

DAWT: Densely Annotated Wikipedia Texts Across Multiple Languages

Nemanja Spasojevic, Preeti Bhargava, Guoning Hu
Lithium Technologies | Klout
San Francisco, CA
{nemanja.spasojevic, preeti.bhargava, guoning.hu}@lithium.com

ABSTRACT

In this work, we open up the DAWT dataset - Densely Annotated Wikipedia Texts across multiple languages. The annotations include labeled text mentions mapping to entities (represented by their Freebase machine ids) as well as the type of the entity. The data set contains total of 13.6M articles, 5.0B tokens, 13.8M mention entity co-occurrences. DAWT contains 4.8 times more anchor text to entity links than originally present in the Wikipedia markup. Moreover, it spans several languages including English, Spanish, Italian, German, French and Arabic. We also present the methodology used to generate the dataset which enriches Wikipedia markup in order to increase number of links. In addition to the main dataset, we open up several derived datasets including mention entity co-occurrence counts and entity embeddings, as well as mappings between Freebase ids and Wikidata item ids. We also discuss two applications of these datasets and hope that opening them up would prove useful for the Natural Language Processing and Information Retrieval communities, as well as facilitate multi-lingual research.

Keywords

Wiki, Wikipedia, Freebase, Freebase annotations, Wikipedia annotations, Wikification, Named Entity Recognition, Entity Disambiguation, Entity Linking

1. INTRODUCTION

Over the past decade, the amount of data available to enterprises has grown exponentially. However, a majority of this data is unstructured or free-form text, also known as Dark Data¹. This data holds challenges for Natural Language Processing (NLP) and information retrieval (IR) tasks unless the text is semantically labeled. Two NLP tasks that are particularly important to the IR community are:

¹https://en.wikipedia.org/wiki/Dark_data

©2017 International World Wide Web Conference Committee (IW3C2), published under Creative Commons CC BY 4.0 License. WWW'17 Companion, April 3–7, 2017, Perth, Australia. ACM 978-1-4503-4914-7/17/04. <http://dx.doi.org/10.1145/3041021.3053367>



- Named Entity Recognition (NER) - task of identifying an *entity mention* within a text,
- Entity Disambiguation and Linking (EDL) - task of linking the mention to its correct entity in a Knowledge Base (KB).

These tasks play a critical role in the construction of a high quality information network which can be further leveraged for a variety of IR and NLP tasks such as text categorization, topical interest and expertise modeling of users [20, 21]. Moreover, when any new piece of information is extracted from text, it is necessary to know which real world entity this piece refers to. If the system makes an error here, it loses this piece of information and introduces noise. As a result, both these tasks require high quality labeled datasets with densely extracted mentions linking to their correct entities in a KB.

Wikipedia has emerged as the most complete and widely used KB over the last decade. As of today, it has around 5.3 million English articles and 38 million articles across all languages. In addition, due to its open nature and availability, Wikipedia Data Dumps have been adopted by academia and industry as an extremely valuable data asset. Wikipedia precedes other OpenData projects like Freebase [8] and DBpedia [10] which were built on the foundation of Wikipedia. The Freebase Knowledge Graph is the most exhaustive knowledge graph capturing 58 million entities and 3.17 billion facts. The Wikipedia and Freebase data sets are large in terms of:

- information comprehensiveness,
- wide language coverage,
- number of cross-article links, manually curated cross entity relations, and language independent entity identifiers.

Although these two data sets are readily available, Wikipedia link coverage is relatively sparse as only the first *entity mention* is linked to the entity's Wikipedia article. This sparsity may significantly reduce the number of training samples one may derive from Wikipedia articles which, in turn, reduces the utility of the dataset. In this work, we primarily focus on creating the DAWT dataset that contains denser annotations across Wikipedia articles. We leverage Wikipedia and Freebase to build a large data set of annotated text where entities extracted from Wikipedia text are mapped to Freebase ids. Moreover, this data set spans multiple languages. In addition to the main dataset, we open up several derived datasets for mention occurrence counts, entity occurrence counts, mention entity co-occurrence counts and entity Word2Vec. We also discuss two applications of these datasets and hope that opening them up would prove use-

ful for the NLP and IR communities as well as facilitate multi-lingual research.

