

VoxEL: A Benchmark Dataset for Multilingual Entity Linking

Henry Rosales-Méndez, Aidan Hogan and Barbara Poblete

IMFD Chile & Department of Computer Science, University of Chile
{hrosales, ahogan, bpoblete}@dcc.uchile.cl

Abstract. The Entity Linking (EL) task identifies entity mentions in a text corpus and associates them with corresponding entities in a given knowledge base. While traditional EL approaches have largely focused on English texts, current trends are towards language-agnostic or otherwise multilingual approaches that can perform EL over texts in many languages. One of the obstacles to ongoing research on multilingual EL is a scarcity of annotated datasets with the same text in different languages. In this work we thus propose VOXEL: a manually-annotated gold standard for multilingual EL featuring the same text expressed in five European languages. We first motivate and describe the VOXEL dataset, using it to compare the behaviour of state of the art EL (multilingual) systems for five different languages, contrasting these results with those obtained using machine translation to English. Overall, our results identify how five state-of-the-art multilingual EL systems compare for various languages, how the results of different languages compare, and further suggest that machine translation of input text to English is now a competitive alternative to dedicated multilingual EL configurations.

Keywords: Multilingual, Entity Linking, Information Extraction

Resource type: Dataset

Permanent URL: <https://dx.doi.org/10.6084/m9.figshare.6539675>

1 Introduction

The Entity Linking (EL) task identifies entity mentions in a text corpus and associates them with corresponding entities in a Knowledge Base (KB). In this way, we can leverage the information of publicly available KBs about real-world entities to achieve a better understanding of their semantics and also of natural language. For instance, in the text *“in the world of pop music, there is Michael Jackson and there is everybody else”* quoted from The New York Times, we can link the mention *Michael Jackson* with its corresponding entry in, e.g., the Wiki-data KB [35] (`wd:Q2831`), or the DBpedia KB [16] (`dbr:Michael_Jackson`)¹ allowing us to leverage, thereafter, the information in the KB about this entity to support semantic search, relationship extraction, text enrichment, entity summarisation, or semantic annotation, amongst other applications.

¹ Throughout, we use prefixes according to <http://prefix.cc>.

One of the major driving forces for research on EL has been the development of a variety of ever-expanding KBs that describe a broad selection of notable entities covering various domains (e.g., Wikipedia, DBpedia, Freebase, YAGO, Wikidata). Hence, while traditional Named Entity Recognition (NER) tools focused on identifying mentions of entities of specific types in a text, EL further requires disambiguation of which entity in the KB is being spoken about; this remains a challenging problem. On the one hand, name variations – such as “*Michael Joseph Jackson*”, “*Jackson*”, “*The King of Pop*” – mean that the same KB entity may be referred to in a variety of ways by a given text. On the other hand, ambiguity – where the name “*Michael Jackson*” may refer to various (other) KB entities, such as a journalist (wd:Q167877), a football player (wd:Q6831558), an actor (wd:Q6831554), amongst others – means that an entity mention in a text may have several KB candidates associated with it.

Many research works have addressed these challenges of the EL task down through the years. Most of the early EL systems proposed in the literature were monolingual approaches focusing on texts written in one single language, in most cases English (e.g., [18,12]). These approaches often use resources of a specific language, such as Part-Of-Speech taggers and WordNet², which prevent generalisation or adaptation to other languages. Furthermore, most of the labelled datasets available for training and evaluating EL approaches were English only (e.g., AIDA/CoNLL [12], DBpedia Spotlight Corpus [18], KORE 50 [13]).

However, as the EL area has matured, more and more works have begun to focus on languages other than English, including multilingual approaches that are either language agnostic [6,5,8,22] – relying only on the language of labels available in the reference KB – or that can be configured for multiple languages [21,30]. Recognising this trend, a number of multilingual datasets for EL were released, such as for the 2013 TAC KBP challenge³ and the 2015 SemEval Task 13 challenge⁴. Although such resources are valuable for multilingual EL research – where in previous work [28] we presented an evaluation of EL systems comparing two languages from the SemEval dataset – they have their limitations, key amongst which are their limited availability (participants only⁵), a narrow selection of languages, and differences in text and annotations across languages that makes it difficult to compare the performance in each language. More generally, the EL datasets available in multiple languages – and languages other than English – greatly lag behind what is available for English.

Contributions: In this paper, we propose the VOXEL dataset: a manually-annotated gold standard for EL considering five European languages, namely German, English, Spanish, French and Italian. This dataset is based on an online source of multilingual news, where we selected and annotated 15 corresponding news articles for these five languages (75 articles in total). Additionally, we

² <https://wordnet.princeton.edu>; April 1st, 2018

³ <https://tac.nist.gov/2013/KBP/>; April 1st, 2018

⁴ <http://alt.qcri.org/semeval2015/task13/>; April 1st, 2018

⁵ We have managed to acquire the SemEval dataset, but unfortunately we were not able to acquire the TAC–KBP dataset: our correspondence was not responded to.

created two versions of VOXEL: a *strict* version where entities correspond to a restricted definition of entity, as a mention of a person, place or organisation (based on traditional MUC/NER definitions), and a *relaxed* version where we considered a broader selection of mentions referring to entities described by Wikipedia. Based on the VOXEL dataset, using the GERBIL evaluation framework [34], we present results for various EL systems, allowing us to compare not only across systems, but also across languages. As an additional contribution, we compare the performance of EL systems configurable for a given language with the analogous results produced by applying state-of-the-art machine translation (Google translate) to English and then applying EL configured for English. Our findings show that most systems perform best for English text. Furthermore, machine translation of input text to English achieves comparable – and often better – performance when compared with dedicated multilingual EL approaches.

