

# A survey of current Link Discovery frameworks

**Editor(s):** Natasha Noy, Google, USA

**Solicited review(s):** Arnab Dutta, University of Mannheim, Germany; Miriam Fernández, Open University, UK; Kavitha Srinivas, IBM, USA

Markus Nentwig<sup>a,\*</sup>, Michael Hartung<sup>a</sup>, Axel-Cyrille Ngonga Ngomo<sup>b</sup> and Erhard Rahm<sup>a</sup>

<sup>a</sup> *Database Group, University of Leipzig, Augustusplatz 10, 04109 Leipzig, Germany*

*E-mails: [nentwig@informatik.uni-leipzig.de](mailto:nentwig@informatik.uni-leipzig.de), [hartung@informatik.uni-leipzig.de](mailto:hartung@informatik.uni-leipzig.de), [rahm@informatik.uni-leipzig.de](mailto:rahm@informatik.uni-leipzig.de)*

<sup>b</sup> *AKSW, University of Leipzig, Augustusplatz 10, 04109 Leipzig, Germany*

*E-mail: [ngonga@informatik.uni-leipzig.de](mailto:ngonga@informatik.uni-leipzig.de)*

**Abstract.** Links build the backbone of the Linked Data Cloud. With the steady growth in size of datasets comes an increased need for end users to know which frameworks to use for deriving links between datasets. In this survey, we comparatively evaluate current Link Discovery tools and frameworks. For this purpose, we outline general requirements and derive a generic architecture of Link Discovery frameworks. Based on this generic architecture, we study and compare the features of state-of-the-art linking frameworks. We also analyze reported performance evaluations for the different frameworks. Finally, we derive insights pertaining to possible future developments in the domain of Link Discovery.

## 1. Introduction

Over the last years, the Linked Open Data (LOD) Cloud has been the most well-known incarnation of the Linked Data Principles. The intention behind this set of interlinked datasets is to create the initial seed for the machine-readable extension of the current Web dubbed the Data Web. While partly very large datasets are being added to the LOD Cloud on a regular basis (e.g., Linked TCGA [53]), they are only sparsely linked with other datasets. Recent studies show that 44% of the LOD datasets are not connected to other datasets at all [55]. This problem is of major importance as links are central for manifold applications including federated queries [52] and answering complex questions [56,60]. The main reason for this blatant lack of links in the LOD Cloud lies in the creation of links being a very tedious process when carried out manually. This is especially true when dealing with large knowledge bases which contain a very large

number of resources. For example, creating links between DBpedia<sup>1</sup> (4.5 million resources) and Linked-GeoData<sup>2</sup> (1+ million resources) would last several decades if checking whether two resources should be linked lasted 1 ms.

Several software tools and frameworks have already been developed to address the link discovery problem especially to identify semantically equivalent objects in different data sources. The basic intuition behind most of these approaches is to reduce the link discovery problem to a similarity computation problem: Given two sets of resources  $S$  and  $T$ , the goal is to automatically find pairs of resources in  $S \times T$  that should be linked with each other, e.g., according to a `owl:sameAs` relationship. Two main problems arise when dealing with link discovery in this manner: achieving both a high effectiveness and a high efficiency of the linking process. A high effectiveness requires finding (almost) all links between two given sources without deriving incorrect links. Achiev-

---

\*Corresponding author.  
E-mail: [nentwig@informatik.uni-leipzig.de](mailto:nentwig@informatik.uni-leipzig.de).

---

<sup>1</sup><http://dbpedia.org>.

<sup>2</sup><http://linkedgeo.org>.

ing this goal requires finding a suitable link configuration or specification [23,35] specifying the similarity condition(s) two resources  $s \in S$  and  $t \in T$  have to comply with in order to count as a being in the input relation. Even when given a suitable link specification, we have to address the efficiency problem since a naïve implementation which compares all elements of  $S$  with all elements of  $T$  would have a complexity of  $O(|S| \cdot |T|)$ .

Link discovery and the related problems of entity resolution or object matching are being studied extensively. A large number of techniques have already been described in several surveys and books, e.g., [7,15,63]. In contrast to these works, we focus on surveying and comparing the currently available link discovery tools and frameworks. The goal is thus to survey the state-of-the-art in existing solutions which could be applied to solve specific linking tasks. Our comparison is based on numerous criteria derived from major requirements as well as from the steps of a generic link discovery workflow that we will present in the following sections. The workflow takes into account the newest developments in this research area including support for learning-based configurations and human interaction. We will first present a functional comparison of eleven current frameworks. We will consider published performance evaluations for the considered tools including the outcome of instance-level benchmarks of the Ontology Evaluation Alignment Initiative (OAEI). We will try to assess the used evaluation criteria and comparability of the achieved results.

We expect the presented criteria and methodology to be useful to comparatively evaluate additional tools. We plan to continuously extend and update the tool comparison under <http://aksw.org/projects/linkinglod>.

## 2. Problem statement and requirements

### 2.1. Link Discovery problem

The Link Discovery (LD) problem can be described as follows: Given two sets of resources  $S$  and  $T$  (for example about movies) and a relation  $\mathcal{R}$  (e.g., `owl:sameAs` or `dbo:producer`), find all pairs  $(s, t) \in S \times T$  such that  $\mathcal{R}(s, t)$  holds. The result is represented as a set of links called a *mapping*:  $M_{S,T} = \{(a_i, \mathcal{R}, b_j) | a_i \in A, b_j \in B\}$ . Optionally a similarity score ( $sim \in [0, 1]$ ) computed by an LD tool can be added to the entries of mappings to express the con-

fidence of a computed link. In this case, links can be represented as quadruples  $(a_i, \mathcal{R}, b_j, sim(a_i, b_j))$ .

Solving the LD problem is challenging due to the typically large volume and semantic heterogeneity of datasets making it difficult to meet major requirements such as high effectiveness and high efficiency. These and further requirements are part of the LD problem and will be discussed in the next subsection. LD has many similarities with the problem of entity resolution (also called deduplication, reference reconciliation or object matching) that has already been extensively addressed [7,11,28]. In particular, similar techniques for evaluating the similarity between objects and for improving the efficiency can be applied. Still there are significant differences between LD and entity resolution that have led to the development of specific tools for LD. Most entity resolution approaches focus on homogeneous datasets of relatively simple, structured objects, described by a set of single-valued attributes (see for example the benchmark datasets in [29]). By contrast, the resources for LD can be heterogeneous and highly interrelated within the datasets. In particular, resources such as DBpedia or LinkedGeoData usually abide by an ontology, which describes the properties that resources of a certain type can have as well as the relations between the classes that the resources instantiate. Thus, the LD process usually involves an ontology and an instance matching part (see general workflow in Fig. 1). Furthermore, entity resolution techniques focus on finding semantically equivalent objects while LD aims at identifying diverse relations (including `owl:sameAs` as well as domain-specific relations).

### 2.2. Requirements

As mentioned before, supporting a high effectiveness and efficiency are two main requirements for a LD framework. In the following we pose further requirements and desiderata such as low manual effort for configuration and tuning, support for online LD as well as the provision of a powerful infrastructure.

*Effectiveness*: A LD tool should generate mappings of high-quality w.r.t. common measures such as precision, recall and F-measure. Hence, results should be precise, i.e., the links generated by a given framework should be correct (precision). A LD tool should also generate as many as possible links to ensure completeness. In summary, only links between resources that really belong together should be produced. This aim is usually achieved by a combination of different LD

methods. Systems may support rather simple matching techniques such as string similarity comparisons for labels (e.g., [46]) but also complex ones, e.g., by considering the semantic neighborhood of a resource or by reusing already available links [24,42]. Furthermore, a LD tool should support different link types [61]. In our comparison we will evaluate which LD methods are supported by an LD tool and which effectiveness could be demonstrated in benchmark evaluations.

*Efficiency:* A LD tool should be fast and scalable to large datasets, e.g., with hundreds of thousands or millions of resources. A naive, non-scalable approach evaluates all possible pairs of resources (Cartesian product) resulting in a quadratic complexity. Hence, a main efficiency goal is to reduce the search space so that the evaluation of irrelevant pairs of resources is largely avoided. Another general optimization approach is parallel LD on multiple cores or multiple nodes in a cluster. This includes the utilization of modern hardware and infrastructures such as graphical processing units (GPU) or Hadoop-based clusters [35].

*Low Configuration and Tuning Effort:* Achieving a high effectiveness generally demands complex link specifications with the combined use of multiple similarity measures and adequate settings for configuration parameters such as similarity thresholds. Manually specifying such configurations is very difficult and time-consuming so that this effort should largely be reduced by automated approaches. This can be achieved by learning-based methods, e.g., by supervised approaches using training data of matching or non-matching pairs of resources. Alternatively, the LD framework can analyze the datasets, e.g., to select suitable similarity measures or properties to evaluate. In order to really reduce the manual configuration effort, the automated approaches should not introduce a significant extra configuration, e.g., for providing training data or specifying new tuning parameters.

*Online and Offline LD:* In addition to a classical offline execution of LD, applications such as mashups or on-demand query systems demand an online LD to integrate data from several data sources at runtime. Hence, a LD tool should support such a runtime or ad-hoc LD, e.g., by providing an appropriate API. Typically, the number of resources to be linked in this way is small thereby facilitating a sufficiently fast execution.

*Powerful infrastructure:* The support for LD discussed in the previous desiderata requires a set of powerful and easy-to-use tools. In particular, a LD tool should come with flexible libraries with different sim-

ilarity functions, support different performance optimizations and provide different possibilities to access data sources for LD and a graphical user interface to display and configure the workflow. Furthermore, the specified LD workflow should be executable on different platforms, preferably with parallel processing. Besides, mechanisms for collaborative work in groups or crowd-sourcing should be provided to more easily overcome problems like labeling of training data or the generation of gold standards. Overall, a tool should be designed domain-independent but it should be possible to flexibly customize it for specific LD tasks, e.g., linking geographical resources or knowledge from the life sciences.