2. PROBLEM STATEMENT

The Wikification problem was introduced by Mihalcea et al. [13], where task was to introduce hyperlinks to the correct wikipedia articles for a given mention. In Wikipedia, only the first mention is linked or annotated. In our task, we focus on densifying the annotations i.e. denser hyperlinks from mentions in Wikipedia articles to other Wikipedia articles. The ultimate goal is to have high-precision hyperlinks with relatively high recall that could be further used as ground truth for other NLP tasks.

For most of the supervised Machine Learning or NLP tasks, one of the challenges is gathering ground truth at scale. In this work, we try to solve the problem of generating a labeled data set at large scale with the following constraints:

- The linked entity ids need to be unified across different languages. In Freebase, the machine id is same across different languages and hence, we annotate Wikipedia with Freebase machine ids,
- The dataset needs to be comprehensive (with large number of entities spanning multiple domains),
- The labels should be precise.

3. CONTRIBUTIONS

Our contributions in this work are:

- We extract a comprehensive inventory of mentions spanning several domains.
- We densify the entity links in the Wikipedia documents by 4.8 times.
- The DAWT dataset covers several more languages in addition to English such as Arabic, French, German, Italian, and Spanish.
- Finally, we open up this dataset and several other derived datasets (such as mention occurrence counts, entity occurrence counts, mention entity co-occurrence counts, entity word2vec and mappings between Freebase ids and Wikidata item ids) for the benefit of the IR and NLP communities.

4. KNOWLEDGE BASE

Our KB consists of about 1 million Freebase² machine ids for entities. These were chosen from a subset of all Freebase entities that map to Wikipedia entities. We prefer to use Freebase as our KB since in Freebase, the same id represents a unique entity across multiple languages. For a more general use, we have also provided the mapping from Freebase id to Wikipedia link and Wikidata item id (see Section 6.4). Due to limited resources and usefulness of the entities, our KB contains approximately 1 million most important entities from among all the Freebase entities. This gives us a good balance between coverage and relevance of entities for processing common social media text. To this end, we calculate an entity importance score [2] using linear regression with features capturing popularity within Wikipedia links, and importance of the entity within Freebase. We used signals such as Wiki page rank, Wiki and Freebase incoming

²Freebase was a standard community generated KB until June 2015 when Google deprecated it in favor of the commercially available Knowledge Graph API.

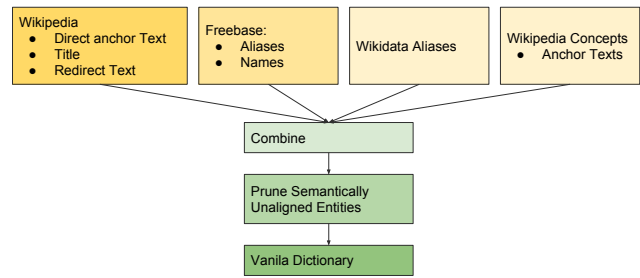


Figure 1: Candidate Dictionary Generation Overview

and outgoing links, and type descriptors within our KB etc. We use this score to rank the entities and retain only the top 1 million entities in our KB.

In addition to the KB entities, we also employ two special entities: **NIL** and **MISC**. **NIL** entity indicates that there is no entity associated with the mention, eg. mention ‘the’ within the sentence may link to entity **NIL**. This entity is useful especially when dealing with stop words and false positives. **MISC** indicates that the mention links to an entity which is outside the selected entity set in our KB.

5. DAWT DATA SET GENERATION

Our main goals when building the DAWT data set were to maintain high precision and increase linking coverage. As shown in Figure 1, we first generate a list of candidate phrases mapping to the Wikipedia articles by combining:

- Wiki articles (direct anchor texts, titles of pages, redirect text to wiki pages)
- Freebase Aliases, and Freebase Also Known As fields related to entities
- Wikipedia Concepts (English anchor texts)

The initial candidate lists are pruned to remove outlier phrases that do not semantically align with the rest of the list. As a semantic alignment metrics of two phrases, we used a combination of Jaccard similarity (both token and 3-gram character), edit distance (token and character), and largest common subsequence. We averaged the metrics and for each candidate in the list, we calculated average alignment against all the other candidates. As final step we remove all candidates, if any, with the lowest alignment scores.