2 Preliminaries

We first introduce some preliminaries relating to EL. Let E be a set of entity identifiers in a KB; these are typically IRIs, such as `wd:Q2831`, `dbr:Michael_Jackson`. Given an input text, the EL process can be conceptualised in terms of two main phases. First, Entity Recognition (ER) establishes a set of entity mentions M , where each such mention is a sub-string referring to an entity, annotated with its start position in the input text, e.g., (37, “*Michael Jackson*”). Second, for each mention $m \in M$ recognised by the first phase, Entity Disambiguation (ED) attempts to establish a link between m and the corresponding identifier $e \in E$ for the KB entity to which it refers. The second disambiguation phase can be further broken down into a number of (typical) sub-tasks, described next:

Candidate entity generation: For each mention $m \in M$, this stage selects a subset of the most probable KB entities $E_m \subseteq E$ to which it may refer.

There are two high-level approaches by which candidate entities are often generated. The first is a dictionary-based approach, which involves applying keyword or string matching between the mention m and the label of entities from E . The second is an NER-based approach, where traditional NER tools are used to identify entity mentions (potentially) independently of the KB.

Candidate entity ranking: This stage is where the final disambiguation is made: the candidate entities E_m for each mention m are ranked according to some measure indicating their likelihood of being the reference for m . The measures used for ranking each entity $e \in E_m$ may take into account features of the candidate entity e (e.g., centrality), features of the candidate link (m, e) (e.g., string similarity), features involving e and candidates for neighbouring mentions E'_m (e.g., graph distance in the KB), and so forth. Ranking may take the form of an explicit metric that potentially combines several measures, or may be implicit in the use of machine-learning methods that classify candidates, or that compute an optimal assignment of links.

Unlinkable mention prediction: The target KBs considered by EL are often, by their nature, incomplete. In some applications, it may thus be useful

to extract entity mentions from the input text that do not (yet) have a corresponding entity in the KB. These are sometimes referred to as *emerging entities*, are typically produced by NER candidate generation (rather than a dictionary approach), and are assigned a label such as *NIL* (Not In Lexicon).

It is important to note that while the above processes provide a functional overview of the operation of most EL systems, not all EL systems follow this linear sequence of steps. Most systems perform recognition first, and once the mentions are identified the disambiguation phase is initiated [18,21]. However, other approaches may instead apply a unified process, building models that create feedback between the recognition and disambiguation steps [7]. In any case, the output of the EL process will be a set of links of the form (m, e) , where the mention m in the text is linked to the entity e in the KB, optionally annotated with a confidence score – often called a *support* – for the link.

3 Related Work

We now cover related works in the context of multilingual EL, first discussing approaches and systems, thereafter discussing available datasets.

3.1 Multilingual EL Systems

In theory, any EL system can be applied to any language; as demonstrated in our previous work [28], even a system supporting only English may still be able to correctly recognise and link the name of a person such as *Michael Jackson* in the text of another language, assuming the alphabet remains the same. Hence, the notion of a multilingual EL system can become blurred. For example language-agnostic systems – systems that require no linguistic components or resources specific to a language – can become multilingual simply by virtue of having a reference KB with labels in a different – or multiple different – language(s).

Here we thus focus on EL systems that have published evaluation results over texts from multiple languages⁶, thus demonstrating proven multilingual capabilities. We summarise such systems in Table 1, where we provide details on the year of the main publication, the languages evaluated, as well as denoting whether or not entity recognition is supported⁷, and whether or not a demo, source code or API is currently available.⁸ As expected, a high-level inspection of the table shows that English is the most popularly-evaluated (and thus we surmise supported) language, followed by European languages such as German, Spanish, French, Dutch and Italian. We also highlight that most of the multilingual EL approaches included in the table have emerged since 2010.

⁶ This excludes systems such as Apache Stanbol, OpenCalais, PoolParty, etc.

⁷ Some systems assume that mentions have previously been extracted from the text and are given as input, thereafter focusing only on the disambiguation process.

⁸ We presented an earlier version of such a table in previous work [28].

Table 1. Overview of multilingual EL approaches; the italicised approaches will be incorporated as part of our experiments.

Name	Year	Evaluated Languages	ER	Demo	Src	API
KIM [25]	2004	EN,FR,ES	✓	✓	✗	✓
<i>TagME</i> [8]	2010	DE,EN,NL	✓	✓	✗	✓
SDA [3]	2011	EN,FR	✓	✗	✗	✗
ualberta [10]	2012	EN,ZH	✓	✗	✗	✗
HITS [7]	2012	EN,ES,ZH	✓	✗	✗	✗
<i>THD</i> [6]	2012	DE,EN,NL	✓	✓	✓	✓
<i>DBpedia Spotlight</i> [18,5]	2013	DA,DE,EN,ES,FR,HU,IT,NL,RU	✓	✓	✓	✓
Wang-Tang [37]	2013	EN,ZH	✓	✗	✗	✗
AGDISTIS [33]	2014	DE,EN	✗	✓	✓	✓
<i>Babelfy</i> [21]	2014	DE,EN,ES,FR,IT	✓	✓	✗	✓
<i>FREME</i> [30]	2016	DE,EN	✓	✗	✓	✓
WikiME [32]	2016	AR,DE,EN,ES,FR,HE,IT,TA,TH,TL,TR,UR,ZH	✓	✓	✗	✗
FEL [23]	2017	EN,ES,ZH	✓	✗	✓	✗
FOX [31]	2017	DE,EN,ES,FR,NL	✓	✓	✓	✓
MAG [22]	2017	DE,EN,ES,FR,IT,JA,NL	✗	✓	✓	✓

We will later conduct experiments using the GERBIL evaluation framework [34], which allows for invoking and integrating the results of a variety of public APIs for EL, generating results according to standard metrics in a consistent manner. Hence, in our later experiments, we shall only consider those systems with a working REST-API made available by the authors of the system. In addition, we will manually label our VOXEL system according to Wikipedia, with which other important KBs such as DBpedia, YAGO, Freebase, Wikidata, etc., can be linked; hence we only include systems that support such a KB linked with Wikipedia. Note that GERBIL automatically takes care of mapping coreferent identifiers across KBs (and even across languages in cases such as DBpedia with different KB identifiers for different languages and cross-language links).

