngomo, Daniel Gerber, and Philipp Cimiano. Template-based question answering over rdf data. In *Proceedings of the 21st international conference on World Wide Web*, pages 639–648, 2012.
194. Jacopo Urbani, Spyros Kotoulas, Jason Maassen, Frank van Harmelen, and Henri Bal. Owl reasoning with webpie: calculating the closure of 100 billion triples. In *ESWC2010*, 2010.
195. Johanna Völker and Mathias Niepert. Statistical schema induction. In *ESWC (1)*, pages 124–138, 2011.
196. Johanna Völker and Sebastian Rudolph. Fostering web intelligence by semi-automatic OWL ontology refinement. In *Web Intelligence*, pages 454–460. IEEE, 2008.
197. Johanna Völker, Denny Vrandečić, York Sure, and Andreas Hotho. Learning disjointness. In *ESWC 2007*, volume 4519 of *LNCS*, pages 175–189. Springer, 2007.
198. Julius Volz, Christian Bizer, Martin Gaedke, and Georgi Kobilarov. Discovering and maintaining links on the web of data. In *ISWC 2009*, pages 650–665. Springer, 2009.
199. D. Walker and R. Amsler. The use of machine-readable dictionaries in sublanguage analysis. *Analysing Language in Restricted Domains*, 1986.
200. Yair Wand and Richard Y. Wang. Anchoring data quality dimensions in ontological foundations. *Communications of the ACM*, 39(11):86–95, 1996.
201. Gang Wang, Yong Yu, and Haiping Zhu. Pore: Positive-only relation extraction from wikipedia text. In *ISWC07*, volume 4825 of *LNCS*, pages 575–588, Berlin, Heidelberg, 2007. Springer Verlag.
202. Richard Y. Wang and Diane M. Strong. Beyond accuracy: what data quality means to data consumers. *Journal of Management Information Systems*, 12(4):5–33, 1996.
203. Hiroaki Watanabe and Stephen Muggleton. Can ILP be applied to large dataset? In *ILP 2009*, 2009.
204. William Winkler. The state of record linkage and current research problems. Technical report, Statistical Research Division, U.S. Bureau of the Census, 1999.
205. William Winkler. Overview of record linkage and current research directions. Technical report, Bureau of the Census - Research Report Series, 2006.
206. Dekai Wu, Grace Ngai, and Marine Carpuat. A stacked, voted, stacked model for named entity recognition. In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003 - Volume 4, CONLL '03*, pages 200–203, Stroudsburg, PA, USA, 2003. Association for Computational Linguistics.
207. Harris Wu, Mohammad Zubair, and Kurt Maly. Harvesting social knowledge from folksonomies. In *Proceedings of the seventeenth conference on Hypertext and hypermedia, HYPERTEXT '06*, pages 111–114, New York, NY, USA, 2006. ACM.
208. Yulan Yan, Naoaki Okazaki, Yutaka Matsuo, Zhenglu Yang, and Mitsuru Ishizuka. Unsupervised relation extraction by mining wikipedia texts using information from the web. In *ACL, ACL '09*, pages 1021–1029, 2009.
209. Yang Yu and Jeff Heflin. Extending functional dependency to detect abnormal data in RDF Graphs. In *ISWC*, volume 7031, pages 794–809. Springer Berlin Heidelberg, 2011.

210. Amrapali Zaveri, Dimitris Kontokostas, Mohamed A. Sherif, Lorenz Bühmann, Mohamed Morsey, Sören Auer, and Jens Lehmann. User-driven quality evaluation of dbpedia. In *To appear in Proceedings of 9th International Conference on Semantic Systems, I-SEMANTICS '13, Graz, Austria, September 4-6, 2013*, pages 97–104. ACM, 2013.
211. Amrapali Zaveri, Jens Lehmann, Sören Auer, Mofeed M. Hassan, Mohamed A. Sherif, and Michael Martin. Publishing and interlinking the global health observatory dataset. *Semantic Web Journal*, Special Call for Linked Dataset descriptions(3):315–322, 2013.
212. Amrapali Zaveri, Anisa Rula, Andrea Maurino, Ricardo Pietrobon, Jens Lehmann, and Sören Auer. Quality assessment methodologies for linked open data. Under review, available at <http://www.semantic-web-journal.net/content/quality-assessment-methodologies-linked-open-data>.
213. GuoDong Zhou and Jian Su. Named entity recognition using an hmm-based chunk tagger. In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics*, ACL '02, pages 473–480, Morristown, NJ, USA, 2002. Association for Computational Linguistics.