

Building and Evaluating Universal Named-Entity Recognition English corpus

Diego Alves^[0000-0001-8311-2240], Gaurish Thakkar^[0000-0002-8119-5078], and
Marko Tadić^[0000-0001-6325-820X]

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

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

Abstract. This article presents the application of the Universal Named Entity framework to generate automatically annotated corpora. By using a workflow that extracts Wikipedia data and meta-data and DBpedia information, we generated an English dataset which is described and evaluated. Furthermore, we conducted a set of experiments to improve the annotations in terms of precision, recall, and F1-measure. The final dataset is available and the established workflow can be applied to any language with existing Wikipedia and DBpedia. As part of future research, we intend to continue improving the annotation process and extend it to other languages.

Keywords: named entity recognition · data extraction · multilingual
nlp

1 Introduction

Named entity recognition and classification (NERC) is an important field inside Natural Language Processing (NLP), being a crucial task of information extraction from texts. It was first defined in 1995 in the 6th Message Understanding Conference (MUC-6)[5] and since then has been used in multiple NLP applications such as events and relations extraction, question answering systems and entity-oriented search.

As shown by Alves et al.[1], NERC corpora and tools present an immense variety in terms of annotation hierarchy and formats. NERC hierarchy structure is usually locally defined considering the final NLP application where it will be used at. If certain types like "Person", "Location", and "Organization" are present in almost every NERC system, some corpora are composed of more complex types of annotation. It is the case, for example, of Portuguese Second HAREM[7], Czech Named Entity Corpus 2.0[18] and Romanian Ronec[6]. Multilingual alternatives also exist, it is the case of spaCy software[8] proposing two different single-level hierarchies composed of either 18 or 4 NERC types following OntoNotes 5.0[19] and Wikipedia[15] respectively.

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

Unlike Part-of-Speech and Dependency Parsing tagging which have Universal Dependencies¹, there is no universal alternative for NERC in terms of annotation framework and multilingual repository following the same standards.

Hence, we use the Universal Named Entity (UNER) framework which is composed of a complex NERC hierarchy inspired by Sekine’s work[17] and propose a process which parses data from Wikipedia², extract named entities through hyperlinks, aligns them with DBpedia³[11] entity classes and translates them into UNER types and subtypes. This process can be applied to any language present in both Wikipedia and DBpedia, therefore generating multilingual NERC corpora following the same procedure and hierarchy.

Thus, UNER is useful for multilingual NLP tasks which need recognition and classification of named entities going beyond classical NERC hierarchy involving only a few types. UNER data can be used in its totality or can be easily adapted to specific needs. For example, considering the classic “Location” type usually present in NERC corpora, UNER can be used to get more detailed information: if an entity is more specifically a country, mountain, island, etc.

This paper presents the UNER hierarchy and its workflow for data extraction and annotation. It details the application of the proposed process on English language with qualitative and quantitative evaluation of the automatically annotated data. It also presents the evaluation of different alternatives implemented for the improvement of the generated dataset.

This article is organized as follows: In Section 2, we present the state-of-the-art concerning NERC automatic data generation workflows; in Section 3, we describe UNER framework and hierarchy, and in Section 4 the details of the data extraction and annotation workflow and evaluation of the generated dataset. In Section 5, we report the experiments that were conducted for improving the dataset quality in terms of Precision and Recall. Section 6 is dedicated to the discussion of the results, and in Section 7, we present our conclusions and possible future directions for research.

2 Related Work

UNER framework was first introduced by us in a previous article [2], where we defined its hierarchy. In this article, the framework was revised and a workflow for automatic text annotation was developed and applied to generate an annotated corpus in English with the respective evaluation.

Deep learning has been employed in NERC systems in recent years improving stat-of-the-art performance, and, therefore, increasing the need of quality annotated datasets as stated by Yadav and Bethard[21] and Li et al.[12]. These authors have provided a large overview of existing techniques for extracting and classifying named entities using machine and deep learning methods.

¹ <https://universaldependencies.org>

² <https://www.wikipedia.org/>

³ <https://www.dbpedia.org/>

The problem of composing new NERC datasets has been the object of the study proposed by Lawson et al.[10]. Manual annotation of large corpora is usually very costly, the authors propose, then, the usage of Amazon Mechanical Turk as a low-cost alternative. However, this method still depends on a specific budget for the task and can be very time-consuming. It also depends on the availability of annotators for each specific language which may be problematic if the aim is to generate large multilingual corpora.