With these criteria in mind, we experiment with the following systems:

TagME (2010) uses analyses of anchor texts in Wikipedia pages to perform EL [8]. The ranking stage is based primarily on two measures: *commonness*, which describes how often an anchor text is associated with a particular Wikipedia entity; and *relatedness*, which is a co-citation measure indicating how frequently candidate entities for different mentions are linked from the same Wikipedia article. TagME is language agnostic: it can take advantage of the Wikipedia Search API to apply the same conceptual process over different language versions of Wikipedia to support multilingual EL.

THD (2012) is based on three measures [6]: *most frequent senses*, which ranks candidates for a mention based on the Wikipedia Search API results for that mention; *co-occurrence*, which is a co-citation measure looking at how often candidate entities for different mentions are linked from the same paragraphs

in Wikipedia; and *explicit semantic analysis*, which uses keyword similarity measures to relate mentions with a concept. These methods are language agnostic and applicable to different language versions of Wikipedia.

DBpedia Spotlight (2013) was first proposed to deal with English annotations [18], based on keyword and string matching functions ranked by a probabilistic model based on a variant of a TF-IDF measure. DBpedia Spotlight is largely language-agnostic, where an extended version later proposed by Daiber et al. [5] leverages the multilingual information of the Wikipedia and DBpedia KBs to support multiple languages.

Babelfy (2014) performs EL with respect to a custom multilingual KB BabelNet⁹ constructed from Wikipedia and WordNet, using machine translation to bridge the gaps in information available for different language versions of Wikipedia [21]. Recognition is based on POS tagging for different languages, selecting candidate entities by string matching. Ranking is reduced to finding the densest subgraph that relates neighbouring entities and mentions.

FRÉME (2016) delegates the recognition of entities to the Stanford-NER tool, which is trained over the anchor texts of Wikipedia corpora in different languages. Candidate entities are generated by keyword search over local indexes, which are then ranked based on the number of matching anchor texts in Wikipedia linking to the corresponding article of the candidate entity [30].

With respect to FOX, note that while it meets all of our criteria, at the time of writing, we did not succeed in getting the API to run over VOXEL without error; hence we do not include this system. We also omit AGDISTIS and MAG from our selection because they do not perform recognition, requiring a prior identification of the entities in the input text (finding a suitable NER tool/model is not straightforward for some of the languages in our dataset).

3.2 Multilingual EL Datasets

In order to train and evaluate EL approaches, labelled datasets – annotated with the correct entity mentions and their respective KB links – are essential. In some cases these datasets are labelled manually, while in other cases labels can be derived from existing information, such as anchor texts. In Table 2 we survey the labelled datasets most frequently used by EL approaches (note that sentence counts were not available for some datasets).

We can see that the majority of datasets provide text in one language only – predominantly English – with the exceptions being as follows:

SemEval 2015 Task 13: is built over a biomedical, math, computer and social domain and is designed to support EL and WSD at the same time, containing annotations to Wikipedia, BabelNet and WordNet [20].

DBpedia Abstracts: provides a large-scale training and evaluation corpora based on the anchor texts extracted from the abstracts (first paragraph) of Wikipedia pages in seven languages [2].¹⁰

⁹ <http://babelnet.org/>; April 1st, 2018

¹⁰ <http://wiki-link.nlp2rdf.org/abstracts/>; April 1st, 2018

Table 2. Survey of dataset for EL task. For multilingual datasets, the quantities shown refer to the English data available. We present metadata about the relaxed and strict version of our dataset by VOxEL_R and VOxEL_S respectively. (Abbreviations: $|D|$ number of documents, $|S|$ number of sentences, $|E|$ number of entities, **Mn** denotes that all entities were manually annotated.)

Dataset	$ D $	$ S $	$ E $	Mn	Languages
AIDA/CoNLL-Complete [12]	1393	22,137	34,929	✓	EN
KORE50 [13]	50	50	144	✓	EN
IITB [15]	103	1,781	18,308	✓	EN
ACE2004 [26]	57	-	306	✗	EN
AQUAINT [26]	50	533	727	✗	EN
MSNBC [4]	20	668	747	✗	EN
DBpedia Spotlight [18]	10	58	331	✓	EN
N3-RSS 500 [27]	1	500	1000	✓	EN
Reuters 128 [27]	128	-	881	✓	EN
Wes2015 [36]	331	-	28,586	✓	EN
News-100 [27]	100	-	1656	✓	DE
Thibaudet [1]	1	3,807	2,980	✗	FR
Bergson [1]	1	4,280	380	✗	FR
SemEval 2015 Task 13 [20]	4	137	769	✓	EN,ES,IT
DBpedia Abstracts [2]	39,132	-	505,033	✗	DE,EN,ES,FR,IT,JA,NL
MEANTIME [19]	120	597	2,790	✓	EN,ES,IT,NL
VoxEL_R	15	94	674	✓	DE,EN,ES,FR,IT
VoxEL_S	15	94	204	✓	DE,EN,ES,FR,IT

MEANTIME: consists of 120 news articles from WikiNews¹¹ with manual annotations of entities, events, temporal information and semantic roles [19].¹²

With respect to DBpedia Abstracts, while offering a very large multilingual corpus, the texts across different languages vary, as do the documents available; while such a dataset could be used to compare different systems for the same languages, it could not be used to compare the same systems for different languages. Furthermore, there are no guarantees for the completeness of the annotations since they are anchor texts/links extracted from Wikipedia; hence the dataset is best suited as a large collection of positive (training) examples, in a similar manner to how TagME [8] and FRENDE [30] use anchor texts.

