

# MAG: A Multilingual, Knowledge-base Agnostic and Deterministic Entity Linking Approach

Diego Moussallem  
AKSW Research Group  
University of Leipzig  
Germany

moussallem@informatik.uni-leipzig.de

Michael Röder  
Data Science Group  
University of Paderborn  
Germany  
michael.roeder@upb.de

Ricardo Usbeck  
Data Science Group  
University of Paderborn  
Germany  
ricardo.usbeck@upb.de

Axel-Cyrille Ngonga Ngomo  
Data Science Group  
University of Paderborn  
Germany  
axel.ngonga@upb.de

## ABSTRACT

Entity linking has recently been the subject of a significant body of research. Currently, the best performing approaches rely on trained mono-lingual models. Porting these approaches to other languages is consequently a difficult endeavor as it requires corresponding training data and retraining of the models. We address this drawback by presenting a novel multilingual, knowledge-base agnostic and deterministic approach to entity linking, dubbed MAG. MAG is based on a combination of context-based retrieval on structured knowledge bases and graph algorithms. We evaluate MAG on 23 data sets and in 7 languages. Our results show that the best approach trained on English datasets (PBOH) achieves a micro F-measure that is up to 4 times worse on datasets in other languages. MAG on the other hand achieves state-of-the-art performance on English datasets and reaches a micro F-measure that is up to 0.6 higher than that of PBOH on non-English languages.

## CCS CONCEPTS

• **Information systems** → **Information extraction**; • **Computing methodologies** → **Natural language processing**;

## KEYWORDS

Entity Linking; Multilingual; Named Entity Disambiguation

### ACM Reference format:

Diego Moussallem, Ricardo Usbeck, Michael Röder, and Axel-Cyrille Ngonga Ngomo. 2017. MAG: A Multilingual, Knowledge-base Agnostic and Deterministic Entity Linking Approach. In *Proceedings of K-CAP 2017: Knowledge Capture Conference, Austin, TX, USA, December 4–6, 2017*, 8 pages. <https://doi.org/10.1145/3148011.3148024>

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

K-CAP 2017: Knowledge Capture Conference, December 4–6, 2017, Austin, TX, USA  
© 2017 Copyright held by the owner/author(s). Publication rights licensed to Association for Computing Machinery.  
ACM ISBN 978-1-4503-5553-7/17/12...\$15.00  
<https://doi.org/10.1145/3148011.3148024>

## 1 INTRODUCTION

More than one exabyte of data is added to the Web every day.<sup>1</sup> Automatic extraction of knowledge from this data demands the use of efficient Natural Language Processing (NLP) techniques such as text aggregation, text summarization and knowledge extraction. One of the most important NLP tasks is Entity Linking (EL), also known as Named Entity Disambiguation (NED). The goal here is as follows: Given a piece of text, a reference knowledge base  $K$  and a set of entity mentions in that text, map each entity mention to the corresponding resource in  $K$ .

Several challenges have to be addressed when dealing with EL. For example, an entity can have a large number of surface forms (SF) (also known as labels) due to synonymy, acronyms and typos. For example, New York City, NY and Big Apple are all labels for the same entity. Moreover, multiple entities can share the same name due to homonymy and ambiguity. For example, both the state and the city of New York are called New York. Despite the complexity of the endeavor, EL approaches have achieved increasingly better results over the past few years by relying on trained machine learning models (see [45] for an overview). A portion of these approaches claim to be multilingual and rely on using cross-lingual dictionaries. However, our experiments (see Section 4) show that the underlying models being trained on English corpora make them prone to failure when migrated to a different language. For example, while PBOH [15] achieves an average micro F-measure of 0.69 on corpora in English, it only achieves an average micro F-measure of 0.31 on other languages.

We alleviate this problem by presenting MAG, a novel multilingual EL approach. MAG (Multilingual AGDISTIS) is based on concepts similar to those underlying AGDISTIS [44] but goes beyond this approach by relying on time-efficient graph algorithms combined with language-independent features to link entities to a given reference knowledge base. Hence, MAG is knowledge-base agnostic, i.e., it can be deployed on any reference knowledge base  $K$ . Our approach is also deterministic and does not rely on any trained model. Hence, it can be deployed on virtually any language.

The main contributions of this paper can be summarized as follows:

<sup>1</sup><https://tinyurl.com/kx6baxx>

- We present a novel multilingual and deterministic approach for EL which combines lightweight and easily extensible graph-based algorithms with a new context-based retrieval method.
- MAG features an innovative candidate generation method which relies on various filter methods and search types for a better candidate selection.
- We provide a thorough evaluation of our overall system on 23 data sets using the GERBIL platform [45]. Our results show that MAG achieves state-of-the-art performance on English. In addition, MAG outperforms all state-of-the-art approaches on 6 non-English data sets.

MAG was implemented within the AGDISTIS framework.<sup>2</sup> The version of MAG used in this paper and also all experimental results are publicly available<sup>3,4</sup>.

## 2 RELATED WORK

EL approaches can be subdivided into 2 different classes: 1) English-only approaches which are based and evaluated on English data sets and 2) multilingual approaches.

**English-only** In 2013, Van Erp et al. [9] proposed an approach, dubbed NERD-ML, for entity recognition tailored for extracting entities from tweets. This approach relies on entity type classification over a rich feature vector composed of a set of linguistic components. Next to that, a plethora of other approaches including [3, 5, 13, 17, 25, 28, 36, 50] have been developed. These approaches mostly rely on graph algorithms and/or on machine-learning techniques. However, most of them do not offer a webservice nor an implementation that is publicly available and are thus difficult to compare.

**Multilingual** A large number of multilingual approaches have been developed over the years. DBpedia Spotlight [26] combines Named Entity Recognition (NER) and NED models based on a vector-space representation of entities and the use of cosine similarity for performing the disambiguation task. Hoffart et al. [19] present AIDA which is based on the YAGO2 Knowledge Base (KB). Gad-Elrab et al. [14] have implemented AIDA for other languages. KEA [40] is based on a fine-granular context model that takes into account heterogeneous text sources as well as text created by automated multimedia analysis. Dojchinovski et al. [8] presents Entity-classifier.eu which is based on hypernyms and a Wikipedia-based entity classification system which identifies salient words. The input is transformed to a lower dimensional representation keeping the same quality of output for all sizes of input text. In 2014, Zhang et al. [49] presented x-Lisa, a three-step pipeline based on cross-lingual Linked Data lexica that harnesses the multilingual Wikipedia. Navigli et al. [29] proposed Babelfy, which is based on random walks and a densest subgraph algorithm and relies on the BabelNet semantic network [30] as background knowledge. Usbeck et al. [44] presented AGDISTIS, a KB-agnostic entity disambiguation approach based on string similarity measures and the graph-based HITS algorithm. WAT [34] is the successor of TagME [11] which includes a re-design of all TagME components, namely, the spotter, the