A generic method for extracting entities from Wikipedia articles was proposed in Bekavac and Tadić[3] and includes Multi-Word Extraction of named entities using local regular grammars. Therefore, for each targeted language, a new set of rules must be defined. Another automatic multilingual solutions has been proposed by Ni et al.[14] and Kim et al.[9] using either annotation projection on comparable corpora or Wikipedia metadata on parallel datasets. Both methods, however, still require manual annotations that are language-dependent and cannot be applied universally. Furthermore, Weber & Vieira[13] use a similar process to the one presented in this article for annotating Wikipedia texts using DBpedia information. However, their focus is on Portuguese only with a very simple NERC hierarchy.

The idea of using Wikipedia metadata to annotate multilingual corpora has also been proposed by Nothman et al.[15] for English, German, Spanish, Dutch, and Russian. Despite the multilingual approach, it also requires manually annotated text.

3 UNER Dataframe Description

As mentioned in the previous section, the UNER hierarchy was introduced by Alves et al.[1]. It was built upon the 4-level NERC hierarchy proposed by Sekine[17], which was chosen as it presents a high conceptual hierarchy. The changes comparing both structures have been detailed by the authors. The proposed UNER hierarchy is also composed of 4 levels. Level 0 being the root node from which derives all the other levels. Level 1 consists of three main classes: *Name*, *Time Expression* and *Numerical Expression*. Level 2 is composed of 29 named-entity categories which can be detailed in a third level with 95 types. Additionally, level 4 contains 129 subtypes (Alves et al.[1]).

This first version of the UNER hierarchy, therefore, encompasses 215 labels which can contain up to four levels of granularity depending on how detailed is the named-entity type. UNER labels are composed of tags from each level separated by a hyphen "-". As level 0 is the root and common for all entities, it is not present in the label. For example:

- UNER label *Name-Event-Natural-Phenomenon-Earthquake* is composed of level 1 *Name*, level 2 *Event*, level 3 *Natural Phenomenon* and level 4 *Earthquake*.

The idea of using both Wikipedia data and metadata associated with DBpedia information to generate UNER annotated datasets compelled us to revise

the first proposed UNER hierarchy. The main reason is that the automatic annotation process is based on a list of equivalences between UNER labels and DBpedia classes.

By generating the list of equivalences, it was noticeable that not all UNER labels would have a DBpedia equivalent class. This is case for majority of Time and Numerical expressions. These cases will have to be dealt with by other automatic methods in future work.

Therefore, for this article, we consider version 2 of UNER, presented in the GitHub webpage of the project⁴. It is composed of 124 labels and its hierarchy is detailed in Table 1.

Table 1. Description of the number of nodes per level inside UNER v2 hierarchy.

Level	Number of nodes
0 (root)	1
1	3
2	14
3	53
4	88

In the annotation process, we have decided to use the IOB format[16] as it is widely used by many NERC systems as showed by Alves et al.[1]. Therefore, each annotated entity token also receives in the beginning of the UNER label the letter “B” if the token is the first of the entity or “I” if inside. Non-entity tokens receive only the tag “O”.

4 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 subtypes (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, implement the IOB format [16] and gather statistical information concerning the generated corpus.

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

4.1 Process description

1. **UNER/DBpedia Mapping:** This is a mapper that connects each pertinent DBpedia class with a single UNER tag. It was created by the members

⁴ <https://github.com/cleopatra-itn/MIDAS/>

of the project, analysing each DBpedia class and associating it to the most pertinent 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 examples of classes and 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 Ontology⁵, where classes are presented in a structural order which allowed us to define the hierichal levels.

4.2 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 website⁶. Next, we perform text extraction preserving the hyperlinks in the article using WikiExtractor⁷. 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.