Unlike DBpedia Abstracts, the SemEval and MEANTIME datasets contain analogous documents translated to different languages (also known as *paral-*

¹¹ <https://en.wikinews.org/>; April 1st, 2018

¹² <http://www.newsreader-project.eu/results/data/wikinews/>; April 1st, 2018

lel corpora [20]). Our VOXEL dataset complements these previous resources but with some added benefits. Primarily, both the SemEval and MEANTIME datasets exhibit slight variations in the annotations across languages, leading to (e.g.) a different number of entity annotations in the text for different languages; for example SemEval [20] reports 1,261 annotations for English, 1,239 for Spanish, and 1,225 for Italian, while MEANTIME [19] reports 2,790 entity mentions for English, 2,729 for Dutch, 2,709 for Italian and 2,704 for Spanish. On the other hand, VOXEL has precisely the same annotations across languages aligned at the sentence level, and also features datasets labelled under two definitions of entity. More generally, we see VOXEL as complementing these other datasets.¹³

4 The VoxEL Dataset

In this section, we describe the VOXEL Dataset that we propose as a gold standard for EL involving five languages: German, English, Spanish, French and Italian. VOXEL is based on 15 news articles sourced from the VoxEurop¹⁴ website: a European newsletter with the same news articles professionally translated to different languages. This source of text thus obviates the need for translation of texts to different languages, and facilitates the consistent identification and annotation of mentions (and their Wikipedia links) across languages. With VOXEL, we thus provide a high-quality resource with which to evaluate the behaviour of EL systems across a variety of European languages.

While the VoxEurop newsletter is a valuable source of professionally translated text in several European languages, there are sometimes natural variations across languages that – although they preserve meaning – may change how the entities are mentioned. A common example is the use of pronouns rather than repeating a person’s name to make the text more readable in a given language. Such variations would then lead to different entity annotations across languages, hindering comparability. Hence, in order to achieve the same number of sentences and annotations for each new (document), we applied small manual edits to homogenise the text (e.g., replacing a pronoun by a person’s name). On the other hand, sentences that introduce new entities in one particular language, or that deviate too significantly across all languages, are eliminated; fewer than 10% of the sentences from the original source were eliminated.

When labelling entities, we take into consideration the lack of consensus about what is an “*entity*” [14,17,29]: some works conservatively consider only mentions of entities referring to fixed types such as person, organisation and location as entities (similar to the traditional NER/TAC consensus on an entity), while other authors note that a much more diverse set of entities are available in Wikipedia and related KBs for linking, and thus consider any noun-phrase mentioning an entity in Wikipedia to be a valid target for linking [24]. Furthermore, there is a lack of consensus on how overlapping entities – like *New York*

¹³ In previous work we used the SemEval dataset to compare EL systems for English and Spanish texts, where we refer the reader to [28] for more details.

¹⁴ <http://www.voxeurop.eu/>; April 1st, 2018

City Fire Department – should be treated [14,17]; should *New York City* be annotated as a separate entity or should we only cover maximal entities? Rather than take a stance on such questions – which appear application dependant – we instead create two versions of the data: a *strict* version that considers only maximal entity mentions referring to persons, organisations and locations; and a *relaxed* version that considers any noun phrase mentioning a Wikipedia entity as a mention, including overlapping mentions where applicable. For example, in the sentence “*The European Central Bank released new inflation figures today*” the strict version would only include “*European Central Bank*”, while the relaxed version would also include “*Central Bank*” and “*inflation*”.

To create the annotation of mentions with corresponding KB identifiers, we implemented a Web tool¹⁵ that allows a user to annotate a text, producing output in the NLP Interchange Format (NIF) [11], as well as offering visualisations of the annotations that facilitate, e.g., revision. For each language, we provide annotated links targeting the English Wikipedia entry, as well as that language’s version of Wikipedia (if different from English). In case there was no appropriate Wikipedia entry for a mention of a person, organisation or place, we annotate the mention with a `NotInLexicon` marker. These annotations were created by the first author in English, which were then revised by the other authors according to the two labelling guidelines (*strict* and *relaxed*). The first author then extended these annotations to the other languages using the sentence-level correspondence, thereafter verifying that each language has the same number of annotations and the same set of English Wikipedia identifiers for each sentence.

In summary, VOXEL consists of 15 news articles (documents) from the multilingual newsletter VoxEurop, totalling 94 sentences; the central topic of these documents is politics, particularly at a European level. This text is annotated five times for each language, and two times for the strict and relaxed versions, giving a total of 150 annotated documents and 940 sentences. The same number of annotations is given for each language (including by sentence). For the strict version, each language has 204 annotated mentions, while for the relaxed version, each language has 674 annotated mentions. In the relaxed version, 6.2%, 10.8%, 20.3% and 62.7% of the entries correspond to *persons*, *organisations*, *places* and *others* respectively, while in the strict version the entities that fall in the first three classes constitute 16.9%, 28.7% and 54.4% (*others* are excluded by definition under the strict guidelines). Again, this homogeneity of text and annotations across languages was non-trivial to achieve, but facilitates comparison of evaluation results not only across systems, but across languages.