disambiguator, and the pruner. In 2015, Consoli and Recupero [4] presented FRED, a novel machine reader which extends TagMe with entity typing capabilities. Zwicklbauer et al. [51] presented DoSer, an approach akin to AGDISTIS but based on entity embeddings. Later, Ganea et al. [15] introduced PBOH, a probabilistic graphical model which uses pairwise Markov Random Fields. Our related work study suggests that only three of the related approaches are multilingual and deterministic without domain restriction. These are VINCULUM [23], AGDISTIS [44] and QVC [48]. However, VINCULUM, QVC and other recent supervised approaches [33, 38] could not be included in our evaluation because neither their code nor a webservice were made available publicly, and the approaches could not be reconstructed based on the respective papers only.

## 3 THE MAG APPROACH

In this section, we present MAG in detail. Throughout this work, we rely on the following formal definition of EL.

*Definition 3.1.* Entity Linking: Let  $\mathcal{E}$  be a set of entities from a KB and  $\mathcal{D}$  be a document containing potential mentions of entities  $\mathbf{m} = (m_1, \dots, m_n)$ . The goal of an entity linking system is to generate an assignment  $\mathcal{F}$  of mentions to entities with  $\mathcal{F}(\mathbf{m}) \in (\mathcal{E} \cup \{\epsilon\})^n$  for the document  $\mathcal{D}$ , where  $\epsilon$  stands for an entity that is not in the KB.

The EL process implemented by MAG consists of two phases. Several indexes are generated during the offline phase. The entity linking per se is carried out during the online phase and consists of two steps: 1) candidate generation and 2) disambiguation. An overview can be found in Figure 1.

### 3.1 Offline Index Creation

MAG relies on the following five indexes: surface forms, person names, rare references, acronyms and context.

**Surface forms.** MAG relies exclusively on structured data to generate surface forms for entities so as to remain KB-agnostic. For each entity in the reference KB, our approach harvests all labels of the said entity as well as its type and indexes them.<sup>5</sup> Additional SFs can be collected from different sources [2, 44].

**Person names** - This index accounts for the variations in names for referencing persons [21] across languages and domains. Persons are referred to by different portions of their names. For example, the artist Beyoncé Giselle Knowles-Carter is often referred to as Beyoncé or Beyoncé Knowles. In Brazil, she is also known as Beyoncé Carter and Beyoncé G. Knowles. Moreover, languages such as Chinese and Japanese put the family name in front of the given name (in contrast to English, where names are written in the reverse order). We address the problem of labelling persons by generating all possible permutations of the words within the known labels of persons and adding them to the index of names.

**Rare references** - This index is created if textual descriptions are available for the resources of interest (e.g., if resources have a `rdfs:comment` property). A large number of textual entity descriptions provide type information pertaining to the resource at hand, as in the example “Michael Joseph Jackson was an American singer

<sup>2</sup><https://github.com/AKSW/AGDISTIS>

<sup>3</sup><https://github.com/AKSW/AGDISTIS/releases/tag/v1.0-mag>

<sup>4</sup><http://faturl.com/magexp/?open&selected=0>

<sup>5</sup>In our implementation, we relied on predicates such as `rdfs:label` and `rdf:type`.

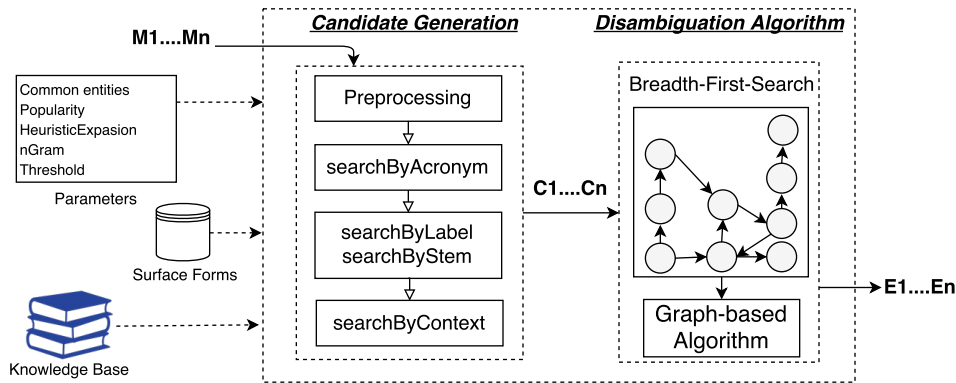


Figure 1: MAG architecture overview.

...<sup>6</sup>. Hence, we use a POS tagger (the Stanford POS tagger [42] in our implementation) on the first line of a resource’s description and collect any noun phrase that contains an adjective. For example, we can extract the supplementary SF American singer for our example. This is similar to [37].

**Acronyms** - Acronyms are used across a large number of domains, e.g., in news (see AIDA and MSNBC data sets). We thus reuse a handcrafted index from STANDS4.<sup>7</sup>

**Context** - Previous works[51] rely on semantic embeddings such as Word2vec [27] to create indexes that model the words surrounding resource mentions in textual corpora. Given that we aim to be KB-agnostic and deterministic, our context index relies on the Concise Bounded Description (CBD)<sup>8</sup> of resources. The literals found in the CBD of each resource are first freed of stop words. Then, each preprocessed string is added as an entry that maps to the said resource.

All necessary information for recreating our indexes or building new indexes for other KBs can be found in our Wiki.<sup>9</sup>

### 3.2 Candidates Generation

The candidate generation and the disambiguation steps occur online, i.e., when MAG is given a document and a set of mentions to disambiguate. The goal of the candidate generation step is to retrieve a tractable number of candidates for each of the mentions. These candidates are later inserted into the disambiguation graph, which is used to determine the mapping between entities and mentions (see Section 3.3).

First, we **preprocess mentions** to improve the retrieval quality. Before using common normalization techniques, we apply a filter for separating acronyms. The acronym filter detects acronyms by looking for strings made up of 5 uppercase letters or less. For example, “PSG” is the acronym of “Paris Saint-Germain”. In the case where a mention is considered an acronym, all further preprocessing steps are skipped. Otherwise, we use basic normalization methods which remove punctuation, symbols and additional white spaces. We use regular expressions to normalize the structure of

strings. For example, our procedure will map strings to the lower case except for the first letter. Therewith, “NEW YORK” is normalized to “New York”. Furthermore, our preprocessing is capable of recognizing mentions such as camel cases “AmyWinehouse” and adds a space between lower-case and upper-case letters (i.e. true casing technique [24]).