⁵ <http://mappings.dbpedia.org/server/ontology/classes/>

⁶ <https://dumps.wikimedia.org/>

⁷ <https://github.com/attardi/wikiextractor>

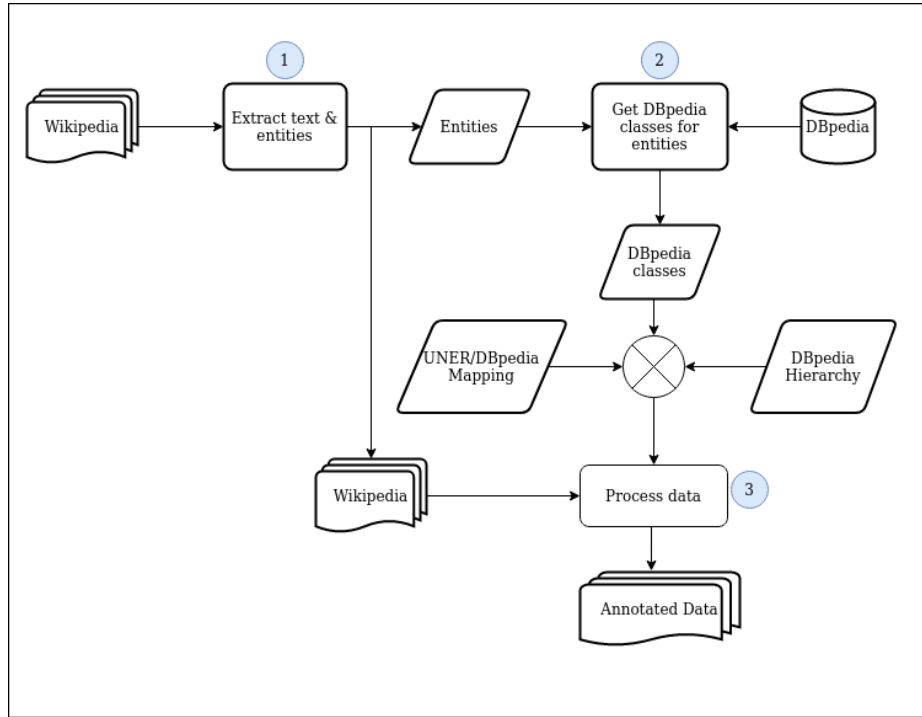


Fig. 1. Main process steps for Wikipedia data extraction and DBpedia/UNER annotations. Squares represent data and diamonds represent processing steps.

2. **Wikipedia-DBpedia entity linking:** For all the unique named-entities from the dumps, we query the DBpedia endpoint using a SPARQL query with SPARQLWrapper⁸ 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.
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 a 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, a 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.

⁸ <https://rdflib.dev/sparqlwrapper/>

4.3 Post-processing steps

The post-processing steps correspond to three different 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[16] 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.

The whole process and post-processing steps were applied to English language, generating the UNER English corpus which is described and evaluated in the following section. This baseline corpus is the base for the improvement experiments presented in later sections.

4.4 UNER English Corpus (Baseline)

General Information The English Wikipedia [20] is composed of 6,188,204 articles. After applying the main process of the proposed workflow, we obtained annotated text files divided into folders. The size of the English UNER corpus is presented in the following table.

Table 2. Corpus Size.

Corpus	Size	Folders	Files
English UNER	3.3 GB	172	17,150

Statistical information concerning the corpus is obtained by applying the post-processing steps previously described. Table 3 presents the main statistics about number of tokens and entities. Inside UNER English Corpus, 8.9% of tokens are entities.

As presented in section 3, the UNER hierarchy used for annotating the English Wikipedia texts is composed of 124 different multi-leveled labels with equivalences to DBpedia classes. However, baseline UNER English corpus contains 99 different UNER tags (80%).

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*

Table 3. Corpora Annotation Statistics.

	English UNER Corpus
Total Number of Tokens	325,395,838
Number of Non-Entity Tokens	320,719,350
Number of Entity Tokens	31,676,488
Number of Entities	15,101,318
Number of Different Entities	630,519

- *Location*: correspond to all subtypes inside the UNER types *Name-Location*.
- *Organization*: correspond to all subtypes inside the UNER type *Name-Organization*

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

Table 4. Corpora Annotation Statistics in terms of number of occurrences of the most used NERC classes (and % of all entities occurrences).

	English UNER Corpus
Person	4,200,313 (27.8%)
Location	2,613,248 (17.3%)
Organization	3,489,813 (23.1%)

These main classes correspond to 68.2% of NEs in the generated corpus.

Qualitative evaluation The proposed process requires the identification of the DBpedia classes associated with the respective tokens (via hyperlink) and the translation to UNER using (UNER/DBpedia equivalences).