5 Experiments

We now use our proposed VOXEL dataset to conduct experiments in order to explore the behaviour of state-of-the-art EL systems for multilingual settings. In particular, we are interested in the following questions:

¹⁵ <https://github.com/henryrosalesmendez/NIFify>; April 1st, 2018

Table 3. GERBIL Evaluation of EL systems with Micro Recall (mR), Precision (mP) and F_1 (mF). A value “–” indicates that the system does not support the corresponding language. The results in bold are the best for that metric, system and dataset variant comparing across the five languages (i.e., the best in each row, split by Relax/Strict).

		Relaxed					Strict				
		DE	EN	ES	FR	IT	DE	EN	ES	FR	IT
Babelfy _R	mP	0.840	0.649	0.835	0.824	0.810	0.932	0.785	0.929	0.889	0.907
	mR	0.461	0.522	0.549	0.488	0.451	0.676	0.735	0.710	0.632	0.578
	mF	0.595	0.578	0.662	0.613	0.579	0.784	0.759	0.805	0.739	0.706
Babelfy _S	mP	0.903	0.722	0.916	0.912	0.884	0.942	0.816	0.923	0.912	0.894
	mR	0.181	0.219	0.210	0.200	0.192	0.558	0.524	0.593	0.563	0.583
	mF	0.301	0.336	0.342	0.328	0.316	0.701	0.638	0.722	0.697	0.706
DBspot	mP	0.731	0.745	0.691	0.658	0.682	0.781	0.854	0.690	0.691	0.800
	mR	0.508	0.577	0.399	0.360	0.488	0.544	0.602	0.382	0.406	0.549
	mF	0.600	0.650	0.506	0.466	0.569	0.641	0.706	0.492	0.512	0.651
FREME	mP	0.762	0.803	0.655	0.737	0.857	0.750	0.871	0.660	0.739	0.858
	mR	0.161	0.267	0.175	0.129	0.213	0.426	0.764	0.553	0.416	0.652
	mF	0.266	0.400	0.276	0.219	0.342	0.543	0.814	0.602	0.532	0.740
TagME	mP	0.635	0.754	–	–	0.494	0.875	0.946	–	–	0.742
	mR	0.232	0.488	–	–	0.182	0.652	0.784	–	–	0.509
	mF	0.340	0.592	–	–	0.266	0.747	0.857	–	–	0.604
THD	mP	0.831	0.806	–	–	–	0.857	0.809	–	–	–
	mR	0.109	0.253	–	–	–	0.352	0.647	–	–	–
	mF	0.194	0.386	–	–	–	0.500	0.719	–	–	–

- **RQ1:** How does the performance of systems compare for multilingual EL?
- **RQ2:** For which of the five languages are the best results achieved?
- **RQ3:** How would a method based on machine translation to English compare with directly configuring the system for a particular language?

In order to address **RQ1** and **RQ2**, we ran the multilingual EL systems Babelfy, DBpedia Spotlight, FREME, TagME and THD over both versions of VOXEL in all five languages. These experiments were conducted with the GERBIL [34] EL evaluation framework, which provides unified access to the public APIs of multiple EL tools, abstracting different input and output formats using the NIF vocabulary, translating identifiers across KBs, and allowing to apply standard metrics to measure the performance of results with respect to a labelled dataset. GERBIL calls these systems via their REST APIs maintaining default (non-language) parameters, except for the case of Babelfy, for which we analyse two configurations: one that applies a more liberal interpretation of entities to include conceptual entities (Babelfy_R), and another configuration that applies a stricter definition of entities (Babelfy_S), where the two configurations correspond loosely with the relaxed/strict versions of our dataset.

The results of these experiments are shown in Table 3, where we present micro-measures for Precision (mP), Recall (mR) and F_1 (mF), with all systems,

for all languages, in both versions of the dataset.¹⁶ From first impressions, we can observe that two systems – TagME and THD – cannot be configured for all languages, where we leave the corresponding results blank.

With respect to **RQ1**, for the Relaxed version, the highest F_1 scores are obtained by Babelfy_R (0.662: ES) and DBpedia Spotlight (0.650: EN). On the other hand, the highest F_1 scores for the Strict version are TagME (0.857: EN) and Babelfy_R (0.805: ES). In general, the F_1 scores for the Strict version were higher than those for the Relaxed version: investigating further, the GERBIL framework only considers annotations to be false positives when a different annotation is given in the labelled dataset at an overlapping position; hence fewer labels in the Strict dataset will imply fewer false positives overall, which seems to outweigh the effect of the extra true positives that the Relaxed version would generate. Comparing the best Strict/Relaxed results for each system, we can see that Babelfy_R, DBpedia Spotlight and FREDM have less of a gap between both, meaning that they tend to annotate a broader range of entities; on the other hand, Babelfy_S and THD are more restrictive in the entities they link.

With respect to **RQ2**, considering all systems, we can see a general trend that English had the best results overall, with the best mF for DBpedia Spotlight, FREDM and TagME. For THD, German had higher precision but much lower recall; a similar result can be seen for FREDM in Italian in the Relaxed version. On the other hand, Babelfy generally had best results in German and Spanish, where, in fact, it often had the *lowest* precision in English.

With respect to possible factors that explain such differences across languages, there are variations between languages that may make the EL task easier or harder depending on the features used; for example, systems that rely on capitalisation may perform differently for Spanish, which uses less capitalisation, (e.g., “*Jungla de cristal*”: a Spanish movie title in sentence case); and German, where all nouns are capitalised. Furthermore, the quality of EL resources available for different languages – in terms of linguistic components, training sets, contextual corpora, KB meta-data, etc. – may also vary across languages.