### 3. LD workflow

Current LD frameworks mostly apply workflows consisting of several steps to perform LD. In most cases, these workflows are instantiations of the generic workflow shown in Fig. 1. This workflow is a generalization of the architecture given by analysing the latter on compared LD frameworks (starting in Section 4). The input of the workflow includes the two datasets (source, target) to be linked, configuration parameters and optional background knowledge resources. The input data may be provided in the form of RDF/OWL dumps or in the form of a SPARQL endpoint for query-based data access. Linking may be restricted to a subset of a data source, e.g., instances of a particular class, as for example a geographic data source contains settlements and there is no need to compare these with actors from a more generic data source such as DBpedia. The configuration input may either be a complete linking specification (e.g., rules for comparing resources) or selected parameters such as similarity thresholds. Training data required for learning-based linking is another kind of configuration input. Optionally, tools can make use of further knowledge resources such as dictionaries or previously determined mappings for reuse. The output of the workflow is the set of found links or correspondences representing a mapping between the source and target datasets.

The generic workflow itself has three main phases: preprocessing, matching (similarity computation) and postprocessing. Preprocessing in turn deals with two important tasks: finalizing the linking specification (configuration) and improving runtime efficiency, e.g., by reducing the search space for similarity computations in the main match phase. Preprocessing may also

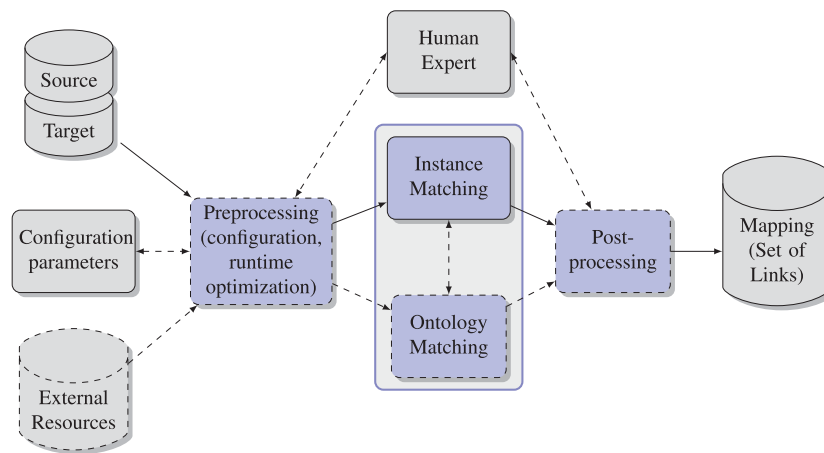


Fig. 1. General workflow LD Frameworks (steps with dashed borders are optional).

include preparatory steps to transform and clean the input data, e.g., to remove stop words or resolve abbreviations. While matching is completely automatic there may be user interaction for preprocessing, e.g., to label training data for learning-based linking, and post-processing, e.g., to verify computed links with a lower confidence.

In the following we describe the two preprocessing steps about configuration and runtime optimization as well as the match and postprocessing phases and their implementation alternatives in more detail. In the tool evaluation we will study which of the different options are applied.

### 3.1. Configuration

LD is typically based on evaluating the similarity of resources according to one or several criteria. Each criterion is based on a specific similarity measure or similarity function and compares either properties or the semantic context of resources. For example two movies may be linked by a `owl:sameAs` property based on the similarity of their titles, their release years and the set of actors who starred in them. Specifying a linking configuration thus entails the specification of the elements (properties, context) to evaluate as well as the similarity measures to apply (e.g., a 3-gram string similarity, Jaccard similarity for sets or numerical difference) and a way to derive a combined linking decision from the individual similarity values, e.g., based on similarity thresholds to meet.

According to [28], different similarity values may be combined either numerically or using rule-based or workflow-based approaches. *Numerical approaches*

aggregate different similarity values, e.g., by taking a weighted average, and apply a single similarity threshold to the aggregated value. *Rule-based approaches* use so-called *match rules* to derive a match or link decision. Such rules define logical combinations of conditions, e.g., 3-gram similarity for title > 0.9 and equal release year. *Workflow-based approaches* are less common and assume the iterative calculation of different similarity values during the match phase to determine a link decision. For example, one could first calculate the string similarity for a selected property and then apply a more expensive context-based similarity measure (e.g., for the set of movie actors) only for pairs of resources with a high similarity for the first criterion [34,59].

A manual definition of effective linking specifications such as match rules is difficult to achieve in many cases even for domain experts. Hence, it is desirable to automate at least some of the decisions such as selecting the properties or the similarity measures to evaluate. This is achieved by *adaptive LD approaches* that analyze characteristics of the input data to achieve a partially automated specification of the linking configuration [3,18,37,40,64].

Alternatively, *learning-based approaches* can be applied to semi-automatically or automatically derive a linking specification. The proposed learning approaches for this purpose are mostly supervised, i.e., they depend on suitable training data consisting of pairs of resources which are labeled as matching (linking) or non-matching. The learned classification model may be based on different learning techniques such as decision trees, SVM or genetic algorithms. Labeling training data is often a manual step

requiring the interaction of humans. The manual labeling effort may be kept feasible by crowdsourcing. Alternatively, the amount of training can be limited by *active learning* where user feedback is only requested for a smaller amount of controversial pairs where a similarity function cannot find a clear linking decision. Learning-based approaches may also be unsupervised, thereby avoiding the need for training data. However, these approaches may still require the specification of critical parameters such as suitable similarity or distance measures and threshold values [37,43].

### 3.2. Runtime optimization

The main approach to optimize the runtime for LD during preprocessing is a reduction of the search space to avoid that the Cartesian Product of the input datasets  $S \times T$  needs to be evaluated. This is mainly supported by two complementary approaches called blocking and filtering. *Blocking* partitions the datasets into multiple partitions or blocks such that links are only determined between resources of the same partition. There are several approaches with disjoint or overlapping partitions for this, e.g., standard blocking (based on a predefined blocking key) or canopy clustering [4,7]. Furthermore, multiple blocking keys may be applied to partition the input according to several criteria so that the likelihood of finding all links is improved. The blocking key is commonly based on attribute or property values, e.g., one could partition movies according to the first three letters of the movie title or according to last name of the movie director. Resources can also be partitioned based on their associated ontology concepts if both data sources have comparable concepts, e.g., the genre of movies. We will call such an approach *concept-based blocking*.

*Filtering* utilizes details of the linking configuration, such as the similarity measure or similarity threshold, to filter pairs of records that cannot meet the similarity condition. For example, token-based string similarity measures such as the Jaccard or Dice similarity can only exceed a certain threshold if the input strings are of similar length and share a certain number of tokens [5]. Preprocessing can support the efficient execution of such filters in the match phase, e.g., by creating a token index.

Blocking and filtering can jointly be applied, e.g., to reduce the number of comparisons for partition-wise linking. Furthermore, both approaches can be utilized in combination with parallel LD [27,35].

### 3.3. Match approaches

The main phase of the LD workflow applies the linking specification and evaluates the specified similarity measures on the pairs of resources that still need to be considered according to the used blocking or filter methods. An LD tool typically has a library of different match techniques (or matchers) that apply a similarity measure on the resources to link with each other. These matchers have been categorized as either element- or structure-based [14,51] depending on whether they evaluate simple resource elements such as atomic property values (literals) or whether they consider the context of resources (e.g., related instances or the ontological context). Element-level matchers are most common and can be based on similarity measures for strings (n-gram, TF/IDF, edit distance, etc.) [6], numbers or domain-specific data types such as geographical coordinates. They are typically applied on matching of comparable properties of resources that have been specified as part of the linking specification (either manually or automatically). Similarity computation may also utilize different kinds of background knowledge such as general-purpose or domain-specific dictionaries and thesauri.

Structure- or context-based matchers are more sophisticated and aim at deriving the similarity of resources from the similarity of their context. There is a large spectrum of possible approaches depending on what context and which similarity computation is applied. For example, some approaches use so-called anchor links between highly similar resources as a seed to iteratively find matching entities in the sets of their related entities [20,24]. The search for matches can also be confined to instances of equivalent or related classes thereby utilizing the ontological context.

A promising LD approach is to utilize already existing links and mappings to find new links. Based on the transitivity of the equality relation one can compose several `owl:sameAs` links to derive new `owl:sameAs` links. Effective strategies for such a composition of mappings and links have been proposed and evaluated in [17]. Public mapping repositories such as BioPortal [49] or LinkLion [31] support the publication of links and thus their reuse for determining new links.

### 3.4. Postprocessing

In the final phase the results of the matchers need to be combined and the links need to be selected from

Table 1  
Considered LD tools

| System/initial publication | Year | Institution                        | Learning-based | OAEI IM participation | Support for pure ontology matching |
|----------------------------|------|------------------------------------|----------------|-----------------------|------------------------------------|
| RiMOM [58]                 | 2004 | Univ. of Tsinghua, China           |                | ✓                     | ✓                                  |
| KnoFuss [44]               | 2007 | Open Univ. Milton Keynes, UK       | ✓              |                       |                                    |
| AgreementMaker [8]         | 2009 | Univ. of Illinois at Chicago, USA  |                | ✓                     | ✓                                  |
| Silk [61]                  | 2009 | FU Berlin, Germany                 | ✓              |                       |                                    |
| CODI [42]                  | 2010 | Univ. of Mannheim, Germany         |                | ✓                     | ✓                                  |
| LIMES [32]                 | 2011 | Univ. of Leipzig, Germany          | ✓              |                       |                                    |
| LogMap [24]                | 2011 | Univ. of Oxford, UK                |                | ✓                     | ✓                                  |
| SERIMI [3]                 | 2011 | Delft Univ. of Techn., Netherlands |                | ✓                     |                                    |
| Zhishi.links [46]          | 2011 | Shanghai Jiao Tong Univ., China    |                | ✓                     |                                    |
| SLINT+ [41]                | 2012 | Nat. Inst. of Informatics, Japan   |                | ✓                     |                                    |
| RuleMiner [45]             | 2012 | Shanghai Jiao Tong Univ., China    | ✓              |                       |                                    |

Notes: Sorted by year of initial publication

the set of candidate links according to the linking specification, e.g., by applying a match rule or a learned classification model. The resulting links may be further refined or repaired to avoid inconsistencies, such as the violation of ontological or application-specific constraints. For example, one could request a 1:1 mapping so that each instance is linked with at most one instance of the other input dataset. Hence, postprocessing could enforce this restriction by selecting the best link per instance, e.g., with the highest computed confidence value. Human feedback is generally helpful during postprocessing to verify the correctness of computed links.

#### 4. Functional comparison

In this section, we provide a functional comparison of eleven state-of-the-art frameworks for LD based on the requirements and the general LD workflow discussed in the previous sections. The selection of tools was further based on the following criteria:

- participation in the OAEI instance matching benchmark track with relatively good performance; or
- learning-based approach for LD and published evaluation results.

