

Building Multilingual Corpora for a Complex Named Entity Recognition and Classification Hierarchy using Wikipedia and DBpedia

Diego Alves¹[0000-0001-8311-2240], Gaurish Thakkar¹[0000-0002-8119-5078],
Gabriel Amaral²[0000-0002-4482-5376], Tin Kuculo³[0000-0001-6874-8881], and
Marko Tadić¹[0000-0001-6325-820X]

¹ Faculty of Humanities and Social Sciences, University of Zagreb, Zagreb 10000,
Croatia

{dfvalio,marko.tadic}@ffzg.hr, gthakkar@m.ffzg.hr

² King's College London, London, United Kingdom gabriel.amaral@kcl.ac.uk

³ L3S Research Center, Leibniz University Hannover, Hannover, Germany
kuculo@l3s.de

Abstract. With the ever-growing popularity of the field of NLP, the demand for datasets in low resourced-languages follows suit. Following a previously established framework, in this paper¹, we present the UNER dataset, a multilingual and hierarchical parallel corpus annotated for named-entities. We describe in detail the developed procedure necessary to create this type of dataset in any language available on Wikipedia with DBpedia information. The three-step procedure extracts entities from Wikipedia articles, links them to DBpedia, and maps the DBpedia sets of classes to the UNER labels. This is followed by a post-processing procedure that significantly increases the number of identified entities in the final results. The paper concludes with a statistical and qualitative analysis of the resulting dataset.

Keywords: named-entity · multilingualism · data extraction

1 Introduction

Named entity recognition and classification (NERC) is an essential Natural Language Processing (NLP) task involved in many applications like interactive question answering, summarizing, relation extraction, and text mining. It was originally defined in the 6th Message Understanding Conference (MUC-6) as the identification of “Person”, “Location” and “Organization” types [3]. As shown by Alves et al. [1], NERC corpora usually respond to the specific needs of local projects and differ in terms of the complexity of the types’ hierarchy and format.

NERC systems that depend on machine learning models need to be trained with a large amount of language-specific annotated data to attain high accuracy.

¹ Copyright © 2021 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

Nevertheless, annotating this type of data is expensive, time-consuming, and requires specific annotators, training to guarantee good inter and intra-annotator agreements [5].

The complexity of generating NERC annotated data increases when multilingualism and a complex hierarchy of types are involved [12]. The Universal Named Entity Recognition (UNER) framework proposes a hierarchy composed of 3 levels²: the first one covers eight broad categories, the second is composed of 47 named-entity types, and 15 of which are detailed in a third level with a total of 69 subtypes [1]. Therefore, our challenge was to generate automatically annotated data (silver standard) following the UNER hierarchy.

Our goal is to parse data from Wikipedia³ corpora in multiple languages, extract named entities through hyperlinks, align them with entity classes from DBpedia⁴ [6] and translate them into UNER types and subtypes. The idea is to propose a NERC dataset creation workflow that works for languages covered by both Wikipedia and DBpedia, including under-resourced languages.

This paper presents the data extraction and annotation workflow and its application on English⁵ and Croatian⁶ languages under the licence CC BY-NC-SA 4.0. It is organized as follows: In Section 2, we present the state-of-the-art concerning NERC automatic data generation; in Section 3, we describe the details of the data extraction and annotation workflow, and in Section 4, we provide statistical analysis and qualitative evaluation of the annotation. Section 5 is dedicated to the discussion of the results, and in Section 6, we present our conclusions and possible future directions for research, including applications of the UNER corpora.

2 Related Work

Due to the importance of the NERC task, the challenge of creating a quality annotated dataset has been the object of many studies.

Yadav and Bethard [15] and Li et al. [7] provided complete reviews on NERC methods, however, focusing on data extraction using existing training corpora, not detailing procedures for generating new data. Nevertheless, both articles agree on the fact that a large amount of quality data is required for this task.