Regarding **RQ3**, we present another experiment to address the question of the efficacy of using machine translations. First we note that, although works in related areas – such as cross-lingual ontology matching [9] – have used machine translation to adapt to multilingual settings, to the best of our knowledge, no system listed in Table 1 uses machine translations over the input text (though systems such as Babelfy do use machine translations to enrich the lexical knowledge available in the KB). Hence we check to see if translating a text to English using a state-of-the-art approach – Google Translate¹⁷ – and applying EL over the translated English text would fare better than applying EL directly over the target language; we choose one target language to avoid generating results for a quadratic pairing of languages, and we choose English since it was the only language working for/supported by all systems in Table 3.

¹⁶ The GERBIL results are available at https://users.dcc.uchile.cl/~hrosales/ISWC2018_experiment_GERBIL.html

¹⁷ <https://translate.google.com/>; April 1st, 2018

A complication for these translation experiments is that while VOXEL contains annotations for the texts in their original five languages, including English, it does not contain annotations for the texts translated to English. While we considered manually annotating such documents produced by Google Translate, we opted against it partly due to the amount of labour it would again involve, but more importantly because it would be specific to one translation service at one point in time: as these translation services improve, these labelled documents would quickly become obsolete. Instead, we apply evaluation on a per-sentence basis, where for each sentence of a text in a non-English language, we translate it and then compare the set of annotations produced against the set of manually-annotated labels from the original English documents; in other words, we check the annotations produced by sentence, rather than by their exact position. This is only possible because in the original VOXEL dataset, we defined a one-to-one correspondence between sentences across the five different languages.

Note that since GERBIL requires labels to have a corresponding position, we thus needed to run these experiments locally outside of the GERBIL framework. Hence, for a sentence s , let A denote the IRIs associated with manual labels for s in the original English text, and let B denote the IRIs annotated by the system for the corresponding sentence of the translated text; we denote true positives by $A \cap B$, false positives by $B - A$, and false negatives by $A - B$.¹⁸

In Table 4, we show the results of this second experiment, focusing this time on the Micro- F_1 (mF) score obtained for each system over the five languages of VOXEL, again for the relaxed and strict versions. For each system, we consider three experiments: (1) the system is configured for the given language and run over text for the given language, (2) the system is configured for English and run over the text translated from the given language, (3) the system is configured for English and run over the text in the given language without translation. We use the third experiment to establish how the translation to English – rather than the system configuration to English – affects the results. First we note that without using positional information to check false positives (as per GERBIL), the results change from those presented in Table 3; more generally, the gap between the Relaxed and Strict version is reduced.

With respect to **RQ3**, in Table 4, for each system, language and dataset variant, we underline which of the three configurations performs best. For example, in DBpedia Spotlight, all values on the (EN,EN_{*t*}) line – which denotes applying DBpedia Spotlight configured for English over text translated to English – are underlined, meaning that for all languages, prior translation to English outperformed submitting the text in its original language to DBpedia Spotlight configured for that language.¹⁹ In fact, for almost all systems, translating the input text to English generally outperforms using the available language configurations of the respective EL systems, with the exception of Babelify, where the available multilingual settings generally outperform a prior translation to

¹⁸ To compute Precision, Recall and F_1 , we do not require true negatives.

¹⁹ ... it also implies that it outperforms running English EL on text in the original language, though this is hardly surprising and just presented for reference.

Table 4. Micro F_1 scores for systems performing EL with respect to the VoxEL dataset. For each system and each non-English language, we show the results of three experiments: first, for $(_,_)$ the system is configured for the same language as the input text; second, for (EN,EN_t) , the system is configured for English and applied to text translated to English from the original language (EN,EN) ; third, for $(\text{EN},_)$, the system is configured for English and run for the text in the current (original) language. Below the name of each system, we provide the relaxed and strict results for the English text. Underlined results indicate the best of the three configurations for the given system, language and dataset variant (e.g., the best for the columns of three values). The best result for each system across all variations (excluding English input) is bolded.

		Relaxed				Strict			
		DE	ES	FR	IT	DE	ES	FR	IT
Babelfy _R (0.545,0.319)	$(_,_)$	<u>0.523</u>	0.541	0.493	<u>0.504</u>	0.344	0.362	0.309	0.365
	(EN,EN_t)	<u>0.507</u>	<u>0.515</u>	<u>0.505</u>	<u>0.501</u>	<u>0.298</u>	<u>0.298</u>	<u>0.314</u>	<u>0.301</u>
	$(\text{EN},_)$	0.215	0.170	<u>0.195</u>	0.140	0.253	0.239	<u>0.220</u>	0.179
Babelfy _S (0.308,0.567)	$(_,_)$	0.279	<u>0.325</u>	0.290	<u>0.311</u>	<u>0.572</u>	<u>0.611</u>	<u>0.583</u>	0.616
	(EN,EN_t)	<u>0.311</u>	0.309	<u>0.322</u>	0.303	0.518	0.523	0.559	0.532
	$(\text{EN},_)$	<u>0.201</u>	0.179	<u>0.189</u>	0.137	0.376	0.372	0.395	0.258
DBpedia Spotlight (0.466,0.707)	$(_,_)$	0.400	0.331	0.240	0.342	0.510	0.477	0.481	0.653
	(EN,EN_t)	<u>0.441</u>	<u>0.454</u>	<u>0.464</u>	<u>0.449</u>	<u>0.696</u>	<u>0.694</u>	<u>0.721</u>	0.729
	$(\text{EN},_)$	0.209	<u>0.161</u>	<u>0.180</u>	<u>0.188</u>	0.374	0.259	0.326	0.323
FREME (0.407,0.708)	$(_,_)$	0.282	0.302	0.268	0.373	0.483	0.583	0.479	0.726
	(EN,EN_t)	<u>0.404</u>	<u>0.403</u>	<u>0.401</u>	<u>0.408</u>	<u>0.701</u>	<u>0.713</u>	<u>0.692</u>	<u>0.711</u>
	$(\text{EN},_)$	0.166	0.183	0.196	0.222	0.190	0.338	0.342	0.374
TagME (0.462,0.327)	$(_,_)$	0.414	–	–	–	0.272	–	–	–
	(EN,EN_t)	<u>0.431</u>	0.450	0.441	0.439	0.330	0.333	0.321	0.336
	$(\text{EN},_)$	0.188	0.181	0.200	0.148	0.212	0.202	0.197	0.164
THD (0.392,0.625)	$(_,_)$	0.241	–	–	–	0.546	–	–	–
	(EN,EN_t)	<u>0.394</u>	<u>0.392</u>	<u>0.386</u>	<u>0.387</u>	<u>0.597</u>	<u>0.620</u>	<u>0.595</u>	0.623
	$(\text{EN},_)$	0.207	0.175	0.217	0.174	0.251	0.332	0.403	0.352