The second step of the candidate generation, the **candidate search**, is divided into three parts:

**By Acronym** - If a mention is considered an acronym by our preprocessing, we expand the mention with the list of possible names from the acronym index mentioned above. For example, “PSG” is replaced by “Paris Saint-Germain”.

**By Label** - This search relies on our SF index. First, MAG retrieves candidates for a mention using exact matches to their respective principal reference. For example, the mention “Barack Obama” and the principal reference of the former president of the USA, which is also “Barack Obama”, match exactly. In cases where we find a string similarity match with the main reference of 1.0, the remaining steps are skipped. If this search does not return any candidates, MAG starts a new search using a trigram similarity threshold  $\sigma$  over the SF index. In cases where the set of candidates is still empty, MAG stems the mention and repeats the search. For example, MAG stems “Northern India” to “North India” to account for linguistic variability. The stemming process is not initialized in the preprocessing step since it is a technique which can induce additional errors [39]. For both search types, i.e. search by acronyms and labels, we apply trigram similarity for retrieving possible candidates.

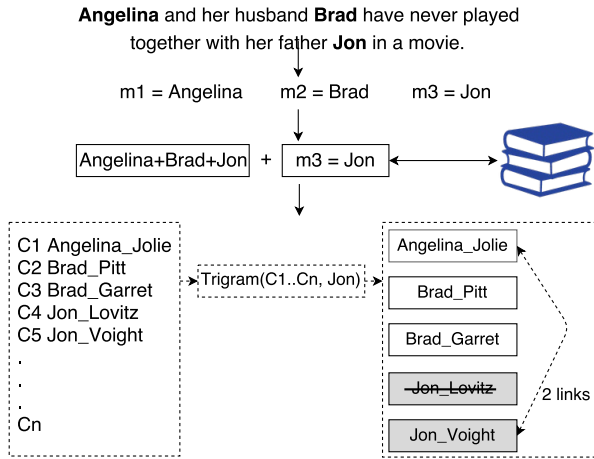
**By Context** - Here, we use two **post-search filters** to find possible candidates from the context index. Before applying both filters, MAG extracts all entities contained in the input document. These entities are used as an addition while searching a mention in the context index. This search relies on TF-IDF [35] which reflects the importance of a word or string in a document corpus relative to the relevance in its index. Afterwards, MAG first filters unlikely candidates by applying trigram similarity. Second, MAG retrieves all direct links among the remaining candidates in the KB. Our approach uses the number of connections to find highly related entity sets for a specific mention. This is similar to finding a dense subgraph [19]. Figure 2 illustrates an example which

<sup>6</sup>See [rdfs:comment of http://dbpedia.org/resource/Michael\\_Jackson](https://www.dbpedia.org/resource/Michael_Jackson).

<sup>7</sup>See <http://www.abbreviations.com/>

<sup>8</sup><https://www.w3.org/Submission/CBD/>

<sup>9</sup><https://github.com/AKSW/AGDISTIS/wiki/3-Running-the-webservice>



**Figure 2: Search using the context index. White boxes on the right side depict candidates discarded by the trigram filter.**

contains three ambiguous entities, namely “Angelina”, “Brad” and “Jon”. Regarding the mention “Jon”, MAG searches the context index using “[Angelina + Brad + Jon] + Jon]” as a query. MAG keeps only “Jon\_Lovitz” and “Jon\_Voight” after trigram filtering. Only “Jon\_Voight”, the father of “Angelina\_Jolie”, has direct connections with the other candidates and is thus chosen.

To improve the quality of candidates, the popularity of a given entity is a good ranking factor. If MAG’s configuration uses this factor, the number of candidates retrieved from the index is increased and then the result is sorted. After this sorting step, MAG returns the top 100 candidates. The popularity is calculated using Page Rank [32] over the underlying KB. If we are unable to leverage Page Rank on certain KB’s, we fall back to a heuristic of inlinks and outlinks. MAG’s candidate generation process is shown in Algorithm 1.<sup>10</sup>

### 3.3 Entity Disambiguation Algorithm

After the candidate generation step, the computation of the optimal candidate to mention assignment starts by constructing a disambiguation graph  $G_d$  with depth  $d$  similar to the approach of AGDISTIS.

*Definition 3.2.* Knowledge Base: We define KB  $K$  as a directed graph  $G_K = (V, E)$  where the nodes  $V$  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 a 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  $G_0$  the algorithm expands the graph using Breadth-First-Search (BFS) technique in order to find hidden paths among candidates. The extension of a graph is  $G_i = (V_i, E_i)$  to a graph  $\rho(G_i) = G_{i+1} = (V_{i+1}, E_{i+1})$  with  $i = 0, \dots, d$ . The  $\rho$  (BFS) operator iterates  $d$  times on the input graph  $G_0$  to compute the initial disambiguation graph  $G_d$ . After constructing  $G_d$ , we need to identify the correct candidate node for a given mention. Here,

<sup>10</sup>The in-document, co-reference resolution is based on earlier works. [44]

---

#### Algorithm 1: Candidates Generation.

---

```

Data: Mention  $m_i$ ,  $\sigma$  trigram similarity threshold
Result:  $C$  candidates found
 $C, \bar{C} \leftarrow \emptyset$ ;
mention  $\leftarrow$  Preprocessing( $m$ ); mention  $\leftarrow$  Co-reference( $m$ );
if containsAcronym( $m$ ) then
   $\bar{C} \leftarrow$  searchByAcronym( $m$ );
else
   $\bar{C} \leftarrow$  searchByLabel( $m$ );
for  $c \in \bar{C}$  do
  if  $\neg c$ .matches([0-9]+) then
    if trigramSimilarity( $c, m$ ) = 1.0 then
       $C \leftarrow C \cup c$ ; return;
    else if trigramSimilarity( $c, m$ )  $\geq \sigma$  then
       $C \leftarrow C \cup c$ ;
   $\bar{C} \leftarrow$  searchByContext( $m$ );
for  $c \in \bar{C}$  do
  if trigramSimilarity( $c, m$ )  $\geq \sigma$  then
    for  $c \in \bar{C}$  do
      if  $c$ .directLinks( $\bar{C}$ )  $\geq 0$  then
         $C \leftarrow C \cup c$ ;

```

---

we rely on HITS [20] or Page Rank [32] as disambiguation graph algorithms. This choice comes from a comparative study of the differences between both [7].