For example, in the candidate set {‘USA’, ‘US’, ‘our’, ...} for entity *USA*, phrase ‘our’ does not align with rest of the cluster and is filtered out. In addition to the candidate dictionary, we also calculate co-occurrence frequencies, based on direct links from Wikipedia article markup, between any 2 entities appearing within the same sentence.

To generate the DAWT dataset, we do the following. For each supported language, and for each Wiki page in the language:

1. Iterate over the Wiki article and extract the set of directly linked entities.
2. Calculate all probable co-occurring entities with the set of directly linked entities from Step 1.
3. Iterate over the Wiki article and map all phrases to their set of candidate entities.
4. Resolve phrases whose candidates have been directly linked from Step 1.
5. For the remaining unresolved references, choose candidates, if any, with the highest probable co-occurrence with the directly linked entities.

Language	Article	Unique Mention Count	Unique Entity Count	Unique Mention Entity pairs	Total Mention Entity links	Total CPU time (days)
en	5,303,722	5,786,727	1,276,917	6,956,439	360,323,792	139.5
es	2,393,366	901,370	224,695	1,038,284	62,373,952	17.9
it	1,467,486	799,988	211,687	931,369	47,659,715	14.2
fr	1,750,536	1,670,491	423,603	1,952,818	93,790,881	28.6
de	1,818,649	2,168,723	426,556	2,438,583	103,738,278	20.2
ar	889,007	394,024	186,787	433,472	12,387,715	1.6

Table 1: DAWT Data Set Statistics

As a last step, the hyperlinks to Wikipedia articles in a specific language are replaced with links to their Freebase ids to adapt to our KB.

The densely annotated Wikipedia articles have on an average 4.8 times more links than the original articles. A detailed description of the data set, per language, along with total CPU time taken for annotation is shown in Table 1. All experiments were run on a 8-core 2.4GHz Xeon processor with 13 GB RAM. As evident, since English had the highest count of documents as well as entities and mentions, it took the maximum CPU time for annotation.

An example of the densely annotated text in JSON format is given below:

Listing 1: Annotated example in JSON file format

```

1 {
2   "tokens": [{
3     "raw_form": "Vlade"
4   }, {
5     "raw_form": "Divac"
6   }, {
7     "raw_form": "is"
8   }, {
9     "raw_form": "a"
10  }, {
11   "raw_form": "retired"
12  }, {
13   "raw_form": "Serbian"
14  }, {
15   "raw_form": "NBA"
16  }, {
17   "raw_form": "player",
18   "break": ["SENTENCE"]
19  }],
20  "entities": [{
21    "id_str": "01vpr3",
22    "type": "PERSON",
23    "start_position": 0,
24    "end_position": 1,
25    "raw_form": "Vlade Divac"
26  }, {
27    "id_str": "077qn",
28    "type": "LOCATION",
29    "start_position": 5,
30    "end_position": 5,
31    "raw_form": "Serbian"
32  }, {
33    "id_str": "05jvx",
34    "type": "ORGANIZATION",
35    "start_position": 6,
36    "end_position": 6,
37    "raw_form": "NBA"
38  }],

```

Mention	Occurrence Count
Apple	16104
apple	2742
Tesla	822

Table 2: Mention Occurrence Counts

Entity	Occurrence Count
0k8z-Apple Inc	39624
014j1m-Apple (fruit)	8727
05dly-Nikola Tesla	2777

Table 3: Entity Occurrence Counts

```

39   "id": "wiki_page_id:en:322505:01vpr3:
40   Vlade_Divac"
}

```

6. DERIVED DATASETS

We also derive and open several other datasets from the DAWT dataset which we discuss here.