To overcome the problem of cost, Lawson et al. [5] proposed the usage of Amazon Mechanical Turk to annotate a specific corpus composed of e-mails. However, this method is still time-consuming and requires specific training and bonus strategies. It also relies on the availability of annotators for each specific language if the aim is to generate multilingual corpora.

A generic method for extracting MWEs from Wikipedias was proposed in Bekavac and Tadić [2] and this covers also Multi Word Extraction named entities

² <https://tinyurl.com/sb3u9ve>

³ <https://www.wikipedia.org/>

⁴ <https://wiki.dbpedia.org/>

⁵ Link for English UNER Corpus v.1: <https://tinyurl.com/y2taxs8b>

⁶ Link for Croatian UNER Corpus v.1: <https://tinyurl.com/y4tlz4a2>

using local regular grammars. An automatic multilingual solution is presented by Ni et al. [9]. Their approach was to use annotation projection on comparable corpora or to project distributed representations of words (embeddings) from a target to a source language and, therefore, use the NER model of the source language to annotate the target one without retraining. Although with promising results, these methods require at least one quality annotated corpus (source language), which is not the case in our workflow, and that can be problematic for complex hierarchies such as UNER.

Kim et al. [4] propose an automatic method (semi-CRF model) to label multilingual data with named entity tags using Wikipedia metadata and parallel sentences (English-foreign language) extracted from this database. This method still requires manual annotation of articles-pairs for the CRF training and the final corpus is restricted only to the encountered parallel sentences. Additionally, using Wikipedia metadata, Nothman et al. [10] present a workflow used for the generation of silver-standard NERC corpora for English, German, Spanish, Dutch, and Russian. The idea was to transform the links between articles into new annotations by projecting the target article’s classifications onto the anchor text. This method also requires previous manual annotation of a considerable number of Wikipedia articles.

A complete automatic workflow is proposed by Weber & Vieira [8] for Portuguese NERC corpus generation by using Wikipedia and DBpedia information. The proposed method is the basis of our study, but we extend it to multiple languages and to a more complex NERC hierarchy.

3 Process for Data Extraction and Annotation

The workflow we have developed allows the extraction of texts and metadata from Wikipedia (for any language present in this database), followed by the identification of the DBpedia classes via the hyperlinks associated with certain tokens (entities) and the translation to UNER types and sub-types (these last two steps being language independent).

Once the main process of data extraction and annotation is over, the workflow proposes post-processing steps to improve the tokenization and implement the IOB format [11]. Statistical information concerning the generated corpus is gathered, and missing UNER entities are automatically identified.

The whole workflow is presented in detail in the project GitHub web-page⁷ together with all scripts that have been used, and that can be applied to any other Wikipedia language.

3.1 Components

The following items are important components that were used for the dataset creation. These are mappers that map entities from the source to a target class or hierarchy.

⁷ <https://github.com/cleopatra-itn/MIDAS>

1. **UNER/DBpedia Mapping:** This is a mapper that connects each pertinent DBpedia class with a single UNER tag. A single extracted named entity might have more than one DBpedia class. For example, entity *2015 European Games* have the following DBpedia classes with the respective UNER equivalences:

- **dbo:Event** – *Name-Event-Historical-Event*
- **dbo:SoccerTournament** – *Name-Event-Occasion-Game*
- **dbo:SocietalEvent** – *Name-Event-Historical-Event*
- **dbo:SportsEvent** – *Name-Event-Occasion-Game*
- **owl:Thing** – *NULL*

The value on the left represents a DBpedia class and its UNER equivalent is on the right side of the class. It maps all the DBpedia classes to UNER equivalent classes.

2. **DBpedia Hierarchy:** This mapper assigns priorities to each DBpedia class. This is used to select a single DBpedia class from the collection of classes that are associated with an entity. Following are classes are their priorities.

- **dbo:Event** – 2
- **dbo:SoccerTournament** – 4
- **dbo:SocietalEvent** – 2
- **dbo:SportsEvent** – 4
- **owl:Thing** – 1

For entity *2015 European Games*, the DBpedia class **SoccerTournament** presides over the other classes as it has a higher priority value. If the extracted entity has two assigned classes with the same hierarchy value the first from the list is chosen as the final one. All the DBpedia classes were assigned with a hierarchy value according to DBpedia Ontology⁸.