**HITS** uses hub and authority scores to define a recursive relationship between nodes. An authority node is a node that many hubs link to and a hub is a node that links to many authorities. The authority values are equal to the sum of the hub scores of each node that points to it. The hub values are equal to the sum of the authority scores of each node that it points to. According to previous work [44], we chose 20 iterations for HITS which suffice to achieve convergence in general.

**Page Rank** has a wide range of implementations. We implemented the general version in accordance with [32]. Thus, we defined the possibility of jumping from any node to any other node in the graph during the random walk with a probability  $\alpha = (1 - w) = 0.15$ . We empirically chose 50 Page Rank iterations which has shown to be a reasonable number for EL [51]. We assigned a standard weight  $w = 0.85$  for each node. Finally, the sum is calculated by spreading the current weight divided by outgoing edges.

Independent of the chosen graph algorithm, the highest candidate score among the set of candidates  $C$  is chosen as correct disambiguation for a given mention  $m_i$ . The entire process is presented in Algorithm 2. Note, MAG also considers emergent entities [18] and assigns a new URI to them.<sup>11</sup>

*Definition 3.3.* Emergent entity: If the candidate generation step fails to retrieve any candidate from the target KB, we assume the mention belongs to an emergent entity.

<sup>11</sup><https://www.w3.org/TR/cooloris/>

**Algorithm 2:** Disambiguation Algorithm based on HITS and Page Rank.

---

**Data:**  $M = \{m_1, m_2 \dots m_n\}$  mentions,  $d$  depth,  $k$  number of iterations

**Result:**  $C = \{c_1, c_2 \dots c_n\}$  identified candidates for named entities

$E \leftarrow \emptyset;$

$V \leftarrow \text{insertCandidates}(N, \sigma);$

$G \leftarrow (V, E);$

$G \leftarrow \text{breadthFirstSearch}(G, d);$

**if HITS then**

```

  HITS( $G(V, E), k$ );
  sortAccordingToAuthorityValue( $V$ );
  for  $m_i \in M$  do
    for  $v \in V$  do
      if  $v$  is a candidate for  $m_i$  then
        store( $m_i, v$ );
        break;

```

**if Page Rank then**

```

  Page Rank( $G(V, E), k, \alpha$ );
  sortSumValue( $V$ );
  for  $m_i \in M$  do
    for  $v \in V$  do
      if  $v$  is a candidate for  $m_i$  then
        store( $m_i, v$ );
        break;

```

---

## 4 EVALUATION

### 4.1 Goals

The aim of our evaluation is three-fold. First, we aim to measure the performance of MAG on 17 data sets and compare it to the state of the art for EL in English. Second, we evaluate MAG’s portability to other languages. To this end, we compare MAG and the multilingual state of the art using 6 data sets from different languages. For both evaluations we use HITS and Page Rank. Third, we carry out a fine-grained evaluation providing a deep analysis of MAG using the method proposed in [47]. Throughout our experiments, we used DBpedia as reference KB.

### 4.2 Experimental setup

For our evaluation, we rely on the GERBIL platform [45] focusing on the **Disambiguation to KB (D2KB)** experiment type. The task is to map a set of given mentions to entities from a given KB or to  $\epsilon$  (meaning the resource cannot be found in the reference KB).

All data sets (see Table 1 for an overview) are integrated into GERBIL for the sake of comparability.<sup>12</sup> **ACE2004** originates from [36] and is a subset of the ACE co-reference documents. The annotations were obtained through crowd-sourcing where annotators linked the first mention of each reference to Wikipedia. **AIDA/CoNLL** is divided into 3 chunks: Training, TestA and TestB and exclusively contains annotations based on named entities. This manually annotated data set was used to evaluate AIDA, and stems from the

CoNLL 2003 shared task [41]. **AQUAINT** contains annotations of the first mention of each entity in its news-wire documents [28]. **Spotlight** was released along with DBpedia Spotlight [26]. The manually annotated data set contains short texts of named entities and common entities such as cancer and home. **IITB** was created in 2009 and has the highest entity/document-density of all corpora [22]. **KORE50**’s aim is to stress test EL systems through difficult disambiguation tasks using highly ambiguous mentions using hand-crafted sentences.<sup>13</sup> The **Microposts 2014** data sets were created for the "Making Sense of Microposts" challenge and contains only tweets. **MSNBC** was introduced in 2007 by [5]. The data set contains news documents with rare SF and a distinctive lexicalization. **N<sup>3</sup> Reuters-128** comprises 128 news articles which were sampled from the Reuters-21578 news articles data set randomly and annotated manually by domain experts. **N<sup>3</sup> RSS-500** consists of data scrapped from 1,457 RSS feeds [16]. The list includes all major worldwide newspapers and a wide range of topics. The corpus was annotated manually by domain experts. **OKE 2015** was used in the OKE challenge [31]. The data sets were curated manually and are divided into 3 subsets. **N<sup>3</sup> news.de** is a real-world data set collected from 2009 to 2011. It contains documents from the German news portal news.de. **DBpedia Abstracts** is a large, multilingual corpus generated from enriched Wikipedia data of annotated Wikipedia abstracts from six languages [1].<sup>14</sup>

**Table 1: Data set statistics.**

Corpus	Language	Topic	Documents	Entities
ACE2004	English	news	57	253
AIDA/CoNLL-Complete	English	news	1393	34929
AIDA/CoNLL-Test A	English	news	216	5917
AIDA/CoNLL-Test B	English	news	231	5616
AIDA/CoNLL-Training	English	news	946	23396
AQUAINT	English	news	50	747
Spotlight Corpus	English	news	58	330
IITB	English	mixed	103	18308
KORE 50	English	mixed	50	144
Microposts2014-Test	English	tweets	1055	1256
Microposts2014-Train	English	tweets	3395	3822
MSNBC	English	news	20	747
N <sup>3</sup> Reuters-128	English	news	128	880
N <sup>3</sup> RSS-500	English	mixed	500	1000
OKE 2015 Task 1 evaluation set	English	mixed	101	664
OKE 2015 Task 1 example set	English	mixed	3	6
OKE 2015 Task 1 training set	English	mixed	96	338
N <sup>3</sup> news.de	German	news	53	627
Dutch Abstract	Dutch	mixed	39,300	385,259
French Abstract	French	mixed	38,197	346,448
Spanish Abstract	Spanish	mixed	37,663	452,628
Italian Abstract	Italian	mixed	36,432	310,775
Japanese Abstract	Japanese	mixed	38,823	316,982

<sup>13</sup><http://www.yovisto.com/labs/ner-benchmarks/><sup>14</sup>We reduced our test set to the first subset of provided abstracts for each language due to evaluation platform limits and display their characteristics in Table 1.<sup>12</sup>The TAC-KBP data sets could not be included in our evaluation because of their redistribution license.