Table 1 lists the considered frameworks with their originating organization, their first LD-related publication and further criteria that allows a rough grouping of the tools. Seven of the tools have participated in the instance matching contest of the OAEI. The remaining four frameworks (Silk, LIMES, KnoFuss and RuleMiner) support among others learning-based ap-

proaches for determining linking specifications. A further criterion indicates that four of the seven tools of the first group have support for pure ontology matching in addition to instance matching. In fact, these frameworks (RiMOM, AgreementMaker, LogMap and CODI) mostly started with ontology matching and supported instance matching later. Due to the generality of the LD workflow and the given requirements the followed approach for tool evaluation and comparison can be easily applied to further LD frameworks.

For the more detailed comparison of the tools we summarize their main features in Tables 2 and 3 for the mentioned two groups of 7 + 4 systems. The considered criteria belong to the following categories largely following the steps of the introduced LD workflow:

- Supported input formats.
- Configuration approach.
- Runtime optimizations.
- Match approaches.
- Postprocessing.
- Support for parallel processing.
- User interface (GUI support) and interaction.
- General availability.

In the following subsections we will discuss these aspects for the different frameworks. Finally we will summarize our observations from the functional comparison and relate these to the posed requirements.

##### 4.1. Data input

Nine of the eleven tools accept the input datasets in RDF file format while two frameworks (AgreementMaker, SERIMI) need to retrieve the data exclusively from SPARQL endpoints. While SPARQL endpoints

Table 2  
Characteristics of proposed LD frameworks

|                             | RiMOM            | AgreementMaker       | CODI                                      | LogMap                                    | SERIMI     | Zhishi.links             | SLINT+             |
|-----------------------------|------------------|----------------------|---|---|------------|--------------------------|--------------------|
| Data Input                  | RDF, OWL         | SPARQL               | RDF, OWL                                  | RDF, OWL                                  | SPARQL     | RDF                      | RDF                |
| Supported linktypes         | owl:sameAs       | owl:sameAs           | owl:sameAs                                | owl:sameAs                                | owl:sameAs | owl:sameAs               | owl:sameAs         |
| Configuration               | adaptive         | manual               | manual                                    | manual                                    | adaptive   | manual                   | adaptive           |
| - matcher combination       | weighted average | weighted combination | weighted average                          | weighted average                          | -          | weighted combination     | weighted average   |
| Runtime optimization        |                  |                      |   |   |            |                          |                    |
| - Blocking                  | -                | -                    | -   | -   | -          | -                        | -                  |
| - Filtering                 | indexing         | indexing             | -   | indexing                                  | -          | indexing                 | indexing           |
| String similarity measures  | ✓                | ✓                    | ✓   | ✓   | ✓          | ✓                        | ✓                  |
| Further similarity measures | -                | -                    | -   | -   | -          | geographical coordinates | inverted disparity |
| Structure matcher           | -                | semantic similarity  | iterative anchor-based mapping generation | iterative anchor-based mapping generation | -          | semantic similarity      | -                  |
| Use of                      |                  |                      |   |   |            |                          |                    |
| - external dictionaries     | ?*               | ?*                   | -   | ?*  | -          | -                        | -                  |
| - existing mappings         | -                | -                    | -   | -   | -          | -                        | -                  |
| Post-processing             | -                | -                    | Coherence checks                          | Inconsistency repair                      | -          | -                        | -                  |
| Parallel processing         | -                | -                    | -   | -   | -          | MapReduce                | -                  |
| GUI/web interface/API       | - / - / -        | ✓ / ? / -            | - / - / -                                 | ✓ / ✓ / -                                 | - / - / -  | - / - / -                | - / - / -          |
| Download Tool/Source        | ✓ / -            | - <sup>1</sup> / -   | ✓ / ✓                                     | ✓ / ✓                                     | ✓ / ✓      | ✓ / -                    | ✓ / -              |
| Open Source project         | -                | -                    | ✓   | ✓   | ✓          | -                        | -                  |

Notes: “-” means not existing, “?” unclear from publication, “\*” supported in respective ontology matching framework, <sup>1</sup> no answer on form submission

support a flexible and dynamic data access they can cause availability and performance problems. In addition to RDF, CODI, LogMap and RiMOM additionally support OWL input files. Access to SPARQL endpoints is also supported by the learning-based tools Silk, LIMES and KnoFuss. Dynamic data access with SPARQL typically uses a restriction to certain classes (e.g., books, settlements) thereby limiting the data volume and search space for finding links. While all frameworks are generic and can thus deal with data from different domains and for different applications some tools have also specifically been used for general web data, e.g., to evaluate a real e-commerce dataset [36] or to support question answering tasks combining Linked Data and web data [30].

Surprisingly, a large number of the considered frameworks does not seem to rely on external background knowledge such as dictionaries or already known links and mappings (except for the use of selected links for training supervised approaches to learn link specifications). This is in strong contrast to ontol-

ogy matching where virtually all current tools utilize dictionaries such as WordNet as background knowledge [50]. The tools RiMOM, AgreementMaker and LogMap also utilize such dictionaries for their ontology matching but apparently not for linking instance data. A possible reason for this situation is the lack of suitable knowledge resources supporting linking at the instance level. Only the LD tool Zhishi.links did use a manually created synonym list, mainly for resolving abbreviations such as (Corp. – Corporation), (NY – New York) [46].

#### 4.2. Configuration

Most frameworks can only determine owl:sameAs links or equivalent instances. LIMES and Silk also support additional link types which need to be manually specified by the tool user.

Four frameworks rely on a purely manually specified linking configuration (CODI, LogMap, AgreementMaker, Zhishi.links). For several matchers the re-

Table 3  
 Characteristics of learning-based LD frameworks

|                                | KnoFuss   | Silk  | LIMES  | RuleMiner  |
|--------------------------------|---|---|--|--|
| Data Input                     | RDF, SPARQL   | RDF, SPARQL, CSV  | RDF, SPARQL, CSV   | RDF  |
| Supported linktypes            | owl:sameAs  | owl:sameAs,<br>user-specified others  | owl:sameAs,<br>user-specified others   | owl:sameAs   |
| Configuration                  | manual (match rules),<br>unsupervised learning<br>(genetic programming) | manual (match rules),<br>supervised learning<br>(genetic programming,<br>active learning) | manual (match rules),<br>supervised learning (genetic<br>programming, active<br>learning), unsupervised<br>(genetic programming) | adaptive (match<br>rules), supervised<br>learning<br>(expectation<br>maximization) |
| Runtime optimization           |   |   |  |  |
| - Blocking                     | -   | multi-dimensional   | -  | -  |
| - Filtering                    | indexing  | -   | space tiling   | indexing   |
| String similarity measures     | ✓   | ✓   | ✓  | ✓  |
| Further similarity<br>measures | -   | numeric, date equality  | geographical coordinates,<br>numeric, date equality  | -  |
| Structure matcher              | -   | -   | -  | semantic similarity  |
| Use of                         |   |   |  |  |
| - external dictionaries        | -   | -   | -  | -  |
| - existing mappings            | -   | -   | -  | -  |
| Post-processing                | one-to-one mapping  | -   | Stable marriage,<br>hospital-resident  | -  |
| Parallel Processing            | -   | MapReduce   | (MapReduce)*   | MapReduce  |
| GUI/web interface/API          | - / - / -   | ✓ / ✓ / ✓   | ✓ / ✓ / ✓  | - / - / -  |
| Download Tool/Source           | ✓ / ✓   | ✓ / ✓   | ✓ / -  | - / -  |
| Open Source project            | ✓   | ✓   | -  | -  |

Notes: “-” means not existing, “\*” investigated in [19], but not available in current release

sulting similarity values are combined according to a weighted average approach or a match rule. The learning-based tools KnoFuss, Silk and LIMES also support manually specified match rules. Four tools (RiMOM, SERIMI, SLINT+, RuleMiner) already follow a semi-automatic, adaptive linking specification by analyzing the datasets and identifying the most discriminating properties. For example, if publications have to be matched, the title will be more discriminating than the venue of the publication. SERIMI is limited to only a single property to be selected for matching. Further parameters such as similarity thresholds have to be manually specified.

Silk, LIMES and RuleMiner support supervised learning of a linking specification. Silk and LIMES employ genetic programming with batch or active learning [21,36]. RuleMiner uses an iterative clustering approach maximizing a likelihood function assuming a close to 1:1 mapping of instances from source to target dataset [45]. Genetic programming starts from a set of random link specifications and uses the evolutionary principles of selection and variation to evolve these specifications until a linking condition meets a

predefined optimization criterion (fitness function) or a maximal number of iterations is reached. For supervised learning, manually labeled link candidates are used within the genetic algorithm to find link specifications that come close to the match decisions for the training data. Active learning aims at reducing the labeling effort for training data and applies an interactive labeling of automatically chosen link candidates [21]. Link candidates for active learning are selected to optimize criteria such as entropy or the similarity correlation to unlabeled instances [38].

KnoFuss and LIMES also implement an unsupervised learning of the linking specification [37,43]. The approaches also utilize genetic programming but try to iteratively optimize measures that evaluate indirect quality criteria such as high similarity values and closeness to a 1:1 mapping (assuming duplicate-free data sources) [37,39,43]. In KnoFuss, the candidate linking specifications aggregate the weighted similarity values for several string matchers and require the aggregated similarity value to exceed a certain threshold. The approach thus has to select the matchers, determine their weights, the aggregation func-



tion (e.g., average or max) and the similarity threshold.

#### 4.3. Runtime optimization

Silk is one of the few frameworks implementing an explicit *blocking* to reduce the search space. They support the (manual) specification of multiple blocking keys, i.e., only instances sharing one of the blocking keys must be compared with each other. A multidimensional index is applied to implement this strategy [23]. An implicit blocking is achieved by preselecting in the input specification the classes to be processed but this does not allow to a-priori reduce the search space for the instances of a class which may be numerous.

The main approach to improve runtime in the considered tools is filtering, especially by utilizing inverted index structures. This optimization focuses mostly on a specific property and similarity measure (matcher). For example token-based string similarity measures such as Jaccard require matching values to share several tokens. Hence all pairs without a common token can be excluded from the comparison. An inverted index allows one to quickly determine the instances that still must be considered. LIMES applies this filtering idea for metric spaces by exploiting the triangular inequality to exclude instances from match comparisons [32]. Newer algorithms implemented in LIMES use space tiling to improve the runtime of measures with Minkowski or orthodromic distances [33]. The idea behind space tiling is to portion the spaces implied by the measures so as to compare the elements of the each tile with a small number of other tiles while ensuring that all links can be found.