3.2 Main process

The main process is schematized in the figure below and is divided into three sub-processes.

1. **Extraction from Wikipedia dumps:** For a given language, we obtain its latest dump from the Wikimedia website⁹. Next, we perform text extraction preserving the hyperlinks in the article using WikiExtractor¹⁰. These are hyperlinks to other Wikipedia pages as well as unique identifiers to those named-entities. We extract all the unique hyperlinks and sort them alphabetically. These hyperlinks will be referred to as named-entities henceforth.
2. **Wikipedia-DBpedia entity linking:** For all the unique named-entities from the dumps, we query the DBpedia endpoint using a SPARQL query with SPARQLWrapper¹¹ to identify the various classes associated with the entity. This step produces, for each named-entity from step 1, a set of DBpedia classes it belongs to.

⁸ <http://mappings.dbpedia.org/server/ontology/classes/>

⁹ <https://dumps.wikimedia.org/>

¹⁰ <https://github.com/attardi/wikiextractor>

¹¹ <https://rdflib.dev/sparqlwrapper/>

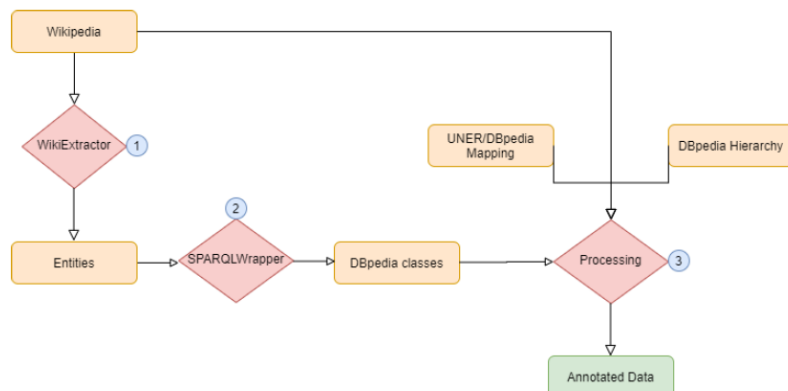


Fig. 1. Main process steps for Wikipedia data extraction and DBpedia/UNER annotations. Squares represent data and diamonds represent processing steps.

3. **Wikipedia-DBpedia-UNER back-mapping:** For every extracted named-entity obtained in step 1, we use the set of classes produced in step 2, along with an UNER/DBpedia mapping schema, to assign UNER classes to each named-entity. For an entity, all the classes obtained from the DBpedia response are mapped to a hierarchy value, highest valued class is resolved and chosen and then it is mapped to UNER class. For constructing the final annotation dataset, we only select those sentences that have at least one single named entity. This reduces the sparsity of annotations and thus reduces the false negatives rate in our test models. This step produces an initial tagged corpus from the whole Wikipedia dump for a specific language.

3.3 Post-processing steps

The post-processing steps correspond to four different Python scripts that provide:

1. The improvement of the tokenization (using regular expressions) by isolating punctuation characters that were connected with words. In addition, it applies the IOB format to the UNER annotations inside the text.
2. The calculation of the following statistic information concerning the generated corpus: Total number of tokens, Number of Non-entity Tokens (tag “O”), Number of Entity Tokens (tags “B” or “I”) and Number of Entities (tag “B”). The script also provides a list of all UNER tags with the number of occurrences of each tag inside the corpus.
3. Listing the entities inside the corpus (tokens and the corresponding UNER tag). Each identified entity appears once in this list, even if it has multiple occurrences in the corpus.

4. The increment of the annotations by an automatic process using the list of entities (output of item 3 above). The script checks if, inside the corpus, entity tokens from the mentioned list were not tagged as a UNER entity. If so, the corresponding UNER tag is associated to them. Entities composed of one single character, two lower-case characters, and digits were discarded. Also, multi-tokens entities were prioritized using the longest match approach.