**4.2.1 Results on English data sets.** For this evaluation, we configured MAG to disambiguate named entities as well as common entities, to use the acronyms index and the popularity scores. According to [44], N-gram at letter level achieved best results when set to 3. The  $\sigma$  trigram threshold was optimal at 0.87<sup>15</sup> and varying the depth of BFS yields an optimal choice of 2. Therefore, we tuned the thresholds to  $d = 2$  and  $\sigma = 0.87$  for performing the evaluation of MAG. The English results are shown in the first part of Table 2.<sup>16</sup>

An analysis of our results shows that although the acronym index is an interesting addition for potential improvements, its contribution amounts only to 0.05% F-measure on average over all data sets. Also, the popularity feature improves the results in almost every data set. It can be explained by the analysis of [47], which demonstrates that most data sets were created using more popular entities as mentions. Thus, this bias eases their retrieval<sup>17</sup>. HITS has shown better results on average than Page Rank.<sup>18</sup> However, Page Rank did show promising results in some data sets (e.g., Spotlight corpus, AQUAINT, and N3-RSS-500).

MAG using HITS outperformed the other approaches on 4 of the 17 data sets while achieving comparable results on others, e.g., ACE2004, MSNBC, and OKE data sets.

Additionally, the performance of MAG can be easily adjusted on each data set by tuning the parameters. For instance, MAG has achieved on *Micropost2014-test* 0.45 F-measure, but the result may increase around 22% to 0.55 F-measure by disabling the context search.<sup>19</sup> This improvement is due to the low entity count (1.81 entities per document) and thus missing links among entities inside a document.

The parameter configuration works not only for huge data sets but also for small data sets or even for disambiguating simple sentences, e.g., "Michael Jordan is a basketball player and Michael Jordan (born 1957), is an American researcher in Machine learning and Artificial intelligence." If the co-reference resolution is turned on, both "Michaels" are linked to the same entity (the basketball player), but if this parameter is turned off, MAG is able to find both correctly.

For the sake of clarity, there is no machine learning in MAG, thus there is no cross-validation set or held-out validation while choosing the parameters.

**4.2.2 Multilingual Results.** Here, we show the easy portability and high quality of MAG for many different languages. Next to German, Italian, Spanish, French and Dutch, we chose Japanese to show the promising potential of MAG across different language systems. MAG's preprocessing NLP techniques are multilingual, thus there is no additional implementation for handling the mentions with different characters. We used the same set of parameters as in the English evaluation but excluded the acronyms as they were only collected for English. Moreover, we performed the Page Rank algorithm over each KB in each respective language to collect

popularity values of their entities. The results displayed in the second part of Table 2 show that MAG, using HITS, outperformed all publicly available state-of-the-art approaches.<sup>20</sup> Also, Page Rank outperformed HITS score on Dutch. The improved performance of MAG is due to its knowledge-base agnostic algorithms and indexing models. For instance, although the mention "Obama" has a high popularity in English, it may have less popularity in Italian or Spanish KBs. Studies about the generation of proper names support this observation [6, 12].

**4.2.3 Fine-Grained Evaluation.** In this analysis, we use an extension of GERBIL done by Waitelonis et al. [47]. This extension provides a fine-grained evaluation which measures the quality of a given EL for linking different types of entities. This extension also considers the assumption of Van Erp et al. [46] that a corpus has a tendency to focus strongly on prominent or popular entities which may cause evaluation problems. Hence, the extension evaluates the capability of a given EL system for finding entities with different levels of popularity thus revealing its degree of bias towards popular entities.

For this evaluation, we used the same set of parameters as for the English evaluation. The fine-grained analysis shows that MAG is better at linking persons than other types of entities. This can be explained by the indexes created by MAG in the offline phase. They collect last names and rare surfaces for entities. In addition, the results show that MAG is not biased towards linking only popular entities as can be seen in Table 3.<sup>21</sup>

## 5 SUMMARY

We presented MAG, a KB-agnostic and deterministic approach for multilingual EL. MAG outperforms the state of the art on all non-English data sets. In addition, MAG achieves a performance similar to the state of the art on English data sets. Nevertheless, an average of 0.63 F-measure places MAG 1st out of 13 annotation systems. As expected, machine learning-based systems retain their advantages due to their tuned training on the provided data sets. However, we are intrigued by our deterministic and knowledge-base independent performance. In this paper, we analyzed the influence of different indexing and searching methods as well as the influence of the data set structure in a fine-grained evaluation. We also provided a context search without relying on machine learning as previously done. Moreover, we showed that current ML-based EL approaches are strongly biased due to their learned model. This behavior can be seen on multilingual data sets. We also deployed and analyzed the influence of acronyms and last names. To the best of our knowledge, no work has investigated this influence on EL and provided a fine-grained evaluation before. In the future, we intend to further investigate disambiguation algorithms and MAG's performance in biomedical and earth science domains using the same data set from Wang et. al [48]. We also aim to evaluate MAG on different languages such as Arabic using other data sets from [43] and export the experiment configurations based on the MEX Vocabulary[10] for reproducibility purposes.

<sup>15</sup>We chose the trigram threshold in accordance with the work of Usbeck. et al. [44] <http://titan.informatik.uni-leipzig.de/rusbeck/agdistis/appendix.pdf>

<sup>16</sup><http://gerbil.aksw.org/gerbil/experiment?id=201701230012> and <http://gerbil.aksw.org/gerbil/experiment?id=201701260017>

<sup>17</sup>see the results without popularity using HITS <http://gerbil.aksw.org/gerbil/experiment?id=201701220014>

<sup>18</sup><http://gerbil.aksw.org/gerbil/experiment?id=201701240030>

<sup>19</sup><http://gerbil.aksw.org/gerbil/experiment?id=201701270035>

<sup>20</sup>For details see <http://faturl.com/multilingualmag/?open>

<sup>21</sup>Detailed information about other data sets can be found here <http://gerbil.s16a.org/gerbil/experiment?id=201701260000>



Table 2: Micro F-measure across approaches. Red entries are the top scores while blue represents the second best scores.