#### 4.4. Matching strategies

All tools support element-level matchers on selected properties based on string similarity measures such as edit distance, n-gram, or Jaccard [6]. Only few tools (Zhishi.links, Silk, LIMES) also support built-in numerical similarity measures (e.g., Euclidean distance) or domain-specific measures such as for geographical coordinates. Except SERIMI, all frameworks can match on more than one property [2]. The similarity values of different matchers are combined according to the linking specification (match rule, weighted average or according to a learned linking specification).

In addition to simple matching on property values five frameworks (CODI, LogMap, AgreementMaker, Zhishi.links, RuleMiner) already apply a structural

matching based on the ontology structure to find links. LogMap and CODI apply an iterative anchor-based matching approach. Within the instances of comparable concepts so-called anchor links are determined first between almost identical instances. Both LogMap and CODI then use information from the ontology to iteratively extend the existing mapping by evaluating the similarity of related instances, either utilizing object-property-assertions [20] or logical reasoning [26]. In LogMap the similarity computation is performed by an algorithm called ISUB [57] that combines three different metrics. CODI simply employs a threshold-based edit distance [47].

The structural matching in AgreementMaker is based on its approach used for ontology matching. Zhishi.links applies a two-step matching approach. Initially it determines property-based similarities. The results are filtered via a threshold and the similarities are then semantically refined based on the similarity of related resources in the ontological context [46]. RuleMiner tries to derive the equivalence decision between instances not only from the similarity of property values but also from references to shared instances [45].

#### 4.5. Postprocessing

The main task of postprocessing is to select the links according to the linking specification, e.g., by applying a match rule taking into account the computed similarity values. Additional verification steps are applied by LogMap and CODI to avoid that inconsistent mappings are determined. These tools also support pure ontology matching where such postprocessing steps are quite common. Specifically, LogMap applies logical reasoning [25] and CODI utilizes logical coherence checks to identify links contradicting ontological restrictions [48]. Furthermore, KnoFuss, LIMES and RuleMiner employ postprocessing strategies to ensure that every instance in the source can only have at most one corresponding instance in the target dataset [36,43].

#### 4.6. Support for parallel LD

For high efficiency and scalability, support for parallel LD is beneficial. In addition to utilizing multiple processors of a single node parallel LD may also use several nodes in a distributed cluster, e.g., running Apache Hadoop with MapReduce. Four of the eleven frameworks already support a MapReduce implementation: LIMES [19], Zhishi.links [46], RuleMiner [45]

and Silk.<sup>3</sup> For Zhishi.links and RuleMiner, MapReduce is mandatory. Unfortunately, the use of MapReduce typically incurs a substantial overhead, e.g., for the disk-based exchange of data between machines, which cannot be compensated by parallel processing for smaller datasets. Hence, the use of MapReduce is mainly viable for very large LD tasks (according to [35] for more than  $10^9$  comparisons). Another promising option is using parallel processing on massively parallel graphic processors (GPUs) as already explored in [35]. A restriction for GPU computations is the limited memory of the GPU. Hence, it is mainly promising for smaller problem sizes, e.g., of up to  $10^6$  comparisons [35]. While these optimizations have already been studied in the context of the mentioned tools they are not always an integral part of the available tool versions as they require a specific infrastructure (Hadoop cluster or GPU).

#### 4.7. User interface and interaction

User interfaces for the eleven frameworks range from simple command line interfaces (with diverging sets of options) over stand-alone installations to web applications. Only four tools (LogMap, AgreementMaker, LIMES, Silk) support a GUI for convenient interactive use (Tables 2 and 3). Furthermore, Silk [22] and LIMES<sup>4</sup> mention the availability of an API to call the LD functionality from other programs.

#### 4.8. Availability for other researchers

As seen in Tables 2 and 3 all tools (except AgreementMaker and RuleMiner) are publicly available; five tools even follow an Open Source strategy.

#### 4.9. Summarizing observations

The considered tools provide a very good general availability providing a rich choice for interested users and researchers. In the following, we will discuss how the described features relate to the requirements for LD frameworks introduced in Section 2.2. We also mention missing features and thus opportunities for future improvement. The discussion may help selecting a tool for use although we cannot make a recommendation for a specific framework. This is also because the main requirements of high effectiveness and high efficiency

would require comparable and meaningful benchmark results. However, this is still an open issue as we will discuss in Section 5.

*Effectiveness:* Effectiveness is mainly influenced by the matchers applied and their configuration and combination. Most of the considered tools only support rather simple property-based matchers; the more advanced structural match techniques are available in five tools. The potential of utilizing already existing links and mappings as well as other background knowledge such as dictionaries is not yet exploited, with the exception of the use of a handcrafted synonym list in Zhishi.links. Support for finding link types other than owl:sameAs is only provided by Silk and LIMES. Except for SERIMI, all frameworks support the combined use of several matchers. Given the difficulty to manually select and configure multiple matchers, adaptive and learning-based configuration approaches may be more effective than manually configured ones although they introduce new difficulties such as the provision of suitable training data.

*Efficiency:* This is mainly addressed by filtering techniques for specific matchers rather than more general blocking approaches to reduce the search space. Parallel processing based on MapReduce is supported by four tools but it is a rather heavy-weight approach requiring a suitable Hadoop cluster environment. Other options such as the use of GPUs or newer Hadoop (in-memory) processing frameworks such as Apache Spark are not yet supported.

*Configuration and tuning effort:* Most tools already support advanced methods for semi-automatic configuration of linking specifications, in four cases based on learning approaches such as genetic programming. The learning-based approaches also allow a manual specification of match rules, thereby providing maximal flexibility.

*Online and Offline LD:* While offline LD is possible with all tools, support for online LD is still limited. Five frameworks can retrieve data at runtime from SPARQL endpoints. Four tools provide a web or graphical user interface to interactively start a LD workflow. From these, only Silk and LIMES allow an interactive configuration via a web interface. An API for external access as desirable to implement online LD in applications such as mashups is only available for Silk and LIMES.

*Powerful infrastructure:* Most frameworks are rather powerful providing many configuration possibilities based on different similarity functions and matchers. As already mentioned four LD frameworks support

<sup>3</sup>[https://www.assembla.com/spaces/silk/wiki/Silk\\_MapReduce](https://www.assembla.com/spaces/silk/wiki/Silk_MapReduce).

<sup>4</sup><http://aksw.org/Projects/LIMES.html>.

Table 4  
OAEI instance matching tasks over the years

|      | Name    | Input Format | Type of problem | Domains                              | LOD Sources                                | Link Type  | Max. # Resources | Tasks |
|------|---------|--------------|-----------------|--------------------------------------|--|------------|------------------|-------|
| 2010 | DI      | RDF          | real            | life sciences                        | diseasome<br>drugbank<br>dailymed<br>sider | equality   | 5,000            | 4     |
|      | IIMB    | OWL          | artificial      | cross-domain                         | Freebase                                   | equality   | 1,416            | 80    |
|      | PR      | RDF, OWL     | artificial      | people<br>geography                  | -  | equality   | 864              | 3     |
| 2011 | DI-NYT  | RDF          | real            | people<br>geography<br>organizations | NYTimes<br>DBpedia<br>Freebase<br>Geonames | equality   | 9,958            | 7     |
|      | IIMB    | OWL          | artificial      | cross-domain                         | Freebase                                   | equality   | 1,500            | 80    |
| 2012 | SB      | OWL          | artificial      | cross-domain                         | Freebase                                   | equality   | 375              | 10    |
|      | IIMB    | OWL          | artificial      | cross-domain                         | Freebase                                   | equality   | 375              | 80    |
| 2013 | RDFT    | RDF          | artificial      | people                               | DBpedia                                    | equality   | 430              | 5     |
| 2014 | id-rec  | OWL          | artificial      | publications                         | ?  | equality   | 2,649            | 1     |
|      | sim-rec | OWL          | artificial      | publications                         | ?  | similarity | 173              | 1     |

Notes: “-” means not existing, “?” unclear from publication

parallel matching using MapReduce. LogMap, AgreementMaker, Silk and LIMES provide GUI support for easy user interaction. The learning-based tools KnoFuss, Silk and LIMES provide the most options for linking configuration and runtime optimization.

## 5. Comparison of evaluation results

In this section, we analyze the published evaluation results for the considered frameworks. Special emphasis is given to results for the Ontology Evaluation Alignment Initiative (OAEI)<sup>5</sup> in the instance matching track aiming on an evaluation of different systems under the same conditions.

Similarly to previous evaluation studies on entity resolution [7,28] we consider the following criteria:

- Format of input data (RDF, OWL, etc.).
- Determined link types.
- Real vs. artificial (synthetic) datasets: artificial datasets are typically created by systematically changing real instances to create similar (matching) instances to identify by the evaluated approaches. This supports the generation of large datasets for scalability experiments.
- Considered data sources and domains.

- Effectiveness: achieved linking quality in terms of precision, recall and F-measure w.r.t. a perfect linking result (gold standard).
- Efficiency: runtime results and scalability to large data volumes.

In the following we first describe the results for OAEI instance matching benchmarks which provide the best possible comparability for the different tools so far. Afterwards we briefly discuss observations from additional evaluations and summarize the main findings.

### 5.1. OAEI benchmark tests

The Ontology Evaluation Alignment Initiative (OAEI) performs yearly contests since 2005 to comparatively evaluate current tools for ontology and instance matching. The original focus has been on ontology matching but since 2009 instance matching has also been a regular evaluation track. As already discussed in the previous section, seven of the eleven tools have already participated in this track. Even three of the four learning-based frameworks used some of the OAEI test cases for their evaluations. Despite this situation, the analysis of the results for the OAEI benchmark is made complicated because the tasks and the participating systems change every year.

Table 4 gives an overview over the OAEI instance matching tasks in five contests from 2010 until 2014.

<sup>5</sup><http://www.ontologymatching.org>.