6.1 Mention Occurrences

This dataset includes the raw occurrence counts for a mention M_i in our corpus and KB. Table 2 shows the raw counts for mentions “Apple”, “apple” and “Tesla”.

6.2 Entity Occurrences

This dataset includes the raw occurrence counts for an entity E_j in our corpus and KB. Table 3 shows the raw counts for entities *Apple Inc.*, *Apple (fruit)* and *Nikola Tesla*. We also generate separate dictionaries for each language. Table 4 shows the different surface form variations and occurrence counts of the same entity across different languages.

6.3 Mention To Entity Co-occurrences

This dataset includes the co-occurrence counts of mentions and entities. This is particularly useful for estimating the prior probability of a mention M_i referring to a candidate entity E_j with respect to our KB and corpora. Table 5 shows the raw and normalized mention entity co-occurrences for the mentions “Apple” and “apple” and different candidate entities. As evident, the probability of mention “Apple” referring to the entity *Apple Inc.* is higher than to the entity *Apple (fruit)*. However, “apple” most likely refers to the entity *Apple (fruit)*. Similarly, the mention “Tesla” most likely refers to the entity *Nikola Tesla*.

6.4 Freebase id to Wikidata id Mappings

In this data set, we use the Freebase machine id to represent an entity. To facilitate studies using Wikidata ids, which are also widely used entity ids in the literatures, we provide a data set that maps individual Freebase ids to Wikidata ids. This data set contains twice as many mappings as

Entity	Language	Surface Form	Occurrence	Normalized Occurrence
0k8z-Apple Inc	English	Apple	35166	71.97%
		Apple s	4749	9.72%
		Apple Inc	2773	5.67%
		apple com	2232	4.57%
	French	Apple Computer	1534	3.14%
		Apple	6305	85.97%
		apple	466	6.35%
		d Apple	228	3.11%
Apple Inc	Apple Inc	109	1.49%	
	Apple Computer	96	1.31%	
014j1m-Apple (fruit)	English	apple	4215	51.72%
		Apple	1499	18.39%
		apples	1409	17.29%
		Apples	188	2.31 %
		Malus	139	1.71 %
	French	pomme	2188	53.80%
		pommes	1354	33.29%
		Pomme	394	9.69%
Pommes	81	1.99%		
pommerai	13	0.32%		
05d1y-Nikola Tesla	English	Tesla	2391	66.79 %
		Nikola Tesla	1043	29.13%
		Nikola Tesla s	71	1.98%
		Tesla Nikola	28	0.78%
		Nicola Tesla	13	0.36%
	French	Tesla	587	67.94%
		Nikola Tesla	270	31.25%
		Nicolas Tesla	5	0.58%
Nicola Tesla	1	0.11%		
nikola tesla	1	0.11%		
0dr90d-Tesla Motors	English	Tesla	1559	74.63%
		Tesla Motors	437	20.92%
		Tesla Roadster	22	1.05%
		teslamotors	11	0.53%
		Tesla Motors Inc	9	0.43%
	French	Tesla	110	68.32%
		Tesla Motors	49	30.44%
		voitures Tesla	2	1.24%

Table 4: Entity surface form variation across languages (English and French)

that from Google³. A summary comparison between these two mapping sets are shown in Table 6, which lists the total numbers of mappings in 4 buckets:

- **Same:** A Freebase id maps to a same Wikidata id.
- **Different:** A Freebase id maps to different Wikidata ids.
- **DAWT Only:** A Freebase id only maps to a Wikidata id in DAWT.
- **Google Only:** A Freebase id only maps to a Wikidata id in Google.

Note that the 24,638 different mappings are mainly caused by multiple Wikidata ids mapping to a same entity. For example, Freebase id 01159r maps to Q6110357 in DAWT and Q7355420 in Google, and both Q6110357 and Q7355420 represent the town Rockland in Wisconsin.

6.5 Entity Embeddings

There have been many efforts on learning word embeddings, i.e., vector space representations of words [6, 14, 15].