Language		Tools/Data sets	AGDISTs	AIDA	Babely	DBpedia	DoSer	entityclassifer.eu	FRED	Kea	NERD-ML	PROH	WAT	xLisa	MAG + HITS	MAG + PR
English	ACE2004	0.65	0.70	0.53	0.48	<b>0.75</b>	0.50	0.00	0.66	0.58	<b>0.72</b>	0.66	0.70	0.69	0.60	
	AIDA/CoNLL-Complete	0.55	0.68	0.66	0.50	0.69	0.50	0.00	0.61	0.20	<b>0.75</b>	<b>0.71</b>	0.48	0.59	0.54	
	AIDA/CoNLL-Test A	0.54	0.67	0.65	0.48	0.69	0.48	0.00	0.61	0.00	<b>0.75</b>	<b>0.7</b>	0.45	0.59	0.54	
	AIDA/CoNLL-Test B	0.52	0.69	0.68	0.52	0.69	0.48	0.00	0.61	0.00	<b>0.75</b>	<b>0.72</b>	0.47	0.57	0.52	
	AIDA/CoNLL-Training	0.55	0.69	0.66	0.50	0.69	0.52	0.00	0.61	0.28	<b>0.75</b>	<b>0.71</b>	0.48	0.60	0.55	
	AQUAINT	0.52	0.55	0.68	0.53	<b>0.82</b>	0.41	0.00	0.78	0.60	<b>0.81</b>	0.73	0.76	0.67	0.68	
	Spotlight	0.27	0.25	0.52	0.71	<b>0.81</b>	0.25	0.04	0.74	0.56	<b>0.79</b>	0.67	0.71	0.65	0.66	
	IITB	0.47	0.18	0.37	0.30	0.43	0.14	0.00	<b>0.48</b>	0.43	0.38	0.41	0.27	<b>0.52</b>	0.43	
	KORE50	0.27	<b>0.70</b>	<b>0.74</b>	0.46	0.52	0.30	0.06	0.60	0.31	0.63	0.62	0.51	0.24	0.24	
	MSNBC	0.73	0.69	0.71	0.42	<b>0.83</b>	0.51	0.00	0.78	0.62	<b>0.82</b>	0.73	0.5	0.79	0.75	
	Microposts2014-Test	0.33	0.42	0.48	0.50	<b>0.76</b>	0.41	0.05	0.64	0.52	<b>0.73</b>	0.60	0.55	0.45	0.44	
	Microposts2014-Train	0.42	0.51	0.51	0.48	<b>0.77</b>	0.00	0.31	0.65	0.52	<b>0.71</b>	0.63	0.59	0.49	0.44	
	N3-RSS-500	0.66	0.45	0.44	0.20	0.48	0.00	0.00	0.44	0.38	0.53	0.44	0.45	<b>0.69</b>	<b>0.67</b>	
	N3-Reuters-128	0.61	0.47	0.45	0.33	<b>0.69</b>	0.00	0.41	0.51	0.41	<b>0.65</b>	0.52	0.39	<b>0.69</b>	0.64	
	OKE 2015 Task 1 evaluation set	0.59	0.56	0.59	0.31	0.59	0.00	0.46	<b>0.63</b>	0.61	<b>0.63</b>	0.57	<b>0.62</b>	0.58	0.55	
	OKE 2015 Task 1 example set	0.50	<b>0.60</b>	0.40	0.22	0.55	0.00	<b>0.60</b>	0.55	0.00	0.50	<b>0.60</b>	0.50	<b>0.67</b>	0.50	
	OKE 2015 Task 1 training set	0.62	0.67	0.71	0.25	<b>0.78</b>	0.00	0.61	<b>0.78</b>	<b>0.77</b>	0.76	0.72	0.75	0.72	0.70	
Multilingual	N <sup>3</sup> news.de	0.61	0.52	0.50	0.48	0.56	0.28	0.00	0.61	0.33	0.30	0.59	0.36	<b>0.76</b>	<b>0.63</b>	
	Italian Abstracts	0.22	0.28	<b>0.33</b>	0.00	0.00	0.00	0.00	0.00	0.00	0.20	0.00	0.00	<b>0.80</b>	<b>0.80</b>	
	Spanish Abstracts	0.25	0.33	0.26	0.00	0.24	0.27	0.00	0.47	0.00	0.31	0.33	0.31	<b>0.75</b>	<b>0.68</b>	
	Japanese Abstracts	0.15	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	<b>0.38</b>	0.00	0.00	<b>0.54</b>	<b>0.54</b>	
	Dutch Abstracts	0.33	0.36	0.36	0.28	0.36	0.22	0.00	0.40	0.00	0.5	0.40	0.25	<b>0.66</b>	<b>0.67</b>	
French Abstracts	0.00	0.00	<b>0.28</b>	0.22	0.00	0.25	0.00	0.00	0.00	0.20	<b>0.28</b>	<b>0.28</b>	<b>0.80</b>	<b>0.80</b>		
<b>Average</b>		0.45	0.48	0.50	0.36	0.55	0.24	0.11	0.53	0.31	<b>0.59</b>	0.54	0.45	<b>0.63</b>	<b>0.59</b>	
<b>Standard Deviation</b>		0.19	0.22	0.18	0.19	0.27	0.21	0.21	0.23	0.26	0.20	0.22	0.20	0.13	0.13	

Table 3: Fine-grained micro F1 evaluation.

Filter	IITB	N3-RSS-500	MSNBC	Spotlight	N3-Reuters-128	OKE 2015
Persons	0.95	0.83	0.94	0.84	0.80	0.92
Page Rank 10%	0.73	0.67	0.83	0.74	0.79	0.76
Page Rank 10%-55%	0.72	0.72	0.70	0.69	0.73	0.79
Page Rank 55%-100%	0.73	0.71	0.73	0.75	0.76	0.82
Hitsscore 10%	0.77	0.74	0.77	0.69	0.73	0.76
Hitsscore 10%-55%	0.69	0.66	0.64	0.69	0.79	0.78
Hitsscore 55%-100%	0.71	0.66	0.84	0.74	0.77	0.80

## ACKNOWLEDGMENTS

This work has been supported by the H2020 project HOBbit (GA no. 688227) as well as the EuroStars projects DIESEL (no. 01QE1512C) and QAMEL (no. 01QE1549C) and supported by the Brazilian National Council for Scientific and Technological Development (CNPq) (no. 206971/2014-1). This work has also been supported by the German Federal Ministry of Transport and Digital Infrastructure (BMVI) in the projects LIMBO (no. 19F2029I) and OPAL (no. 19F2028A) as well as by the German Federal Ministry of Education and Research (BMBF) within 'KMU-innovativ: Forschung für die zivile Sicherheit' in particular 'Forschung für die zivile Sicherheit' and the project SOLIDE (no. 13N14456).