English (we may recall that in Table 3, Babelfy performed best for texts other than English). We further note that the translation results are generally competitive with those for the original English text – shown below the name of the system for the Relaxed and Strict datasets – even slightly outperforming those results in some cases. We also observe from the generally poor $(\text{EN},_)$ results that translation is important; in other words, one cannot simply just apply an EL system configured for English over another language and expect good results.

To give a better impression of the results obtained from the second experiment, in Figure 1, for the selected systems, we show the following aggregations: (1) *Calibrated* $(_,_)$: the mean Micro- F_1 score across the four non-English languages with the EL system configured for that language; (2) *Translation* (EN,EN_t) : the mean Micro- F_1 score across the four non-English languages with the text translated to English and the EL system configured for English; (3) English (EN,EN) : the (single) Micro- F_1 score for the original English text. From




Fig. 1. Summary of the Micro-F₁ results over VOXEL Relaxed/Strict for the translation experiments, comparing mean values for setting the EL system to the language of the text (*Calibrated*), translating the text to English first (*Translation*), and the corresponding F₁ score for EL over the original English text (*English*)

this figure, we can see that translation is comparable to native English EL, and that translation often considerably outperforms EL in the original language.

We highlight that using translation to English, the result will be an annotated text in English rather than the original language. However, given that translation is done per-sentence, the EL annotations for the translated English text could potentially be “mapped” back per sentence to the text in the original language; at the very least, the translated English annotations would be a useful reference.

6 Conclusion

While Entity Linking has traditionally focused on processing texts in English, in recent years there has been a growing trend towards developing techniques and systems that can support multiple languages. To support such research, in this paper we have described a new labelled dataset for multilingual EL, which we call VOXEL. The dataset contains 15 new articles in 5 different languages with 2 different criteria for labelling, resulting in a corpus of 150 manually-annotated news articles. In a Strict version of the dataset considering a core of entities, we derive 204 annotated mentions in each language, while in a Relaxed version of the dataset considering a broader range of entities described by Wikipedia, we derive 674 annotated mentions in each language. The VOXEL dataset is distinguished by having a one-to-one correspondence of sentences – and annotated entities per sentence – between languages. The dataset (in NIF) is available online under a CC-BY 4.0 licence: <https://dx.doi.org/10.6084/m9.figshare.6539675>.

We used the VOXEL dataset to conduct experiments comparing the performance of selected EL systems in a multilingual setting. We found that in general, Babelfy and DBpedia Spotlight performed the most consistently across

languages. We also found that with the exception of Babelfy, EL systems performed best over English versions of the text. Next, we compared configuring the multilingual EL system for each non-English language versus applying a machine translation of the text to English and running the system in English; with the exception of Babelfy, we found that the machine translation approach outperformed configuring the system for a non-English language; even in the case of Babelfy, the translation sometimes performed better, while in others it remained competitive. This raises a key issue for research on multilingual EL: state-of-the-art machine translation is now reaching a point where we must ask if it is worth building dedicated multilingual EL systems, or if we should focus on EL for one language to which other languages can be machine translated.

Acknowledgements The work of Henry Rosales-Méndez was supported by CONICYT-PCHA/Doctorado Nacional/2016-21160017. The work was also supported by the Millennium Institute for Foundational Research on Data (IMFD) and by Fondecyt Grant No. 1181896. We also thank Michael Röder for his considerable help with GERBIL.