³<https://developers.google.com/freebase#freebase-wikidata-mappings>

Such representations are very useful in many tasks, such as word analogy, word similarity, and named entity recognition. Recently, Pennington et al. [17] proposed a model, GloVe, which learns word embeddings with both global matrix factorization and local context windowing. They showed that obtained embeddings captured rich semantic information and performed well in the aforementioned tasks.

There are several word-vector data sets available on GloVe’s website⁴. However, they only contain embeddings of individual words and thus have several limitations:

- Language dependent
- Missing entities that cannot be represented by a single word
- May not properly represent ambiguous words, such as "apple", which can be either the fruit or the technology company.

To facilitate research in this direction, we provide an entity-embedding data set that overcomes the above limitations. This data set contains embeddings of Wiki entities with 3

⁴<http://nlp.stanford.edu/projects/glove/>

Mention	Entity	Co-occurrence	Normalized Co-occurrence
Apple	0k8z-Apple Inc.	6738	87.8%
	014j1m-Apple (fruit)	422	5.5%
	019n_t-Apple Records	302	3.9%
	02hwrl-Apple Store	87	1.1%
	02_7z_-Apple Corps.	84	1.1%
apple	014j1m-Apple	1295	85.3%
	01qd72-Malus	157	10.3%
	02bjnm-Apple juice	46	3.0%
	0gjvjk-The Apple Tree	15	1.0%
	0k8z-Apple Inc.	1	0.1%
Tesla	05d1y-Nikola Tesla	327	49.5%
	0dr90d-Tesla Motors	162	24.5%
	036wfx-Tesla (Band)	92	13.9%
	02rx3cy-Tesla (Microarchitecture)	38	5.7%
	03rhvb-Tesla (Unit)	29	4.4%

Table 5: Mention Entity Co-occurrences

Same	2,048,531
Different	24,638
DAWT Only	2,362,077
Google Only	26,413

Table 6: Comparison of Freebase id to Wikidata id Mappings

different vector sizes: 50, 300, and 1000. They were generated with the GloVe model via the following steps:

1. Represent each Wiki document across all languages as a list of entities: There are about 2.2B total entities and 1.8M unique entities in these documents.
2. Use the open source GloVe code⁵ to process these documents: For each vector size, we ran 300 iterations on a GPU box with 24 cores and 60G dedicated memory. Other runtime configurations were the same as default. In particular, we:
 - Truncate entities with total count < 5
 - Set window size to be 15

Execution time was roughly proportional to the vector size. It took about 25 minutes to run 1 iteration when size is 1000. Among the 1.8 M unique entities, the GloVe model were able to generate embeddings for about 1.6 M entities.

To evaluate these embeddings, we compared them with one of the GloVe word embeddings, which was also generated from Wikipedia data⁶, on the word/entity analogy task. This task is commonly used to evaluate embeddings [14, 15, 17] by answering the following question: *Given word/entity X, Y, and Z, what is the word/entity that is similar to Z in the same sense as Y is similar to X?* For example, given word "Athens", "Greece", and "Paris", the right answer is "France".

Here we used the test data provided by Mikolov et al. [14]⁷. This test set contains 5 semantic and 9 syntactic relation types. For each word in this data set, we find the corresponding Freebase id using the mapping between Freebase ids and english Wikidata urls. Thus, we obtain a test set that contains relations between entities. Note that when

⁵<https://github.com/stanfordnlp/GloVe>

⁶The data is available at <http://nlp.stanford.edu/data/glove.6B.zip>

⁷The data set is available at <http://www.fit.vutbr.cz/~imikolov/rnnlm/word-test.v1.txt>

we could not find a Freebase id for a word, all the associated relations were removed from the test set.

We then ran the test on the 5 semantic types. Syntactic relations were excluded from this test because most of the time the task is trivial when one can correctly link words to entities. For example, when both "bird" and "birds" are linked to entity *015p6-Bird*, and "cat" and "cats" are linked to entity *01yrx-Cat*, the analogy among them is obvious without examining the underlining embeddings.