## REFERENCES

- [1] Martin Brümmer, Milan Dojchinovski, and Sebastian Hellmann. 2016. DBpedia Abstracts: A Large-Scale, Open, Multilingual NLP Training Corpus. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016)*. European Language Resources Association (ELRA), Paris, France.
- [2] Volha Bryl, Christian Bizer, and Heiko Paulheim. 2016. Gathering alternative surface forms for DBpedia entities, In *NLP & DBpedia 2015. CEUR workshop proceedings* 1581, 13–24.
- [3] Andrew Chisholm and Ben Hachey. 2015. Entity disambiguation with web links. *Transactions of the Association for Computational Linguistics* 3 (2015), 145–156.
- [4] Sergio Consoli and Diego Reforgiato Recupero. 2015. Using FRED for Named Entity Resolution, Linking and Typing for Knowledge Base Population. In *Semantic Web Evaluation Challenges: Second SemWebEval Challenge at ESWC 2015, Portorož, May 31 - June 4, 2015*. Springer, Slovenia, 40–50.
- [5] Silviu Cucerzan. 2007. Large-Scale Named Entity Disambiguation Based on Wikipedia Data. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*. Association for Computational Linguistics, Prague, Czech Republic, 708–716.
- [6] Robert Dale and Ehud Reiter. 1995. Computational interpretations of the Gricean maxims in the generation of referring expressions, In *Fourteenth International Conference on Computational Linguistics. Cognitive science* 19, 233–263.
- [7] Pooja Devi, Ashlesha Gupta, and Ashutosh Dixit. 2014. Comparative Study of HITS and PageRank Link based Ranking Algorithms. *International Journal of Advanced Research in Computer and Communication Engineering* 3, 2 (2014), 5749–5754.