Table 5  
Tool participation in OAEI instance matching tracks over the years

| Task        | AgreementMaker | SERIMI | CODI | Zhishi.links | LogMap | RiMOM | SLINT+ | LIMES | Silk | KnoFuss |
|-------------|----------------|--------|------|--------------|--------|-------|--------|-------|------|---------|
| 2010 PR     |                |        | ✓    |              |        | ✓     |        | ✓*    | ✓*   | ✓*      |
| 2010 IIMB   |                |        | ✓    |              |        | ✓     |        |       |      |         |
| 2010 DI     |                |        |      |              |        | ✓     |        |       |      |         |
| 2011 IIMB   |                |        | ✓    |              |        |       |        |       |      |         |
| 2011 DI-NYT | ✓              | ✓      |      | ✓            |        |       | ✓*     |       | ✓*   | ✓*      |
| 2012 SB     |                |        |      |              | ✓      |       |        |       |      |         |
| 2012 IIMB   |                |        |      |              | ✓      |       |        |       |      |         |
| 2013 RDFT   |                |        |      |              | ✓      | ✓     | ✓      |       |      |         |
| 2014 id-rec |                |        |      |              | ✓      | ✓     |        |       |      |         |

Notes: RuleMiner did not participate in any of the given tasks. “\*” did not participate in OAEI contest

Most tasks have only been used in one year while others like IIMB have been changed in different years. Most tests are based on artificially changed datasets where values and the structural context of instances have been modified in a controlled way. The tests cover different domains (life sciences, people, geography, etc.) and LOD data sources (DBpedia, Freebase, GeoNames, NYTimes, etc.). Frequently the benchmarks consist of several match (linking) tasks to cover a certain spectrum of complexity. The number of instances is rather small in all tests with a maximal size of a data source of 9,958 or fewer instances. The evaluation focus has been solely on the effectiveness (e.g., F-Measure) while runtime efficiency has not been measured. Almost all tasks focus on identifying equivalent instances (`owl:sameAs` links).

We briefly characterize the different OAEI tasks as follows.

*IIMB and Sandbox (SB)* The IIMB benchmark has been part of the 2010, 2011 and 2012 contests and consists of 80 test cases using synthetically modified datasets derived from instances of 29 Freebase concepts. The tests and number of instances vary from year to year but the tests are generally of a very small size (e.g., at most 375 instances in 2012). The Sandbox (SB) benchmark from 2012 is very similar to IIMB but limited to 10 different test cases [1].

*PR (Persons/Restaurant)* This benchmark is based on real person and restaurant instance data which are artificially modified by adding duplicates and variations of property values. The dataset is relatively small with about 500-600 instances in the restaurant data source and even less in the person data source. [12]

*DI-NYT (Data Interlinking – NYT)* This 2011 benchmark includes seven tasks to link about 10,000 in-

stances from the NYT data source to DBpedia, Freebase and GeoNames instances. The perfect match result contains about 31,000 `owl:sameAs` links to be identified [13].

*RDFT* This 2013 benchmark is also of small size (430 instances) and uses several tests with differently modified DBpedia data. For the first time in the OAEI instance matching track, no reference mapping is provided for the actual evaluation task. Instead, training data with an appropriate reference mapping is given for each test case thereby supporting frameworks relying on supervised learning [9].

*OAEI 2014* Two benchmark tasks have to be performed in 2014, the first one (id-rec) requiring the identification of the same real-world book entities (sameAs links). For this purpose, 1,330 book instances have to be matched with 2,649 synthetically modified instances in the target dataset. Data transformations include changes like the substitution of book titles and labels with keywords as well as language transformations. The second task (sim-rec) requires determining the similarity of pairs of instances which do not reflect the same real-world entities. This addresses common preprocessing tasks, e.g., to reduce the search space for LD. In 2014, the central evaluation platform SEALS [16] is used for instance matching, too. Still, no runtime evaluation is provided for the instance matching task. The sim-rec task [10] is not further evaluated in this paper.

## 5.2. Evaluation results of OAEI tasks

Table 5 shows the participation of the considered tools in the different OAEI contests and benchmarks. Overall, many tools participated only once or twice (AgreementMaker, SERIMI, Zhishi.links, SLINT+)

Table 6  
F-Measure results of the OAEI 2010 benchmark PR (Person/Restaurant)

|                    | RiMOM | CODI | KnoFuss* | Silk* | LIMES (unsupervised)* |
|--------------------|-------|------|----------|-------|-----------------------|
| Person1            | 1.00  | 0.91 | 1.00     | -     | 1.00                  |
| Person2            | 0.97  | 0.36 | 0.99     | -     | 0.94                  |
| Restaurant (OAEI)  | 0.81  | 0.72 | 0.78     | -     | -                     |
| Restaurant (fixed) | -     | -    | 0.98     | 0.99  | 0.82                  |

Notes: "\*" result was achieved outside the OAEI contest

Table 7  
F-Measure results for OAEI 2011 benchmark DI-NYT [13]

|                   | AgreementMaker | SERIMI | Zhishi.links | KnoFuss* | Silk* | Slint+* |
|-------------------|----------------|--------|--------------|----------|-------|---------|
| nyt-dbpedia-loc.  | 0.69           | 0.68   | 0.92         | 0.89     | 0.93  | 0.97    |
| nyt-dbpedia-org.  | 0.74           | 0.88   | 0.91         | 0.92     | -     | 0.95    |
| nyt-dbpedia-peo.  | 0.88           | 0.94   | 0.97         | 0.97     | -     | 0.99    |
| nyt-freebase-loc. | 0.85           | 0.91   | 0.88         | 0.93     | -     | 0.95    |
| nyt-freebase-org. | 0.80           | 0.91   | 0.87         | 0.92     | -     | 0.96    |
| nyt-freebase-peo. | 0.96           | 0.92   | 0.93         | 0.95     | -     | 0.99    |
| nyt-geonames      | 0.85           | 0.80   | 0.91         | 0.90     | -     | 0.99    |
| H-mean            | 0.82           | 0.85   | 0.91         | 0.93     | -     | 0.97    |

Notes: H-mean is calculated manually from the single F-measure values of the appropriate publication, "\*" result was achieved outside the OAEI contest

and several benchmarks have only been evaluated by one or two systems (IIMB 2010, 2011 and 2012, SB, DI 2010, id-rec 2014). The learning-based tools have used the PR and DI-NYT benchmarks but not within the contest so that a direct comparability is not given. This is because outside the contest tools could apply a more intensive tuning and utilize additional information such as training data. Our comparison will thus focus on the benchmarks with most participants: PR, DI-NYT and RDFT.

Table 6 shows the reported F-Measure results for the PR benchmark tasks for matching people and restaurant records. The original reference mapping proved to be erroneous so that it was corrected after the OAEI contest making it difficult to compare the achieved results. Within the contest the RiMOM system could clearly outperform the CODI system. The evaluations outside the competition used the corrected reference mapping and show especially good results for KnoFuss. In general, the small size of the linking problems and the achievable F-Measure of 0.98–1.0 indicate that the benchmark tasks are easy to solve.

The F-Measure results for the DI-NYT benchmark in Table 7) indicate a more diverse situation. From the three frameworks participating in the contest, Zhishi-links achieved the best results with consistent F-Measure values between 0.87 and 0.97 for the seven tasks. By contrast, AgreementMaker and SERIMI per-

Table 8  
F-measure results for test cases of OAEI 2013 benchmark RDFT

|        | LogMap | RiMOM2013   | SLINT+      |
|--------|--------|-------------|-------------|
| test01 | 0.80   | <b>1.00</b> | 0.98        |
| test02 | 0.88   | 0.97        | <b>1.00</b> |
| test03 | 0.84   | <b>0.98</b> | 0.92        |
| test04 | 0.80   | <b>0.96</b> | 0.91        |
| test05 | 0.74   | <b>0.96</b> | 0.88        |

formed somewhat worse due to problems for one or two of the tasks. The results reported for the three systems that did not participate in the contest are generally better. The achievable F-measure results for all tasks are between 0.93 and 0.99 indicating that these tasks are also relatively easy to solve.

F-Measure results for RDFT benchmark from the OAEI 2013 contest are summarized in Table 8. Again, the different tasks could be solved to a large degree with maximal F-Measure values between 0.96 and 1.0. The overall best results are achieved by RiMOM followed by SLINT+ and LogMap. The 2014 id-rec task turned out to be much more challenging. From the two participants, RiMOM again outperformed LogMap with a F-Measure result of only 0.56 vs. 0.10.

In summary, most of the OAEI instance benchmarks so far have been of small size and relatively easy to solve or attracted only few frameworks participating in

the contest. RiMOM could outperform competing systems in three different benchmarks. The frameworks using OAEI benchmarks outside the contest achieved generally very good results that unfortunately are not directly comparable with the results for the frameworks participating in the OAEI contests. Runtime values and thus scalability have not yet been evaluated for OAEI instance matching.

### 5.3. Other evaluations

The learning-based frameworks KnoFuss, LIMES, Silk and RuleMiner did not yet participate in the OAEI contest but evaluated their effectiveness and runtime efficiency with their own evaluations. SLINT+ also has been evaluated beyond the OAEI test cases [40]. The used evaluation datasets are either very broad such as DBpedia or Freebase or come from different domains, e.g., life sciences (e.g., DrugBank, LinkedCT, DailyMed), geography (e.g., GeoNames, GeoNames, LinkedGeoData), and publications (e.g., DBLP, BNB). Unfortunately the evaluation studies typically used different test cases with specific configurations so that the results can hardly be compared with each other. For example, Silk [21] and LIMES [36] both evaluate a LinkedMDB-DBpedia dataset but use varying numbers of entities. Similarly, reported execution times strongly depend on the used hardware configuration so that they mainly serve to show the relative performance of the respective system w.r.t. different data sizes and other configuration parameters.

Several of the non-OAEI evaluation tests focus on scalability by analyzing LD for large datasets [32,40,43,45,62]. One example is the evaluation of RuleMiner in [45] with the largest dataset (GeoNames) of over 8 million instances and a mapping size of 317,433 links. The correctness of computed links was manually checked only for a sample of 1000 links to keep the manual effort manageable. However, a comparative evaluation of the scalability for different tools is still missing.

For genetic programming algorithms, efficiency largely depends on the number of needed iterations. As an example, Silk needed 2,558 s for 25 iterations to link DrugBank with DBpedia but already 21,387 s (factor 8) for 50 iterations [21]. The selection phase of the genetic algorithm also faces a quadratic complexity w.r.t. the data volume. Hence, random sampling is applied to reduce the number of possible candidates for the generation of the next population. Again, runtime and quality of the results compete with each other

as shown in [43] where bigger sampling sizes help to achieve a good F-measure at the expense of increased execution times.