Post-processing steps from 1 to 4 were applied to the UNER Croatian corpus, while for the English one, only steps 1 to 3 due to its size. Results and qualitative analysis are presented in the following section.

4 Data Analysis

We have decided to test the proposed workflow with two different languages: English and Croatian. The Croatian Wikipedia [13] is composed of 223,488 articles, while the English Wikipedia [14] has 6,188,204 articles (almost 28 times bigger). The aim was to check the compatibility of the process with a very well-resourced language and a low-resourced one. We have applied the main process for both languages and post-processing steps from 1-3 to the English corpus. The Croatian corpus has also been passed through the fourth step of post-processing.

4.1 Statistical Information

After applying the main process of the proposed workflow, we obtain, for each language, annotated text files divided into folders. The sizes of the English and Croatian UNER corpora are presented in the following table.

Table 1. Corpora Size.

Corpus	Total Size	Number of folders	Number of files
English UNER	3.3 GB	172	17,150
Croatian UNER	108 MB	5	411

After applying post-processing steps 1-3 in both Croatian and English UNER datasets, we obtain the following statistics concerning both corpora.

Therefore, concerning the English Corpus, 8.9% of tokens are entities, 7.1% in the Croatian one.

As explained previously, the UNER hierarchy is composed of categories, types, and subtypes. UNER includes the most common classes used in NERC (“Person”, “Location”, “Organization”, being more detailed (subtypes):

- “Person”: correspond to UNER “Name-Person-Name”
- “Location”: correspond to all subtypes inside the UNER types “Name-Location”.
- “Organization”: correspond to all subtypes inside the UNER type “Name-Organization”

Table 2. Corpora Annotation Statistics.

	English UNER Corpus	Croatian UNER Corpus
Total Number of Tokens	325,395,838	9,388,224
Number of Non-Entity Tokens	320,719,350	8,436,254
Number of Entity Tokens	31,676,488	951,970
Number of Entities	15,101,318	668,231
Number of Different Entities	630,519	157,418

Therefore, it is possible to analyse the generated corpora in terms of these more generic classes.

Table 3. Corpora Annotation Statistics in terms of number of occurrences of the most used NERC classes (and % of all entities occurrences).

	English UNER Corpus	Croatian UNER Corpus
Person	4,200,313 (27.8%)	179,440 (26.8%)
Location	2,613,248 (17.3%)	248,801 (37.2%)
Organization	3,489,813 (23.1%)	74,564 (11.5%)

These main classes correspond to 68.2% of NEs in the English corpus and 75.5% in the Croatian one. As explained in section 3.1, the annotation of a certain named-entity depends on the existence of hyperlinks. However, these links are not always associated with the tokens if the entity is mentioned repeatedly in the article. Therefore, to improve the corpus generated through the proposed workflow, we have established another post-processing step, which we applied to the Croatian corpus. By using the list of entities present in the corpus, we identify in the text the tokens that should have been tagged as entities, but that did not have the corresponding hyperlink.

In the following table, we present the statistics of the final post-processed Croatian corpus compared to the one obtained before this last step.

Table 4. Croatian Corpus statistics after final post-processing step and delta compared to the previous corpus.

	Post-processed Croatian UNER Corpus	Delta
Number of Non-Entity Tokens	7,858,052	-578,202
Number of Entity Tokens	1,530,172	+578,202
Number of Entities Occurrences	1,217,633	+549,402

The percentage of entity tokens is increased from 7.1% to 16.2%. In the next table, we focus the statistical analysis on the UNER types that can be associated with the classic NERC classes: *Person*, *Location* and *Organisation*.

Table 5. Croatian UNER Corpus Statistics in terms of number of occurrences of the most used NERC classes (and % of all entities occurrences) and delta compared to the previous corpus.