References

1. Brando, C., Frontini, F., Ganascia, J.G.: REDEN: Named Entity Linking in Digital Literary editions using Linked Data sets. *CSIMQ* (7), 60–80 (2016)
2. Brümmer, M., Dojchinovski, M., Hellmann, S.: DBpedia Abstracts: A Large-Scale, Open, Multilingual NLP Training Corpus. In: *LREC* (2016)
3. Charton, E., Gagnon, M., Ozell, B.: Automatic semantic web annotation of named entities. In: *Canadian Conf. on A.I.* pp. 74–85. Springer (2011)
4. Cucerzan, S.: Large-Scale Named Entity Disambiguation Based on Wikipedia Data. *EMNLP-CoNLL* p. 708 (2007)
5. Daiber, J., Jakob, M., Hokamp, C., Mendes, P.N.: Improving efficiency and accuracy in multilingual entity extraction. In: *I-SEMANTICS*. pp. 121–124 (2013)
6. Dojchinovski, M., Kliegr, T.: Recognizing, classifying and linking entities with Wikipedia and DBpedia. *WIKT* pp. 41–44 (2012)
7. Fahrni, A., Göckel, T., Strube, M.: HITS’ Monolingual and Cross-lingual Entity Linking System at TAC 2012: A Joint Approach. In: *TAC* (2012)
8. Ferragina, P., Scaiella, U.: Tagme: on-the-fly annotation of short text fragments (by wikipedia entities). In: *CIKM*. pp. 1625–1628. ACM (2010)
9. Fu, B., Brennan, R., O’Sullivan, D.: Cross-Lingual Ontology Mapping and Its Use on the Multilingual Semantic Web. In: *Multilingual Semantic Web*. pp. 13–20 (2010)
10. Guo, Z., Xu, Y., de Sá Mesquita, F., Barbosa, D., Kondrak, G.: ualberta at TAC-KBP 2012: English and Cross-Lingual Entity Linking. In: *TAC* (2012)
11. Hellmann, S., Lehmann, J., Auer, S., Brümmer, M.: Integrating NLP Using Linked Data. In: *ISWC*. pp. 98–113 (2013)
12. Hoffart, J., et al.: Robust disambiguation of named entities in text. In: *EMNLP*. pp. 782–792. *ACL* (2011)
13. Hoffart, J., Seufert, S., Nguyen, D.B., Theobald, M., Weikum, G.: KORE: keyphrase overlap relatedness for entity disambiguation. In: *CIKM*. pp. 545–554 (2012)

14. Jha, K., Röder, M., Ngomo, A.C.N.: All that glitters is not gold—rule-based curation of reference datasets for named entity recognition and entity linking. In: *ESWC*. pp. 305–320. Springer (2017)
15. Kulkarni, S., Singh, A., Ramakrishnan, G., Chakrabarti, S.: Collective annotation of Wikipedia entities in web text. In: *SIGKDD*. pp. 457–466 (2009)
16. Lehmann, J., et al.: DBpedia—a large-scale, multilingual knowledge base extracted from Wikipedia. *Semantic Web* 6(2), 167–195 (2015)
17. Ling, X., Singh, S., Weld, D.S.: Design challenges for entity linking. *TACL* 3, 315–328 (2015)
18. Mendes, P.N., Jakob, M., García-Silva, A., Bizer, C.: DBpedia spotlight: shedding light on the web of documents. In: *I-SEMANTICS*. pp. 1–8. ACM (2011)
19. Minard, A., Speranza, M., Urizar, R., Altuna, B., van Erp, M., Schoen, A., van Son, C.: MEANTIME, the NewsReader Multilingual Event and Time Corpus. In: *LREC* (2016)
20. Moro, A., Navigli, R.: SemEval-2015 Task 13: Multilingual All-Words Sense Disambiguation and Entity Linking. In: *SemEval@ NAACL-HLT*. pp. 288–297 (2015)
21. Moro, A., Raganato, A., Navigli, R.: Entity linking meets word sense disambiguation: a unified approach. *Trans. of the ACL* 2, 231–244 (2014)
22. Moussallem, D., et al.: MAG: A Multilingual, Knowledge-base Agnostic and Deterministic Entity Linking Approach. In: *K-CAP*. p. 9 (2017)
23. Pappu, A., Blanco, R., Mehdad, Y., Stent, A., Thadani, K.: Lightweight Multilingual Entity Extraction and Linking. In: *WSDM*. pp. 365–374. ACM (2017)
24. Perera, S., Mendes, P.N., Alex, A., Sheth, A.P., Thirunarayan, K.: Implicit Entity Linking in Tweets. In: *ESWC*. pp. 118–132 (2016)
25. Popov, B., et al.: Kim—a semantic platform for information extraction and retrieval. *Natural Language Engineering* 10(3-4), 375–392 (2004)
26. Ratnov, L., Roth, D., Downey, D., Anderson, M.: Local and global algorithms for disambiguation to Wikipedia. In: *NAACL-HLT*. pp. 1375–1384 (2011)
27. Röder, M., et al.: N³-A Collection of Datasets for Named Entity Recognition and Disambiguation in the NLP Interchange Format. In: *LREC*. pp. 3529–3533 (2014)
28. Rosales-Méndez, H., Poblete, B., Hogan, A.: Multilingual Entity Linking: Comparing English and Spanish. In: *LD4IE@ISWC*. pp. 62–73 (2017)
29. Rosales-Méndez, H., Poblete, B., Hogan, A.: What Should Entity Linking link? In: *AMW* (2018)
30. Sasaki, F., Dojchinovski, M., Nehring, J.: Chainable and Extendable Knowledge Integration Web Services. In: *ISWC*. pp. 89–101 (2016)
31. Speck, R., et al.: Ensemble Learning of Named Entity Recognition Algorithms using Multilayer Perceptron for the Multilingual Web of Data. In: *K-CAP*. p. 26 (2017)
32. Tsai, C.T., Roth, D.: Cross-lingual wikification using multilingual embeddings. In: *NAACL-HLT*. pp. 589–598 (2016)
33. Usbeck, R., et al.: AGDISTIS-graph-based disambiguation of named entities using Linked Data. In: *ISWC*. pp. 457–471. Springer (2014)
34. Usbeck, R., et al.: GERBIL: General Entity Annotator Benchmarking Framework. In: *WWW*. pp. 1133–1143 (2015)
35. Vrandečić, D., Krötzsch, M.: Wikidata: a free collaborative knowledgebase. *Communications of the ACM* 57(10), 78–85 (2014)
36. Waitelonis, J., Exeler, C., Sack, H.: Linked Data enabled generalized vector space model to improve document retrieval. In: *NLP & DBpedia @ ISWC* (2015)
37. Wang, Z., Li, J., Tang, J.: Boosting Cross-Lingual Knowledge Linking via Concept Annotation. In: *IJCAI*. pp. 2733–2739 (2013)