Table 7 shows the accuracy (in %) obtained from our entity embeddings with vector sizes of 50, 300, and 1000. In comparison, it also shows the accuracy from GloVe word embeddings with vector sizes of 50, 100, 200, and 300. Entity embeddings have better performance with vector size of 50. As we increase vector size, word embeddings perform significantly better and outperform entity embeddings when the vector size is 200 or higher. The degraded performance of entity embeddings may due to less training data, since our entity embeddings were obtained from 2.2B tokens, where GloVe's word embeddings were obtained from 6B tokens.

7. APPLICATIONS OF THE DATASET

As discussed earlier, the DAWT and other derived datasets that we have described in this paper have several applications for the NLP and IR communities. These include:

7.1 Named Entity Recognition (NER)

This task involves identifying an entity mention within a text and also generating candidate entities from the KB. For this, the Mention Occurrence, Entity Occurrence and the Mention To Entity Co-occurrence datasets described in Sections 6.1, 6.2 and 6.3 are extremely useful. For instance, the raw mention occurrence counts and probabilities can be stored in a dictionary and can be used to extract the mentions in a document. Furthermore, the mention occurrence count of a mention M_i and its co-occurrence count with an entity E_j can be used to calculate the prior probability of the mention mapping to that entity:

$$\frac{\text{count}(M_i \rightarrow E_j)}{\text{count}(M_i)}$$

This can be used to determine the candidate entities for a mention from our KB.

Relation	GloVe Word dimensionality				DAWT Entity dimensionality		
	50	100	200	300	50	300	1000
Capital-World	74.43	92.77	97.05	97.94	93.24	93.95	91.81
City-in-State	23.22	40.10	63.90	72.59	68.39	88.98	87.90
Capital-Common-Countries	80.04	95.06	96.64	97.23	78.66	79.64	71.54
Currency	17.29	30.05	37.77	35.90	43.88	13.56	2.93
Family	71.05	85.09	89.18	91.23	66.96	72.51	75.15
Average	53.21	68.61	76.91	78.98	70.23	69.73	65.87

Table 7: Accuracy of Semantic Analogy

7.2 Entity Disambiguation and Linking (EDL)

This task involves linking the mention to its correct KB entity. For each mention, there may be several candidate entities with prior probabilities as calculated in NER. In addition, other features derived from these datasets can include entity co-occurrences, entity word2vec similarity and lexical similarity between the mention and the entity surface form. These features can be used to train a supervised learning algorithm to link the mention to the correct disambiguated entity among all the candidate entities as done in [1]. The approach used in [1] employs several such context dependent and independent features and has a precision of 63%, recall of 87% and an F-score of 73%.

8. ACCESSING THE DATASET

The DAWT and derived datasets discussed in paper are available for download at this page: https://github.com/klout/opensource/tree/master/wiki_annotation. The DAWT dataset was generated using Wikipedia Data Dumps from January 20th 2017. Statistics regarding the data set are shown in Table 1.

9. RELATED WORK

While a lot of works have focused on building and opening such datasets, very few have addressed all the challenges and constraints that we mentioned in Section 2. Spitkovsky and Chang [22] opened a cross-lingual dictionary (of English Wikipedia Articles) containing 175,100,788 mentions linking to 7,560,141 entities. This dataset, though extremely valuable, represents mention - entity mappings across a mixture of all languages which makes it harder to use for a specific language. In addition, this work used raw counts that, although useful, lack mention context (such as preceding and succeeding tokens etc.) which have a big impact while performing EDL. The Freebase annotations of the ClueWeb corpora dataset⁸ dataset contains 647 million English web pages with an average of 13 entities annotated per document and 456 million documents having at least 1 entity annotated. It does not support multiple languages.

Another related technique for generating such dictionaries is Wikification [12, 16, 4] where mentions in Wikipedia pages are linked to the disambiguated entities' Wikipedia pages. Such techniques rely on a local or global approach. A local approach involves linking observed entities using only their local context eg. by comparing the relatedness of candidate Wiki articles with the mentions [5, 7, 19] while in global approach entities across the entire document are disambiguated together using document context, thus, ensuring consistency of entities across the document [9, 18]. Most recently Cai et al. [3] achieved 89.97% precision and 76.43%

⁸<http://lemurproject.org/clueweb09/FACC1/>

recall, using an iterative algorithm that leverages link graph, link distributions, and a noun phrase extractor.