An analysis of 943 entities randomly selected from UNER English Corpus has been performed to evaluate this step of the workflow. For each one, we have checked the DBpedia associated classes and the final UNER chosen tag. Table 5 presents the results of this evaluation.

Table 5. Evaluation of the annotation step: DBpedia class extraction and translation to the 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%

In the selected sample, 91% of the entities are correctly tagged with UNER tags. Nevertheless, 6% are associated with the correct UNER type but to a

generic subtype. For example, *Bengkulu* should be tagged as *Name-Location-GPE-City* but received the tag *Name-Location-GPE-GPE.Other*. Errors may come from mistakes in the DBpedia classes associated with the tokens or due to the prioritization rules and equivalences defined between DBpedia and UNER:

- *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*.

UNER English Golden dataset Beside the statistical information presented above, a sample from the generated corpus has been selected and corrected using WebAnno[4] by one annotator. The sample corresponds to one entire file from the output folder and contains 519 sentences and 105 different UNER labels (out of 124 from the list of UNER-DBpedia equivalences). The annotations were done by a non-native English speaker who is a member of the project. He followed objective guidelines, and for some specific entities, research using Wikipedia were done. In cases of multi-possible assignments, a final choice was done by the annotator so that each entity would have only one label in the golden set.

Table 6 presents the evaluation results of the baseline annotations of the file used to create the Golden dataset in terms of Precision, Recall, and F1-measure, considering the mean value of all 105 labels for each metric.

Table 6. Precision, Recall, and F1 Measure values of UNER EN dataset considering 519 manually annotated sentences.

Experiment	Precision	Recall	F1-measure
Baseline	61.9	27.2	37.8

As explained previously, 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. This may be one of the main reasons of the low value obtained for recall.

5 Dataset Improvement

Evaluation of the baseline annotated file using the Golden UNER English Corpus shows that the automatic annotation workflow has room for improvement, especially in terms of reducing the number of false negatives. Strategies for completing the annotation using dictionaries and knowledge graph were applied to the English Corpus. The ensemble of experiments and the evaluation is presented in the subsections below.

5.1 Experiment Design

Seven different experiments were conducted:

1. Global Dictionary: From the whole UNER English Corpus, we have established a dictionary of entities and the respective UNER label. As the same entity may appear in the corpus with different UNER tags (due to the associated DBpedia classes), we have selected for each entity the label with the highest number of occurrences. This dictionary is then used to complete the annotations of the corpus. Only entities with length longer than 2 characters were considered and numerical entities were excluded from the dictionary. Final size of the global dictionary is of 826,371 entities.
2. Global Dictionary only with multiple token entities: Similar to the previous experiment but in this case only entities with more than one token were considered. In total, the global dictionary is composed of 665,081 multi-token entities.
3. Local Dictionaries: In this setup we processed every Wikipedia dump file as a single article. Every entity in the article that is linked to UNER is cached into a local lookup dictionary with its text as the key and UNER class as the value. For every subsequent occurrence of the text in the given article we annotated the text with corresponding UNER class. We performed this step with the speculation that entities are more likely to appear within a single article than in a completely unrelated article. For example, *Barrack Obama* as person is more likely to appear in an article describing him as president than as a fictional character which appears in fictional content about him.
4. Global OEKG Dictionary: Open Event Knowledge Graph (OEKG)⁹ is a multilingual event-centric resource. Its instances have specific DBpedia classes, therefore, we intersected all the entries from the global dictionary with elements from OEKG. For each entity, its associated DBpedia class from OEKG was then mapped to UNER. The global OEKG dictionary contains 128,813 entries.
5. Global OEKG Dictionary only with multiple tokens entities: Similar to experiment 4 only in this case only entities with more than one token were considered (110,226 entities in total).
6. Local Dictionaries followed by Global OEKG Dictionary: Combination of experiment 3 with completion of annotations using dictionary established for experiment 4.
7. Local Dictionaries followed by OEKG Dictionary only with multiple tokens entities: Corpus from experiment 3 is completed using dictionary from experiment 5.

In all experiments, dictionaries were ordered from the longest entities to the shortest ones to guarantee that preferably multi-token entities were annotated and not mono-token ones.

⁹ <http://cleopatra-project.eu/index.php/open-event-knowledge-graph/>

5.2 Evaluation

The evaluation was conducted using the Golden Corpus presented previously. The baseline is the correspondent file with automatic annotations as result of the workflow described in section 4.