- [8] Milan Dojchinovski and Tomas Kliegr. 2013. Entityclassifier.eu: Real-time Classification of Entities in Text with Wikipedia. In *Proceedings of the ECMLPKDD'13*.
- [9] Marieke Van Erp, Giuseppe Rizzo, and Raphaël Troncy. 2013. Learning with the Web: Spotting Named Entities on the Intersection of NERD and Machine Learning. In *Proceedings of the Making Sense of Microposts (#MSM2013) Concept Extraction Challenge*.
- [10] Diego Esteves, Diego Moussallem, Ciro Baron Neto, Tommaso Soru, Ricardo Usbeck, Markus Ackermann, and Jens Lehmann. 2015. MEX vocabulary: a lightweight interchange format for machine learning experiments. In *Proceedings of the 11th International Conference on Semantic Systems*. ACM, 169–176.
- [11] Paolo Ferragina and Ugo Scaiella. 2012. Fast and Accurate Annotation of Short Texts with Wikipedia Pages. *IEEE Software* (2012).
- [12] Thiago Castro Ferreira, Emiel Krahmer, and Sander Wubben. 2017. Generating flexible proper name references in text: Data, models and evaluation. In *Proc. EACL*, Vol. 17.
- [13] Matthew Francis-Landau, Greg Durrett, and Dan Klein. 2016. Capturing semantic similarity for entity linking with convolutional neural networks. *arXiv preprint arXiv:1604.00734* (2016).
- [14] Mohamed H Gad-Elrab, Mohamed Amir Yosef, and Gerhard Weikum. 2015. Named entity disambiguation for resource-poor languages. In *Proceedings of the Eighth Workshop on Exploiting Semantic Annotations in Information Retrieval*. ACM, 29–34.
- [15] Octavian-Eugen Ganea, Marina Ganea, Aurelien Lucchi, Carsten Eickhoff, and Thomas Hofmann. 2016. Probabilistic Bag-Of-Hyperlinks Model for Entity Linking. In *Proceedings of the 25th International Conference on World Wide Web (WWW '16)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, 927–938. <https://doi.org/10.1145/2872427.2882988>
- [16] Daniel Gerber, Axel-Cyrille Ngonga Ngomo, Sebastian Hellmann, Tommaso Soru, Lorenz Bühmann, and Ricardo Usbeck. 2013. Real-time RDF extraction from unstructured data streams. In *ISWC*.
- [17] Xianpei Han, Le Sun, and Jun Zhao. 2011. Collective entity linking in web text: a graph-based method. In *Proceedings of the 34th international ACM SIGIR conference on Research and development in Information Retrieval*. ACM, 765–774.
- [18] Johannes Hoffart, Yasemin Altun, and Gerhard Weikum. 2014. Discovering Emerging Entities with Ambiguous Names. In *Proceedings of the 23rd International Conference on World Wide Web (WWW '14)*. ACM, New York, NY, USA, 385–396.
- [19] Johannes Hoffart, Mohamed Amir Yosef, Ilaria Bordino, Hagen Fürstenau, Manfred Pinkal, Marc Spaniol, Bilyana Taneva, Stefan Thater, and Gerhard Weikum. 2011. Robust Disambiguation of Named Entities in Text. In *Conference on Empirical Methods in Natural Language Processing*.
- [20] Jon M. Kleinberg. 1999. Authoritative sources in a hyperlinked environment. *J. ACM* 46, 5 (1999), 604–632.
- [21] Emiel Krahmer and Kees Van Deemter. 2012. Computational generation of referring expressions: A survey. *Computational Linguistics* 38, 1 (2012), 173–218.
- [22] Sayali Kulkarni, Amit Singh, Ganesh Ramakrishnan, and Soumen Chakrabarti. 2009. Collective annotation of Wikipedia entities in web text. In *15th ACM SIGKDD*. 457–466.
- [23] Xiao Ling, Sameer Singh, and Daniel S Weld. 2015. Design challenges for entity linking. In *Transactions of the Association for Computational Linguistics*, Vol. 3. 315–328.
- [24] Lucian Vlad Lita, Abe Ittycheriah, Salim Roukos, and Nanda Kambhatla. 2003. Truecasing. In *Proceedings of the 41st Annual Meeting on Association for Computational Linguistics-Volume 1*. Association for Computational Linguistics, 152–159.
- [25] Gang Luo, Xiaojiang Huang, Chin-Yew Lin, and Zaiqing Nie. 2015. Joint named entity recognition and disambiguation. In *Proc. EMNLP*.
- [26] Pablo N. Mendes, Max Jakob, Andres Garcia-Silva, and Christian Bizer. 2011. DBpedia Spotlight: Shedding Light on the Web of Documents. In *7th International Conference on Semantic Systems (I-Semantics)*.
- [27] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*. 3111–3119.
- [28] David Milne and Ian H Witten. 2008. Learning to link with wikipedia. In *17th ACM CIKM*. 509–518.
- [29] Andrea Moro, Alessandro Raganato, and Roberto Navigli. 2014. Entity Linking meets Word Sense Disambiguation: A Unified Approach. *TACL* (2014).
- [30] Roberto Navigli and Simone Paolo Ponzetto. 2012. BabelNet: The Automatic Construction, Evaluation and Application of a Wide-Coverage Multilingual Semantic Network. *Artificial Intelligence* (2012).
- [31] Andrea-Giovanni Nuzzolese, AnnaLisa Gentile, Valentina Presutti, Aldo Gangemi, Dario Garigliotti, and Roberto Navigli. 2015. Open Knowledge Extraction Challenge. In *Semantic Web Evaluation Challenges*. Communications in Computer and Information Science, Vol. 548. Springer International Publishing, 3–15.
- [32] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. 1999. *The PageRank citation ranking: Bringing order to the web*. Technical Report. Stanford InfoLab.
- [33] Aasish Pappu, Roi Blanco, Yashar Mehdad, Amanda Stent, and Kapil Thadani. 2017. Lightweight Multilingual Entity Extraction and Linking. In *Proceedings of Tenth International Conference on Web Search and Data Mining*.
- [34] Francesco Piccinno and Paolo Ferragina. 2014. From TagME to WAT: a new entity annotator. In *Proceedings of the first international workshop on Entity recognition & disambiguation*. ACM, 55–62.
- [35] Juan Ramos et al. 2003. Using tf-idf to determine word relevance in document queries. In *Proceedings of the first instructional conference on machine learning*.
- [36] L. Ratinov, D. Roth, D. Downey, and M. Anderson. 2011. Local and Global Algorithms for Disambiguation to Wikipedia. In *ACL*. <http://cogcomp.cs.illinois.edu/papers/RRDA11.pdf>
- [37] Michael Röder, Ricardo Usbeck, René Speck, and Axel-Cyrille Ngonga Ngomo. 2015. CETUS – A Baseline Approach to Type Extraction. In *1st Open Knowledge Extraction Challenge at 12th European Semantic Web Conference (ESWC 2015)*.
- [38] Avirup Sil and Radu Florian. 2016. One for All: Towards Language Independent Named Entity Linking. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*.
- [39] Jasmeet Singh and Vishal Gupta. 2016. Text Stemming: Approaches, Applications, and Challenges. *ACM Computing Surveys (CSUR)* 49, 3 (2016), 45.
- [40] Nadine Steinmetz and Harald Sack. 2013. Semantic Multimedia Information Retrieval Based on Contextual Descriptions. In *The Semantic Web: Semantics and Big Data*, Philipp Cimiano, Oscar Corcho, Valentina Presutti, Laura Hollink, and Sebastian Rudolph (Eds.). Lecture Notes in Computer Science, Vol. 7882. Springer Berlin Heidelberg, 382–396.
- [41] Erik F. Tjong Kim Sang and Fien De Meulder. 2003. Introduction to the CoNLL-2003 Shared Task: Language-Independent Named Entity Recognition. In *Proceedings of CoNLL-2003*.
- [42] Kristina Toutanova and Christopher D Manning. 2000. Enriching the knowledge sources used in a maximum entropy part-of-speech tagger. In *Proceedings of the 2000 Joint SIGDAT conference on Empirical methods in natural language processing and very large corpora: held in conjunction with the 38th Annual Meeting of the Association for Computational Linguistics-Volume 13*. Association for Computational Linguistics, 63–70.
- [43] Chen-Tse Tsai and Dan Roth. 2016. Cross-lingual wikification using multilingual embeddings. In *Proceedings of NAACL-HLT*. 589–598.
- [44] Ricardo Usbeck, Axel-Cyrille Ngonga Ngomo, Michael Röder, Daniel Gerber, Sandro Athaide Coelho, Sören Auer, and Andreas Both. 2014. AGDISTIS - Graph-Based Disambiguation of Named Entities Using Linked Data. In *The Semantic Web - ISWC 2014 - 13th International Semantic Web Conference, October 19-23, 2014, Proceedings, Part I*. Riva del Garda, Italy, 457–471.
- [45] Ricardo Usbeck, Michael Röder, Axel-Cyrille Ngonga Ngomo, Ciro Baron, Andreas Both, Martin Brümmer, Diego Ceccarelli, Marco Cornolti, Didier Chérix, Bernd Eickmann, Paolo Ferragina, Christiane Lemke, Andrea Moro, Roberto Navigli, Francesco Piccinno, Giuseppe Rizzo, Harald Sack, René Speck, Raphaël Troncy, Jörg Waitelonis, and Lars Wesemann. 2015. GERBIL: General Entity Annotator Benchmarking Framework. In *Proceedings of the 24th International Conference on World Wide Web, WWW 2015, Florence, Italy, May 18-22, 2015*. Florence, Italy, 1133–1143.
- [46] Marieke van Erp, Pablo Mendes, Heiko Paulheim, Filip Ilievski, Julien Plu, Giuseppe Rizzo, and Jörg Waitelonis. 2016. Evaluating entity linking: An analysis of current benchmark datasets and a roadmap for doing a better job. In *10th International Conference on Language Resources and Evaluation (LREC)*.
- [47] Jörg Waitelonis, Henrik Jürges, and Harald Sack. 2016. Don'T Compare Apples to Oranges: Extending GERBIL for a Fine Grained NEL Evaluation. In *Proceedings of the 12th International Conference on Semantic Systems (SEMANTICS 2016)*. ACM, New York, NY, USA, 65–72.
- [48] Han Wang, Guang Jin Zheng, Xiaogang Ma, Peter Fox, and Heng Ji. 2015. Language and Domain Independent Entity Linking with Quantified Collective Validation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Lisbon, Portugal, 695–704.
- [49] Lei Zhang and Achim Rettinger. 2014. X-LiSA: Cross-lingual Semantic Annotation. *PVLDB* 7, 13 (2014), 1693–1696.
- [50] Lei Zhang, Achim Rettinger, and Patrick Philipp. 2016. Context-Aware Entity Disambiguation in Text Using Markov Chains. In *IEEE/WIC/ACM International Conference on Web Intelligence (WI'16)*. IEEE Computer Society Press, Omaha, Nebraska, USA.
- [51] Stefan Zwicklbauer, Christin Seifert, and Michael Granitzer. 2016. DoSeR - A Knowledge-Base-Agnostic Framework for Entity Disambiguation Using Semantic Embeddings. In *The Semantic Web. Latest Advances and New Domains: 13th International Conference, ESWC 2016, Heraklion, Crete, Greece, May 29 - June 2, 2016, Proceedings*. Springer International Publishing, Cham, 182–198.