Although the problem of entity linking has been well studied for English, it has still not been explored for other languages. McNamee et al. [11] introduced the problem of cross-language entity linking. The main challenge here is that state-of-the-art part-of-speech taggers perform much better on English than on other languages. In addition, both Wikipedia and Freebase have significantly higher quality and coverage of English compared to any other language.

10. CONCLUSION AND FUTURE WORK

In this work, we opened up the DAWT dataset - Densely Annotated Wikipedia Texts across multiple languages. The annotations include labeled text mentions mapping to entities (represented by their Freebase machine ids) as well as the type of the entity. The data set contains total of 13.6M articles, 5.0B tokens, 13.8M mention entity co-occurrences. DAWT contains 4.8 times more anchor text to entity links than originally present in the Wikipedia markup. Moreover, it spans several languages including English, Spanish, Italian, German, French and Arabic. We also presented the methodology used to generate the dataset which enriched Wikipedia markup in order to increase number of links. In addition to the main dataset, we opened up several derived datasets for mention occurrence counts, entity occurrence counts, mention entity co-occurrence counts, entity word2vec as well as mappings between Freebase ids and Wikidata item ids. We also discussed two applications of these datasets and hope that opening them up would prove useful for the NLP and IR communities as well as facilitate multi-lingual research.

In the future, we plan to improve the algorithm that we used for generating DAWT. Also, we plan to migrate from using Freebase ids in our KB to Wikidata item ids.

11. REFERENCES


- [1] P. Bhargava, N. Spasojevic, and G. Hu. High-throughput and language-agnostic entity disambiguation and linking on user generated data. In *Proceedings of WWW 2017 workshop on Linked Data on the Web (LDOW'17)*, 2017.
- [2] P. Bhattacharyya and N. Spasojevic. Global entity ranking across multiple languages. In *Proceedings of the 26th International Conference on World Wide Web*, page to appear, 2017.
- [3] Z. Cai, K. Zhao, K. Q. Zhu, and H. Wang. Wikification via link co-occurrence. In *Proceedings of the 22nd ACM international conference on Conference on information and knowledge management, CIKM '13*, pages 1087–1096, New York, NY, USA, 2013. ACM.