Golden Corpus has 105 different UNER labels, however, the baseline annotated file has only 62. For each possible label, we calculated precision, recall, and F1-measure. The IOB format[16] was applied, therefore, each UNER label can start either with "B" or "I", and non-entity tokens were tagged with "O".

From the 62 labels of the baseline, only 45 presented results different than 0. Therefore, the values present in the following table consider only these tags and represent the mean value of all the tags taken into account. Table 7 presents the metrics obtained for the baseline and each one of the experiments described in the previous sub-section.

Table 7. Precision, Recall, and F1 Measure values of experiments for improving UNER EN dataset.

Experiment	Precision	Recall	F1-measure
Baseline	72.9	32.0	39.2
1	32.1	35.7	27.0
2	47.5	34.8	34.0
3	73.6	29.6	36.8
4	71.1	33.9	40.8
5	73.0	33.4	40.5
6	72.1	32.1	39.1
7	74.0	31.5	38.7

Using the global dictionary (experiment 1) provides the highest value of recall (+3.7 compared to the baseline) but precision is considerably lower (-40.8). Similar situation when the global dictionary is used only with multi-token entities (experiment 2). Other experiments do not decrease precision so drastically and in some cases this metric is even increased. Recall is increased, compared to the baseline, for all experiments except for 3 and 6, 7. The usage of local dictionaries was not an effective solution for improving this evaluation metric.

Best option, considering F1 measure, is the usage of the dictionary verified with OEKG (experiment 4). Precision is slightly lower than the baseline (-1.8) while recall and F1 measure are higher (+1.9 and +1.6 respectively). If we consider only level 3 of UNER hierarchy, the possible tags are: *Disease, Event, Facility, Location, Natural Object, Organization, Person* and *Product*

The evaluation of each experiment considering only this upper level of the UNER hierarchy is presented in table 8. The IOB format was also considered, therefore, UNER labels could be preceded by either "B" or "I" and non-entity tokens were tagged with "O".

Table 8. Precision, Recall, and F1 Measure values of experiments for improving UNER EN dataset considering only level 3 of UNER hierarchy.

Experiment	Precision	Recall	F1-measure
Baseline	76.9	25.1	34.0
1	25.9	31.0	25.2
2	37.2	27.8	31.1
3	76.8	24.5	33.2
4	74.6	27.3	36.0
5	76.6	26.3	35.3
6	74.6	26.9	35.5
7	76.6	25.8	34.7

In this scenario, the highest precision is from the baseline. The Best recall is obtained when the global dictionary is used (experiment 1) but, as it was observed before, in this case precision is heavily impacted compared to the baseline (-51.0). Experiment 4 is the one with the highest F1 measure, same for the previous evaluation where all UNER levels were considered.

6 Discussion

In section 4 we presented the whole process for generating UNER annotated corpora and the application of this workflow to create a dataset for the English language. The evaluation of the entity’s extraction and translation to UNER (table 5) showed that 85% of analysed entities were correctly annotated. However, errors are introduced to the dataset mainly because of the wrong association of the tokens inside Wikipedia text and DBpedia classes. Furthermore, the selection rule in our process (choosing the DBpedia class having the highest granularity inside DBpedia hierarchy) is also a source of mistakes.

The corpus generated by the proposed workflow was also evaluated using the manually annotated Golden dataset. It is noticeable that even with the reduced version of the UNER hierarchy (UNER v2, 124 labels with DBpedia equivalences) used for this task, the whole dataset was only annotated with 99 different labels, from which only 62 presented values of precision and recall different from zero in the baseline sample file when compared to the Golden. Therefore, further changes in the UNER hierarchy should be implemented, for example, Time and Numerical Expressions (already reduced in UNER v2) should be excluded in this step of the annotation and some detailed labels concerning Events, Organizations and Locations can be joined in more generic categories.

As expected, recall is much inferior to precision in all conducted evaluations. This is due to the fact that inside Wikipedia not all entities are linked to DBpedia. This problem was also encountered by Weber & Vieira[13] who used a similar workflow for Portuguese named-entity task with a much simpler hierarchy. The authors did not conduct any intrinsic evaluation of the generated

corpus but, instead, used it to train models and evaluate results with existing Golden sets.