	Post-Processed Croatian UNER Corpus	Delta
Person	286,593 (23.5%)	+107,153
Location	453,985 (37.3%)	+205,184
Organization	172,471 (14.2%)	+97,907

4.2 Qualitative evaluation

The automatic annotation of the text extracted from Wikipedia requires the identification of the DBpedia classes associated with the respective tokens (via hyperlink) and the translation to UNER using (UNER/DBpedia equivalences). In order to evaluate this step, we have performed an analysis of 943 entities randomly selected concerning the English UNER Corpus. For each one, we have checked the DBpedia associated classes and the final UNER chosen tag. The next table presents the results of this evaluation.

Table 6. Evaluation of the annotation step: DBpedia class extraction and translation to UNER hierarchy.

Tag Evaluation	Number of Occurrences	Percentage
Correct	797	85%
Correct but vague	55	6%
Incorrect due to DBpedia	62	7%
Incorrect due to UNER association	29	3%

The analysis of this sample shows that 91% of the entities are correctly tagged with UNER tags. However, 6% are associated to the correct UNER type but to a generic sub-type. For example, *Bengkulu* should be tagged as *Name-Location-GPE-City* but received the tag *Name-Location-GPE-GPE_Other*.

The incorrect tags are due to errors in the DBpedia classes associated with the tokens or due to the equivalence between DBpedia and UNER rules:

- *Buddhism* is associated only to the DBpedia class *EthnicGroup* and, therefore, is wrongly tagged as *Name-Organization-Ethnic_Group_other* while it should be associated to the UNER tag *Name-Product-Doctrine_Method-Religion*.
- *Brit Awards*, due to the prioritization of DBpedia class hierarchy in the choice of UNER tags, is wrongly tagged as *Name-Organization-Corporation-Company* while it should receive the tag *Name-Product-Award*.

It is possible to observe a considerable number of False Negative instances inside the corpora not being processed by the last post-processing step. This is due to the fact that not all entities in Wikipedia extracted articles have hyperlinks.

A qualitative analysis of the final Croatian UNER corpus shows that the final post-processing step considerably reduces the number of False Negative instances. Nevertheless, with the actual set of rules for this final automatic annotation step, we can observe some problems concerning mostly mono-tokens entities, for example, all *jezik* instances are associated with UNER tag *Name-Product-Printing-Magazine* while in the text it may not correspond to a magazine but to the common noun *jezik* (*language in English*). Also, it is important to mention that while the UNER hierarchy proposes a very complex hierarchy in terms of time and numerical expressions, this workflow covers basically UNER types and sub-types of the *Name* category.

5 Discussion

The implemented workflow and the generated corpora show the potential of the method presented by Weber & Vieira [8] comprising the extraction of texts and metadata from Wikipedia and the usage of DBpedia to annotate in terms of NERC. Our approach is broader, proposing a method that can be used for any language present in Wikipedia and using a more complex NERC hierarchy (UNER). Compared to other proposals of generating silver NERC standards [4][9][10], our workflow has the advantage of being fully automatized. However, we can identify that some improvement is needed concerning, more specifically, the last post-processing step. We have succeeded in increasing the number of annotated entities; nevertheless, the accuracy needs to be enhanced. In addition, a manual and more detailed evaluation is necessary to verify the precision and recall of the classification of the entities inside the final annotated texts.

6 Conclusions and Future Directions

In this paper, we describe an automatic process for generating multilingual Named-Entity recognition corpora by using Wikipedia and DBpedia data. We also present the UNER corpus, a hierarchical named-entity corpus in Croatian and English, developed with the proposed method. We show some statistics of the corpus and detail the procedure used to create it, finishing with a qualitative evaluation. In our future work, we plan to extend our corpus to other under-resourced languages while evaluating our workflow’s performance across the languages. Also analysing the different limitations due to the uneven content of Wikipedia across languages. Following this intrinsic evaluation of our dataset, we will train models on the obtained data to extrinsically evaluate the corpus.

7 Acknowledgements

The work presented in this paper has received funding from the European Union’s Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement no. 812997 and under the name CLEOPATRA (Cross-lingual Event-centric Open Analytics Research Academy).