Instance-based linking is similar to entity resolution and the comparative evaluation of entity resolution frameworks faces similar challenges than the evaluation of LD frameworks. The study [29] evaluated several entity resolution tools on several real datasets on publications and product offers of e-commerce websites. While the publication-related match tasks were relatively easy to solve, the two e-commerce match tasks turned out to be especially challenging with a maximal F-Measure of only 60 and 71% for the considered tools. These match tasks have also been used to evaluate further tools including LD frameworks such as LIMES, e.g., in [36,38]). Results in [38] confirm the difficulty of the e-commerce match tasks with achieved F-Measure values ranging below 35%.

### 5.4. Observations and outlook

Despite the laudable effort of the OAEI instance matching tracks the comparable evaluation of existing tools for LD is still a largely open challenge. This is mainly because the participation in the OAEI contest has been limited so far and using the OAEI tasks outside the competition limits the comparability of the achieved results as they are typically based on different prerequisites, e.g., the use of training data. Evaluation results on a single system or approach aim at showing their effectiveness and efficiency rather than providing a neutral comparative evaluation between systems. Given the general availability of LD tools it would be a worthwhile investigation to apply them under the same prerequisites on a set of LD tasks similar than in the entity resolution study [29]. Such a study can be facilitated by using the recently proposed Semantic Publishing Instance Benchmark (SPIMBench) [54] which was initiated by the Linked Database Benchmark Council (LDBC).<sup>6</sup> This benchmark synthetically generates RDF datasets of arbitrary size so that it can be used to evaluate the scalability of LD tools. It also determines the perfect mappings to evaluate match effectiveness.

## 6. Conclusion

We investigated eleven LD frameworks and compared their functionality based on a common set of criteria. The criteria cover the main steps such as the

---

<sup>6</sup><http://www.ldbc.eu/>.

configuration of linking specifications and methods for matching and runtime optimization. We also covered general aspects such as the supported input formats and link types, support for a GUI and software availability as open source. We observed that the considered tools already provide a rich functionality with support for semi-automatic configuration including advanced learning-based approaches such as unsupervised genetic programming or active learning. On the other side, we found that most tools still focus on simple property-based match techniques rather than using the ontological context within structural matchers. Furthermore, existing links and background knowledge are not yet exploited in the considered frameworks. More comprehensive support of efficiency techniques is also necessary such as the combined use of blocking, filtering and parallel processing.

We also analyzed comparative evaluations of the LD frameworks to assess their relative effectiveness and efficiency. In this respect the OAEI instance matching track is the most relevant effort and we thus analyzed its match tasks and the tool participation and results for the last years. Unfortunately, the participation has been rather low thereby preventing the comparative evaluation between most of the tools. Moreover, the focus of the contest has been on effectiveness so far while runtime efficiency has not yet been evaluated. To better assess the relative effectiveness and efficiency of LD tools it would be valuable to test them on a common set of benchmark tasks on the same hardware. Given the general availability of the tools and the existence of a considerable set of match task definitions and datasets this should be feasible with reasonable effort.

## References

- [1] J.-L. Aguirre, K. Eckert, J. Euzenat, A. Ferrara, W.R. van Hage, L. Hollink, C. Meilicke, A. Nikolov, D. Ritze, F. Scharffe, P. Shvaiko, O. Sváb-Zamazal, C.T. dos Santos, E. Jiménez-Ruiz, B. Cuenca Grau and B. Zاپilko, Results of the ontology alignment evaluation initiative 2012, in: *Proc. of the 7th International Workshop on Ontology Matching*, Boston, MA, USA, November 11, 2012, P. Shvaiko et al., eds, CEUR Workshop Proceedings, Vol. 946, CEUR-WS.org, 2012, pp. 73–115.
- [2] S. Araujo, A. de Vries and D. Schwabe, SERIMI results for OAEI 2011, in: *Proc. of the 6th International Workshop on Ontology Matching*, Bonn, Germany, October 24, 2011, P. Shvaiko et al., eds, CEUR Workshop Proceedings, Vol. 814, CEUR-WS.org, 2011, pp. 212–219.
- [3] S. Araújo, J. Hidders, D. Schwabe and A.P. de Vries, SERIMI – resource description similarity, RDF instance matching and interlinking, in: *Proc. of the 6th International Workshop on Ontology Matching*, Bonn, Germany, October 24, 2011, P. Shvaiko et al., eds, CEUR Workshop Proceedings, Vol. 814, CEUR-WS.org, 2011, pp. 246–247.
- [4] R. Baxter, P. Christen and T. Churches, A comparison of fast blocking methods for record linkage, in: *Proc. of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, L. Getoor, T.E. Senator, P.M. Domingos and C. Faloutsos, eds, Vol. 3, ACM, New York, NY, USA, 2003, pp. 25–27. doi:10.1007/978-3-319-11257-2\_20.
- [5] R.J. Bayardo, Y. Ma and R. Srikant, Scaling up all pairs similarity search, in: *Proc. of the 16th International Conference on World Wide Web, WWW 2007*, Banff, Alberta, Canada, May 8–12, 2007, C.L. Williamson, M.E. Zurko, P.F. Patel-Schneider and P.J. Shenoy, eds, ACM, New York, NY, USA, 2007, pp. 131–140. doi:10.1145/1242572.1242591.
- [6] M. Cheatham and P. Hitzler, String similarity metrics for ontology alignment, in: *Proc. of the Semantic Web – ISWC 2013 – 12th International Semantic Web Conference, Part II*, Sydney, NSW, Australia, October 21–25, 2013, H. Alani, L. Kagal, A. Fokoue, P.T. Groth, C. Biemann, J.X. Parreira, L. Aroyo, N.F. Noy, C. Welty and K. Janowicz, eds, Lecture Notes in Computer Science, Vol. 8219, Springer, 2013, pp. 294–309. doi:10.1007/978-3-642-41338-4\_19.
- [7] P. Christen, *Data Matching – Concepts and Techniques for Record Linkage, Entity Resolution, and Duplicate Detection*, Data-Centric Systems and Applications, 1st edn, Springer, 2012. doi:10.1007/978-3-642-31164-2.
- [8] I.F. Cruz, F. Palandri Antonelli and C. Stroe, AgreementMaker: Efficient matching for large real-world schemas and ontologies, *Proceedings of the VLDB Endowment* 2(2) (2009), 1586–1589. doi:10.14778/1687553.1687598.
- [9] B. Cuenca Grau, Z. Dragisic, K. Eckert, J. Euzenat, A. Ferrara, R. Granada, V. Ivanova, E. Jiménez-Ruiz, A.O. Kempf, P. Lambrix et al., Results of the ontology alignment evaluation initiative 2013, in: *Proc. of the 8th International Workshop on Ontology Matching Co-Located with the 12th International Semantic Web Conference (ISWC 2013)*, Sydney, Australia, October 21, 2013, P. Shvaiko et al., eds, CEUR Workshop Proceedings, Vol. 1111, CEUR-WS.org, 2013, pp. 61–100.
- [10] Z. Dragisic, K. Eckert, J. Euzenat, D. Faria, A. Ferrara, R. Granada, V. Ivanova, E. Jiménez-Ruiz, A.O. Kempf, P. Lambrix, S. Montanelli, H. Paulheim, D. Ritze, P. Shvaiko, A. Solimando, C.T. dos Santos, O. Zamazal and B. Cuenca Grau, Results of the ontology alignment evaluation initiative 2014, in: *Proc. of the 9th International Workshop on Ontology Matching Collocated with the 13th International Semantic Web Conference (ISWC 2014)*, Riva del Garda, Trentino, Italy, October 20, 2014, P. Shvaiko, J. Euzenat, M. Mao, E. Jiménez-Ruiz, J. Li and A. Ngonga, eds, CEUR Workshop Proceedings, Vol. 1317, CEUR-WS.org, 2014, pp. 61–104.
- [11] A.K. Elmagarmid, P.G. Ipeirotis and V.S. Verykios, Duplicate record detection: A survey, *IEEE Transactions on Knowledge and Data Engineering* 19(1) (2007), 1–16.
- [12] J. Euzenat, A. Ferrara, C. Meilicke, J. Pane, F. Scharffe, P. Shvaiko, H. Stuckenschmidt, O. Sváb-Zamazal, V. Svátek and C.T. dos Santos, Results of the ontology alignment evaluation initiative 2010, in: *Proc. of the 5th International Workshop on Ontology Matching (OM-2010)*, Shanghai, China, November 7, 2010, P. Shvaiko et al., eds, CEUR Workshop Proceedings, Vol. 689, CEUR-WS.org, 2010, pp. 85–117.