Concerning the improvement experiments, the best-identified option was to use a dictionary fine-tuned from the Open Event Knowledge Graph. This graph allows to identify a more precise specific DBpedia class and therefore helps in improving recall without considerable loss in precision. In this article, we presented the use of this method as a post-processing step to complete the initial annotations of the baseline. However, the usage of OEKG information can be also implemented inside the workflow as a way of improving overall precision.

7 Conclusions and Future Directions

In this paper, we described an automatic workflow for generating multilingual Named-Entity recognition corpora by using Wikipedia and DBpedia data and following the UNER hierarchy. The whole process is available and can be applied to any language having Wikipedia and DBpedia. We also presented the application of the extraction and annotation method used to generate UNER English corpus. The generated dataset has been described and evaluated with a manually annotated Golden set.

Furthermore, an ensemble of experiments were conducted to improve final annotated dataset. We have identified that the best results were obtained by using a dictionary of entities with verification of the associated DBpedia class using Open Event Knowledge Graph: 76.9 for precision, 31.0 for recall and 36.0 for F1-measure. Nevertheless, there is still room for improvement in both recall and F-measure.

In our future work, we plan to continue exploring OEKG by integrating it to extraction and annotation the workflow and not only as a post-processing step. Also, we intend to extend our corpus to other languages, especially under-resourced ones, while evaluating our workflow’s performance across the languages. Moreover, to complete this intrinsic evaluation of our dataset, we plan to evaluate it extrinsically by using the generated datasets to train the machine and deep learning models.

8 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. Alves, D., Thakkar, G., Tadić, M.: Evaluating language tools for fifteen EU-official under-resourced languages. In: Proceedings of the 12th Language Resources and Evaluation Conference. pp. 1866–1873. European Language Resources Association, Marseille, France (May 2020), <https://www.aclweb.org/anthology/2020.lrec-1.230>
 3. 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>
 4. Eckart de Castilho, R., Mújdricza-Maydt, É., Yimam, S.M., Hartmann, S., Gurevych, I., Frank, A., Biemann, C.: A web-based tool for the integrated annotation of semantic and syntactic structures. In: Proceedings of the Workshop on Language Technology Resources and Tools for Digital Humanities (LT4DH). pp. 76–84. The COLING 2016 Organizing Committee, Osaka, Japan (Dec 2016), <https://www.aclweb.org/anthology/W16-4011>
 5. 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>
 6. Dumitrescu, S.D., Avram, A.: Introducing RONEC - the romanian named entity corpus. CoRR **abs/1909.01247** (2019), <http://arxiv.org/abs/1909.01247>
 7. Freitas, C., Carvalho, P., Gonçalo Oliveira, H., Mota, C., Santos, D.: Second HAREM: advancing the state of the art of named entity recognition in Portuguese. In: In Nicoletta Calzolari; Khalid Choukri; Bente Maegaard; Joseph Mariani; Jan Odijk; Stelios Piperidis; Mike Rosner; Daniel Tapias (ed) Proceedings of the International Conference on Language Resources and Evaluation (LREC 2010)(Valletta 17-23 May de 2010) European Language Resources Association. European Language Resources Association (2010)
 8. Honnibal, M., Montani, I., Van Landeghem, S., Boyd, A.: spaCy: Industrial-strength Natural Language Processing in Python (2020). <https://doi.org/10.5281/zenodo.1212303>, <https://doi.org/10.5281/zenodo.1212303>
 9. 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>
 10. 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>
 11. 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/>
 12. 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>

13. 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>
14. 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>
15. 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>
16. 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>
17. Sekine, S.: The Definition of Sekine’s Extended Named Entities. https://nlp.cs.nyu.edu/ene/version7_1_0Beng.html (07 2007), (Accessed on 28/02/2020)
18. Ševčíková, M., Žabokrtský, Z., Krůza, O.: Named entities in Czech: annotating data and developing NE tagger. In: International Conference on Text, Speech and Dialogue. pp. 188–195. Springer (2007)
19. Weischedel, R., Palmer, M., Marcus, M., Hovy, E., Pradhan, S., Ramshaw, L., Xue, N., Taylor, A., Kaufman, J., Franchini, M., El-Bachouti, M., Belvin, R., Houston, A.: OntoNotes Release 5.0 (2013). <https://doi.org/11272.1/AB2/MKJJ2R>, <https://hdl.handle.net/11272.1/AB2/MKJJ2R>
20. 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]
21. 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>