References

1. Alves, D., Kuculo, T., Amaral, G., Thakkar, G., Tadić, M.: Uner: Universal named-entity recognition framework. In: Proceedings of the 1st International Workshop on Cross-lingual Event-centric Open Analytics. pp. 72–79. Association for Computational Linguistics (2020), <https://www.aclweb.org/anthology/W10-0712>
2. Bekavac, B., Tadić, M.: A generic method for multi word extraction from wikipedia. In: Proceedings of the 30th International Conference on Information Technology Interfaces (2008), <https://www.bib.irb.hr/348724>
3. Chinchor, N., Robinson, P.: Appendix E: MUC-7 Named Entity Task Definition (version 3.5). In: Seventh Message Understanding Conference (MUC-7): Proceedings of a Conference Held in Fairfax, Virginia, April 29 - May 1, 1998 (1998), <https://www.aclweb.org/anthology/M98-1028>
4. Kim, S., Toutanova, K., Yu, H.: Multilingual named entity recognition using parallel data and metadata from Wikipedia. In: Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). pp. 694–702. Association for Computational Linguistics, Jeju Island, Korea (Jul 2012), <https://www.aclweb.org/anthology/P12-1073>
5. Lawson, N., Eustice, K., Perkowitz, M., Yetisgen-Yildiz, M.: Annotating large email datasets for named entity recognition with Mechanical Turk. In: Proceedings of the NAACL HLT 2010 Workshop on Creating Speech and Language Data with Amazon’s Mechanical Turk. pp. 71–79. Association for Computational Linguistics, Los Angeles (jun 2010), <https://www.aclweb.org/anthology/W10-0712>
6. Lehmann, J., Isele, R., Jakob, M., Jentzsch, A., Kontokostas, D., Mendes, P.N., Hellmann, S., Morsey, M., van Kleef, P., Auer, S., Bizer, C.: Dbpedia - a large-scale, multilingual knowledge base extracted from wikipedia. *Semantic Web* **6**(2), 167–195 (2015). <https://doi.org/10.3233/SW-140134>, <https://madoc.bib.uni-mannheim.de/37476/>
7. li, J., Sun, A., Han, R., Li, C.: A survey on deep learning for named entity recognition. *IEEE Transactions on Knowledge and Data Engineering* **PP**, 1–1 (03 2020). <https://doi.org/10.1109/TKDE.2020.2981314>
8. Menezes, D.S., Savarese, P., Milidiú, R.L.: Building a massive corpus for named entity recognition using free open data sources (2019), <http://arxiv.org/abs/1908.05758>
9. Ni, J., Dinu, G., Florian, R.: Weakly supervised cross-lingual named entity recognition via effective annotation and representation projection. *CoRR* **abs/1707.02483** (2017), <http://arxiv.org/abs/1707.02483>
10. Nothman, J., Ringland, N., Radford, W., Murphy, T., Curran, J.R.: Learning multilingual named entity recognition from wikipedia. *Artif. Intell.* **194**, 151–175 (2013), <http://dblp.uni-trier.de/db/journals/ai/ai194.html#NothmanRRMC13>
11. Ramshaw, L., Marcus, M.: Text chunking using transformation-based learning. In: Third Workshop on Very Large Corpora (1995), <https://www.aclweb.org/anthology/W95-0107>
12. Sekine, S.: Named entity: History and future (2004), <http://cs.nyu.edu/~sekine/papers/NEsurvey200402.pdf>
13. Wikipedia: Croatian Wikipedia — Wikipedia, the free encyclopedia. <http://en.wikipedia.org/w/index.php?title=Croatian%20Wikipedia&oldid=983689658> (2020), [Online; accessed 10-November-2020]
14. Wikipedia: English Wikipedia — Wikipedia, the free encyclopedia. <http://en.wikipedia.org/w/index.php?title=English%20Wikipedia&oldid=987449701> (2020), [Online; accessed 14-November-2020]

15. Yadav, V., Bethard, S.: A survey on recent advances in named entity recognition from deep learning models. CoRR **abs/1910.11470** (2019), <http://arxiv.org/abs/1910.11470>