- [4] X. Cheng and D. Roth. Relational inference for wikification. *Urbana*, 51:61801, 2013.
- [5] S. Cucerzan. Large-scale named entity disambiguation based on wikipedia data. In *EMNLP-CoNLL*, volume 7, pages 708–716, 2007.
- [6] S. Deerwester, S. T. Dumais, G. W. Furnas, T. K. Landauer, and R. Harshman. Indexing by latent semantic analysis. *Journal of the American society for information science*, 41(6):391, 1990.
- [7] P. Ferragina and U. Scaiella. Tagme: on-the-fly annotation of short text fragments (by wikipedia entities). In *Proceedings of the 19th ACM international conference on Information and knowledge management*, pages 1625–1628. ACM, 2010.
- [8] X. Glorot, A. Bordes, and Y. Bengio. Domain adaptation for large-scale sentiment classification: A deep learning approach. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pages 513–520, 2011.
- [9] S. Kulkarni, A. Singh, G. Ramakrishnan, and S. Chakrabarti. Collective annotation of wikipedia entities in web text. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 457–466. ACM, 2009.
- [10] J. Lehmann, R. Isele, M. Jakob, A. Jentzsch, D. Kontokostas, P. N. Mendes, S. Hellmann, M. Morsey, P. van Kleef, S. Auer, et al. Dbpedia-a large-scale, multilingual knowledge base extracted from wikipedia. *Semantic Web*, 6(2):167–195, 2015.
- [11] P. McNamee, J. Mayfield, D. Lawrie, D. W. Oard, and D. S. Doermann. Cross-language entity linking. In *IJCNLP*, pages 255–263, 2011.
- [12] O. Medelyan, D. Milne, C. Legg, and I. H. Witten. Mining meaning from wikipedia. *International Journal of Human-Computer Studies*, 67(9):716–754, 2009.
- [13] R. Mihalcea and A. Csomai. Wikify!: Linking documents to encyclopedic knowledge. In *Proceedings of the Sixteenth ACM Conference on Conference on Information and Knowledge Management, CIKM '07*, pages 233–242, 2007.
- [14] T. Mikolov, K. Chen, G. Corrado, and J. Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013.
- [15] T. Mikolov, I. Sutskever, K. Chen, G. Corrado, and J. Dean. Distributed representations of words and phrases and their compositionality. In *Proceedings of NIPS, 2013*, 2013.
- [16] D. Milne and I. H. Witten. Learning to link with wikipedia. In *Proceedings of the 17th ACM Conference on Information and Knowledge Management, CIKM '08*, pages 509–518, 2008.
- [17] J. Pennington, R. Socher, and C. D. Manning. Glove: Global vectors for word representation. In *Proceedings of EMNLP, 2014*, page 1532–1543, 2014.
- [18] L. Ratinov, D. Roth, D. Downey, and M. Anderson. Local and global algorithms for disambiguation to wikipedia. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*, pages 1375–1384. Association for Computational Linguistics, 2011.
- [19] B. Skaggs and L. Getoor. Topic modeling for wikipedia link disambiguation. *ACM Transactions on Information Systems (TOIS)*, 32(3):10, 2014.
- [20] N. Spasojevic, P. Bhattacharyya, and A. Rao. Mining half a billion topical experts across multiple social networks. *Social Network Analysis and Mining*, 6(1):1–14, 2016.
- [21] N. Spasojevic, J. Yan, A. Rao, and P. Bhattacharyya. Lasta: Large scale topic assignment on multiple social networks. In *Proc. of ACM Conference on Knowledge Discovery and Data Mining (KDD)*, KDD '14, 2014.
- [22] V. I. Spitzkovsky and A. X. Chang. A cross-lingual dictionary for english wikipedia concepts. In *Language Resources and Evaluation*, 2012.

APPENDIX

A. SAMPLE ANNOTATED WIKIPEDIA TEXTS FROM DAWT

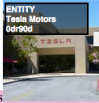

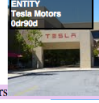
Figure 2 shows samples of the densely annotated Wikipedia pages for the entities *Nikola Tesla* and *Tesla Motors* across English and Arabic.

Nikola Tesla  10 July 1856 7 January 1943 was a **Serbian American** inventor electrical engineer mechanical engineer physicist and futurist best known for his contributions to the design of the modern **alternating current** AC **electricity supply** system Tesla gained experience in **telephony** and **electrical engineering**

(a) Annotated wikipedia article on Nikola Tesla (English)

ميكانيكي ومهندس  مهندس  مخترع  1943  7 يناير  1856  10 يوليو  نيكولا تيسلا  صربيا  قى بلغراد  كهراباني  في عهد الامبراطورية النمساوية  نال الجنسية الأمريكية  في عهد الامبراطورية النمساوية  كرواتيا  ولد في سميلجان  كهراباني 

(b) Annotated wikipedia article on Nikola Tesla (Arabic)

Tesla Motors  Inc is an American automotive and energy storage company that designs manufactures and sells electric car s electric vehicle powertrain components and battery products  Tesla Motors  is a public company that trades on the NASDAQ stock exchange under the symbol TSLA

(c) Annotated wikipedia article on Tesla Motors (English)

يتم تداول أسهمها في بورصة  NO IMAGE  NO IMAGE  NO IMAGE  NO IMAGE  NO IMAGE  NO IMAGE  NO IMAGE  NO IMAGE  NO IMAGE  NO IMAGE  NO IMAGE  NO IMAGE  NO IMAGE  NO IMAGE  NO IMAGE  NO IMAGE  NO IMAGE  NO IMAGE  NO IMAGE  NO IMAGE  NO IMAGE  NO IMAGE  NO IMAGE  NO IMAGE

(d) Annotated wikipedia article on Tesla Motors (Arabic)

Figure 2: Samples of extracted text across different Languages