- [13] J. Euzenat, A. Ferrara, W.R. van Hage, L. Hollink, C. Meilicke, A. Nikolov, D. Ritze, F. Scharffe, P. Shvaiko, H. Stuckenschmidt, O. Sváb-Zamazal and C.T. dos Santos, Results of the ontology alignment evaluation initiative 2011, in: *Proc. of the 6th International Workshop on Ontology Matching*, Bonn, Germany, October 24, 2011, P. Shvaiko et al., eds, CEUR Workshop Proceedings, Vol. 814, CEUR-WS.org, 2011, pp. 85–113.
- [14] J. Euzenat and P. Shvaiko, *Ontology Matching*, Springer, 2007. doi:10.1007/978-3-540-49612-0.
- [15] A. Ferrara, A. Nikolov and F. Scharffe, Data linking for the Semantic Web, *International Journal on Semantic Web and Information Systems (IJSWIS)* 7(3) (2011), 46–76. doi:10.4018/jswis.2011070103.
- [16] R. García-Castro and S.N. Wrigley, SEALS methodology for evaluation campaigns, Technical report, Universidad Politécnica de Madrid, September 2011.
- [17] M. Hartung, A. Groß and E. Rahm, Composition methods for link discovery, in: *Proc. of the Datenbanksysteme für Business, Technologie und Web (BTW), 15. Fachtagung des GI-Fachbereichs “Datenbanken und Informationssysteme” (DBIS)*, Magdeburg, Germany, 11.–15.3.2013, V. Markl, G. Saake, K.-U. Sattler, G. Hackenbroich, B. Mitschang, T. Härder and V. Köppen, eds, LNI, Vol. 214, GI, 2013, pp. 261–277.
- [18] O. Hassanzadeh, K.Q. Pu, S. Hassas Yeganeh, R.J. Miller, L. Popa, M.A. Hernández and H. Ho, Discovering linkage points over web data, *Proceedings of the VLDB Endowment* 6(6) (2013), 445–456. doi:10.14778/2536336.2536345.
- [19] S. Hillner and A.-C. Ngonga Ngomo, Parallelizing LIMES for large-scale link discovery, in: *Proc. the 7th International Conference on Semantic Systems, I-SEMANTICS 2011*, Graz, Austria, September 7–9, 2011, C. Ghidini, A.-C. Ngonga Ngomo, S.N. Lindstaedt and T. Pellegrini, eds, ACM International Conference Proceeding Series, ACM, New York, NY, USA, 2011, pp. 9–16. doi:10.1145/2063518.2063520.
- [20] J. Huber, T. Sztyley, J. Noessner and C. Meilicke, CODI: Combinatorial optimization for data integration – results for OAEI 2011, in: *Proc. of the 6th International Workshop on Ontology Matching*, Bonn, Germany, October 24, 2011, P. Shvaiko et al., eds, CEUR Workshop Proceedings, Vol. 814, CEUR-WS.org, 2011.
- [21] R. Isele and C. Bizer, Active learning of expressive linkage rules using genetic programming, *Journal of Web Semantics* 23 (2013), 2–15. doi:10.1016/j.websem.2013.06.001.
- [22] R. Isele, A. Jentzsch and C. Bizer, Silk server-adding missing links while consuming Linked Data, in: *Proc. of the First International Workshop on Consuming Linked Data*, Shanghai, China, November 8, 2010, O. Hartig, A. Harth and J. Sequeda, eds, CEUR Workshop Proceedings, Vol. 665, CEUR-WS.org, 2010.
- [23] R. Isele, A. Jentzsch and C. Bizer, Efficient multidimensional blocking for link discovery without losing recall, in: *Proc. of the 14th International Workshop on the Web and Databases 2011, WebDB 2011*, Athens, Greece, June 12, 2011, A. Marian and V. Vassalos, eds, 2011.
- [24] E. Jiménez-Ruiz and B. Cuenca Grau, LogMap: Logic-based and scalable ontology matching, in: *Proc. of the Semantic Web – ISWC 2011 – 10th International Semantic Web Conference, Part I*, Bonn, Germany, October 23–27, 2011, L. Aroyo, C. Welty, H. Alani, J. Taylor, A. Bernstein, L. Kagal, N.F. Noy and E. Blomqvist, eds, Lecture Notes in Computer Science, Vol. 7031, Springer, 2011, pp. 273–288. doi:10.1007/978-3-642-25073-6\_18.
- [25] E. Jiménez-Ruiz, B. Cuenca Grau and I. Horrocks, LogMap and LogMapLt results for OAEI 2013, in: *Proc. of the 8th International Workshop on Ontology Matching Co-Located with the 12th International Semantic Web Conference (ISWC 2013)*, Sydney, Australia, October 21, 2013, P. Shvaiko et al., eds, CEUR Workshop Proceedings, Vol. 1111, CEUR-WS.org, 2013, pp. 131–138.
- [26] E. Jiménez-Ruiz, B. Cuenca Grau, Y. Zhou and I. Horrocks, Large-scale interactive ontology matching: Algorithms and implementation, in: *ECAI 2012 – 20th European Conference on Artificial Intelligence. Including Prestigious Applications of Artificial Intelligence (PAIS-2012) System Demonstrations Track*, Montpellier, France, August 27–31, 2012, L. De Raedt, C. Bessière, D. Dubois, P. Doherty, P. Frasconi, F. Heintz and P.J.F. Lucas, eds, Frontiers in Artificial Intelligence and Applications, Vol. 242, IOS Press, 2012, pp. 444–449. doi:10.3233/978-1-61499-098-7-444.
- [27] L. Kolb, A. Thor and E. Rahm, Dedoop: Efficient deduplication with Hadoop, *Proceedings of the VLDB Endowment* 5(12) (2012), 1878–1881. doi:10.14778/2367502.2367527.
- [28] H. Köpcke and E. Rahm, Frameworks for entity matching: A comparison, *Data & Knowledge Engineering* 69(2) (2010), 197–210. doi:10.1016/j.datak.2009.10.003.
- [29] H. Köpcke, A. Thor and E. Rahm, Evaluation of entity resolution approaches on real-world match problems, *Proceedings VLDB Endowment*, 33(1–2) (September 2010), 484–493. doi:10.14778/1920841.1920904.
- [30] J. Lehmann, T. Furche, G. Grasso, A.-C. Ngonga Ngomo, C. Schallhart, A. Sellers, C. Unger, L. Bühmann, D. Gerber, K. Höffner, D. Liu and S. Auer, DEQA: Deep Web extraction for question answering, in: *Proc. of the Semantic Web – ISWC 2012 – 11th International Semantic Web Conference, Part II*, Boston, MA, USA, November 11–15, 2012, P. Cudré-Mauroux, J. Heflin, E. Sirin, T. Tudorache, J. Euzenat, M. Hauswirth, J.X. Parreira, J. Hendler, G. Schreiber, A. Bernstein and E. Blomqvist, eds, Lecture Notes in Computer Science, Vol. 7650, Springer, 2012, pp. 131–147. doi:10.1007/978-3-642-35173-0\_9.
- [31] M. Nentwig, T. Soru, A.-C. Ngonga Ngomo and E. Rahm, LinkLion: A link repository for the Web of Data, in: *The Semantic Web: ESWC 2014 Satellite Events – ESWC 2014 Satellite Events*, Anissaras, Crete, Greece, May 25–29, 2014, V. Presutti, E. Blomqvist, R. Troncy, H. Sack, I. Papadakis and A. Tordai, eds, Lecture Notes in Computer Science, Vol. 8798, Springer, 2014, pp. 439–443, Revised Selected Papers. doi:10.1007/978-3-319-11955-7\_63.
- [32] A.-C. Ngonga Ngomo, LIMES – a time-efficient approach for large-scale link discovery on the Web of Data, in: *IJCAI 2011, Proc. of the 22nd International Joint Conference on Artificial Intelligence*, Barcelona, Catalonia, Spain, July 16–22, 2011, T. Walsh, ed., IJCAI/AAAI, 2011, pp. 2312–2317.
- [33] A.-C. Ngonga Ngomo, Link discovery with guaranteed reduction ratio in affine spaces with Minkowski measures, in: *Proc. of the Semantic Web – ISWC 2012 – 11th International Semantic Web Conference, Part I*, Boston, MA, USA, November 11–15, 2012, P. Cudré-Mauroux, J. Heflin, E. Sirin, T. Tudorache, J. Euzenat, M. Hauswirth, J.X. Parreira, J. Hendler, G. Schreiber, A. Bernstein and E. Blomqvist, eds, Lec-



- ture Notes in Computer Science, Vol. 7649, Springer, 2012, pp. 378–393. doi:10.1007/978-3-642-35176-1\_24.
- [34] A.-C. Ngonga Ngomo, HELIOS – execution optimization for Link Discovery, in: *Proc. of the Semantic Web – ISWC 2014 – 13th International Semantic Web Conference, Part I*, Riva del Garda, Italy, October 19–23, 2014, P. Mika et al., eds, Lecture Notes in Computer Science, Vol. 8796, Springer, 2014, pp. 17–32. doi:10.1007/978-3-319-11964-9\_2.
- [35] A.-C. Ngonga Ngomo, L. Kolb, N. Heino, M. Hartung, S. Auer and E. Rahm, When to reach for the cloud: Using parallel hardware for Link Discovery, in: *Proc. of the Semantic Web: Semantics and Big Data, 10th International Conference, ESWC 2013*, Montpellier, France, May 26–30, 2013, P. Cimiano et al., eds, Lecture Notes in Computer Science, Vol. 7882, Springer, 2013, pp. 275–289. doi:10.1007/978-3-642-38288-8\_19.
- [36] A.-C. Ngonga Ngomo and K. Lyko, EAGLE: Efficient active learning of link specifications using genetic programming, in: *Proc. of the Semantic Web: Research and Applications – 9th Extended Semantic Web Conference, ESWC 2012*, Heraklion, Crete, Greece, May 27–31, 2012, E. Simperl et al., eds, Lecture Notes in Computer Science, Vol. 7295, Springer, 2012, pp. 149–163. doi:10.1007/978-3-642-30284-8\_17.
- [37] A.-C. Ngonga Ngomo and K. Lyko, Unsupervised learning of link specifications: Deterministic vs. non-deterministic, in: *Proc. of the 8th International Workshop on Ontology Matching Co-Located with the 12th International Semantic Web Conference (ISWC 2013)*, Sydney, Australia, October 21, 2013, P. Shvaiko et al., eds, CEUR Workshop Proceedings, Vol. 1111, CEUR-WS.org, 2013, pp. 25–36.
- [38] A.-C. Ngonga Ngomo, K. Lyko and V. Christen, COALA – correlation-aware active learning of link specifications, in: *Proc. of the Semantic Web: Semantics and Big Data, 10th International Conference, ESWC 2013*, Montpellier, France, May 26–30, 2013, P. Cimiano et al., eds, Lecture Notes in Computer Science, Vol. 7882, Springer, 2013, pp. 442–456. doi:10.1007/978-3-642-38288-8\_30.
- [39] A.-C. Ngonga Ngomo, M.A. Sherif and K. Lyko, Unsupervised Link Discovery through knowledge base repair, in: *Proc. of the Semantic Web: Trends and Challenges – 11th International Conference, ESWC 2014*, Anissaras, Crete, Greece, May 25–29, 2014, V. Presutti et al., eds, Lecture Notes in Computer Science, Vol. 8465, Springer, 2014, pp. 380–394. doi:10.1007/978-3-319-07443-6\_26.
- [40] K. Nguyen, R. Ichise and B. Le, Interlinking linked data sources using a domain-independent system, in: *Proc. of the Semantic Technology, Second Joint International Conference, JIST 2012*, Nara, Japan, December 2–4, 2012, H. Takeda, Y. Qu, R. Mizoguchi and Y. Kitamura, eds, Lecture Notes in Computer Science, Vol. 7774, Springer, 2012, pp. 113–128. doi:10.1007/978-3-642-37996-3\_8.
- [41] K. Nguyen, R. Ichise and B. Le, SLINT: A schema-independent Linked Data interlinking system, in: *Proc. of the 7th International Workshop on Ontology Matching*, Boston, MA, USA, November 11, 2012, P. Shvaiko et al., eds, CEUR Workshop Proceedings, Vol. 946, CEUR-WS.org, 2012, pp. 1–12.
- [42] M. Niepert, C. Meilicke and H. Stuckenschmidt, A probabilistic-logical framework for ontology matching, in: *Proc. of the Twenty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2010*, Atlanta, Georgia, USA, July 11–15, 2010, M. Fox and D. Poole, eds, AAAI Press, 2010, pp. 1413–1418.
- [43] A. Nikolov, M. d’Aquin and E. Motta, Unsupervised learning of Link Discovery configuration, in: *Proc. of the Semantic Web: Research and Applications – 9th Extended Semantic Web Conference, ESWC 2012*, Heraklion, Crete, Greece, May 27–31, 2012, E. Simperl et al., eds, Lecture Notes in Computer Science, Vol. 7295, Springer, 2012, pp. 119–133. doi:10.1007/978-3-642-30284-8\_15.
- [44] A. Nikolov, V. Uren and E. Motta, KnoFuss: A comprehensive architecture for knowledge fusion, in: *Proc. of the 4th International Conference on Knowledge Capture (K-CAP 2007)*, Whistler, BC, Canada, October 28–31, 2007, D.H. Sleeman and K. Barker, eds, ACM, New York, NY, USA, 2007, pp. 185–186. doi:10.1145/1298406.1298446.
- [45] X. Niu, S. Rong, H. Wang and Y. Yu, An effective rule miner for instance matching in a Web of Data, in: *21st ACM International Conference on Information and Knowledge Management, CIKM’12*, Maui, HI, USA, October 29–November 02, 2012, X.-w. Chen, G. Lebanon, H. Wang and M.J. Zaki, eds, ACM, New York, NY, USA, 2012, pp. 1085–1094. doi:10.1145/2396761.2398406.
- [46] X. Niu, S. Rong, Y. Zhang and H. Wang, Zhishi.links results for OAEI 2011, in: *Proc. of the 6th International Workshop on Ontology Matching*, Bonn, Germany, October 24, 2011, P. Shvaiko et al., eds, CEUR Workshop Proceedings, Vol. 814, CEUR-WS.org, 2011, pp. 220–227.
- [47] J. Noessner and M. Niepert, CODI: Combinatorial optimization for data integration – results for OAEI 2010, in: *Proceedings of the 5th International Workshop on Ontology Matching (OM-2010)*, Shanghai, China, November 7, 2010, P. Shvaiko et al., eds, CEUR Workshop Proceedings, Vol. 689, CEUR-WS.org, 2010, pp. 142–149.
- [48] J. Noessner, M. Niepert, C. Meilicke and H. Stuckenschmidt, Leveraging terminological structure for object reconciliation, in: *Proc. of the Semantic Web: Research and Applications, 7th Extended Semantic Web Conference, ESWC 2010, Part II*, Heraklion, Crete, Greece, May 30–June 3, 2010, L. Aroyo, G. Antoniou, E. Hyvönen, A. ten Teije, H. Stuckenschmidt, L. Cabral and T. Tudorache, eds, Lecture Notes in Computer Science, Vol. 6089, Springer, 2010, pp. 334–348. doi:10.1007/978-3-642-13489-0\_23.
- [49] N.F. Noy, N.H. Shah, P.L. Whetzel, B. Dai, M. Dorf, N. Griffith, C. Jonquet, D.L. Rubin, M.-A.D. Storey, C.G. Chute and M.A. Musen, BioPortal: Ontologies and data resources with the click of a mouse, *Nucleic Acids Research* **37**(Web-Server-Issue) (2009), 170–173. doi:10.1093/nar/gkp440.
- [50] E. Rahm, Towards large-scale schema and ontology matching, in: *Schema Matching and Mapping*, Z. Bellahsene, A. Bonifati and E. Rahm, eds, Data-Centric Systems and Applications, Springer, 2011, pp. 3–27. doi:10.1007/978-3-642-16518-4\_1.
- [51] E. Rahm and P.A. Bernstein, A survey of approaches to automatic schema matching, *The VLDB Journal* **10** (2001), 334–350. doi:10.1007/s007780100057.
- [52] M. Saleem and A.-C. Ngonga Ngomo, HiBISCuS: Hypergraph-based source selection for SPARQL endpoint federation, in: *Proc. of the Semantic Web: Trends and Challenges – 11th International Conference, ESWC 2014*, Anissaras, Crete, Greece, May 25–29, 2014, V. Presutti et al., eds, Lecture Notes in Computer Science, Vol. 8465, Springer, 2014, pp. 176–191. doi:10.1007/978-3-319-07443-6\_13.

- [53] M. Saleem, S.S. Padmanabhuni, A.-C. Ngonga Ngomo, J.S. Almeida, S. Decker and H.F. Deus, Linked cancer genome atlas database, in: *I-SEMANTICS 2013 – 9th International Conference on Semantic Systems, ISEM '13*, Graz, Austria, September 4–6, 2013, M. Sabou, E. Blomqvist, T. Di Noia, H. Sack and T. Pellegrini, eds, ACM, New York, NY, USA, 2013, pp. 129–134. doi:[10.1145/2506182.2506200](https://doi.org/10.1145/2506182.2506200).
- [54] T. Saveta, E. Daskalaki, G. Flouris, I. Fundulaki, M. Herschel and A.-C. Ngonga Ngomo, Pushing the limits of instance matching systems: A semantics-aware benchmark for linked data, in: *Proc. of the 24th International Conference on World Wide Web Companion, WWW 2015 – Companion Volume*, Florence, Italy, May 18–22, 2015, A. Gangemi, S. Leonardi and A. Panconesi, eds, ACM, New York, NY, USA, 2015, pp. 105–106. doi:[10.1145/2740908.2742729](https://doi.org/10.1145/2740908.2742729).
- [55] M. Schmachtenberg, C. Bizer and H. Paulheim, Adoption of the linked data best practices in different topical domains, in: *Proc. of the Semantic Web – ISWC 2014 – 13th International Semantic Web Conference, Part I*, Riva del Garda, Italy, October 19–23, 2014, P. Mika et al., eds, Lecture Notes in Computer Science, Vol. 8796, Springer, 2014, pp. 245–260. doi:[10.1007/978-3-319-11964-9\\_16](https://doi.org/10.1007/978-3-319-11964-9_16).
- [56] S. Shekarpour, A.-C. Ngonga Ngomo and S. Auer, Question answering on interlinked data, in: *22nd International World Wide Web Conference, WWW '13*, Rio de Janeiro, Brazil, May 13–17, 2013, D. Schwabe, V.A.F. Almeida, H. Glaser, R.A. Baeza-Yates and S.B. Moon, eds, International World Wide Web Conferences Steering Committee Republic and Canton of Geneva, Switzerland, 2013, pp. 1145–1156.
- [57] G. Stoilos, G. Stamou and S. Kollias, A string metric for ontology alignment, in: *Proc. of the Semantic Web – ISWC 2005, 4th International Semantic Web Conference, ISWC 2005*, Galway, Ireland, November 6–10, 2005, Y. Gil, E. Motta, V.R. Benjamins and M.A. Musen, eds, Lecture Notes in Computer Science, Vol. 3729, Springer, 2005, pp. 624–637. doi:[10.1007/11574620\\_45](https://doi.org/10.1007/11574620_45).
- [58] J. Tang, B.-Y. Liang, J. Li and K. Wang, Risk minimization based ontology mapping, in: *Proc. of the Content Computing, Advanced Workshop on Content Computing, AWCC 2004*, ZhenJiang, JiangSu, China, November 15–17, 2004, C.-H. Chi and K.-Y. Lam, eds, Lecture Notes in Computer Science, Vol. 3309, Springer, 2004, pp. 469–480. doi:[10.1007/978-3-540-30483-8\\_58](https://doi.org/10.1007/978-3-540-30483-8_58).
- [59] A. Thor and E. Rahm, MOMA – a mapping-based object matching system, in: *Online Proc. of the CIDR 2007, Third Biennial Conference on Innovative Data Systems Research*, Asilomar, CA, USA, January 7–10, 2007, 2007, pp. 247–258, [www.cidrdb.org](http://www.cidrdb.org).
- [60] C. Unger, L. Bühmann, J. Lehmann, A.-C. Ngonga Ngomo, D. Gerber and P. Cimiano, Template-based question answering over RDF data, in: *Proc. of the 21st World Wide Web Conference 2012, WWW 2012*, Lyon, France, April 16–20, 2012, A. Mille, F.L. Gandon, J. Misselis, M. Rabinovich and S. Staab, eds, ACM, 2012, pp. 639–648. doi:[10.1145/2187836.2187929](https://doi.org/10.1145/2187836.2187929).
- [61] J. Volz, C. Bizer, M. Gaedke and G. Kobilarov, Discovering and maintaining links on the Web of Data, in: *Proc. of the Semantic Web – ISWC 2009, 8th International Semantic Web Conference, ISWC 2009*, Chantilly, VA, USA, October 25–29, 2009, A. Bernstein, D.R. Karger, T. Heath, L. Feigenbaum, D. Maynard, E. Motta and K. Thirunarayan, eds, Lecture Notes in Computer Science, Vol. 5823, Springer, 2009, pp. 650–665. doi:[10.1007/978-3-642-04930-9\\_41](https://doi.org/10.1007/978-3-642-04930-9_41).
- [62] J. Volz, C. Bizer, M. Gaedke and G. Kobilarov, Silk – a link discovery framework for the Web of Data, in: *Proc. of the WWW2009 Workshop on Linked Data on the Web, LDOW 2009*, Madrid, Spain, April 20, 2009, C. Bizer, T. Heath, T. Berners-Lee and K. Idehen, eds, CEUR Workshop Proceedings, Vol. 20, CEUR-WS.org, 2009.
- [63] S. Wölger, K. Siorpaes, T. Bürger, E. Simperl, S. Thaler and C. Hofer, A survey on data interlinking methods, Technical report, STI Innsbruck, March 2011.
- [64] Q. Zheng, C. Shao, J. Li, Z. Wang and L. Hu, RiMOM2013 results for OAEI 2013, in: *Proc. of the 8th International Workshop on Ontology Matching Co-Located with the 12th International Semantic Web Conference (ISWC 2013)*, Sydney, Australia, October 21, 2013, P. Shvaiko et al., eds, CEUR Workshop Proceedings, Vol. 1111, CEUR-WS.org, 2013, pp. 161–168.