

IXA/Cogcomp at SemEval-2023 Task 2: Context-enriched Multilingual Named Entity Recognition using Knowledge Bases

Iker García-Ferrero¹, Jon Ander Campos¹, Oscar Sainz¹,
Ander Salaberria¹, Dan Roth²

¹HiTZ Center - Ixa, University of the Basque Country UPV/EHU
{iker.garciaf,jonander.campos,oscar.sainz,ander.salaberria}@ehu.eus

²University of Pennsylvania
danroth@seas.upenn.edu

Abstract

Named Entity Recognition (NER) is a core natural language processing task in which pre-trained language models have shown remarkable performance. However, standard benchmarks like CoNLL 2003 (Tjong Kim Sang and De Meulder, 2003) do not address many of the challenges that deployed NER systems face, such as having to classify emerging or complex entities in a fine-grained way. In this paper we present a novel NER cascade approach comprising three steps: first, identifying candidate entities in the input sentence; second, linking the each candidate to an existing knowledge base; third, predicting the fine-grained category for each entity candidate. We empirically demonstrate the significance of external knowledge bases in accurately classifying fine-grained and emerging entities. Our system exhibits robust performance in the MultiCoNER2 (Fetahu et al., 2023b) shared task, even in the low-resource language setting where we leverage knowledge bases of high-resource languages.

1 Introduction

The research on the Named Entity Recognition field has mainly focused on recognizing entities from news corpora. Datasets like CoNLL 2003 (Tjong Kim Sang and De Meulder, 2003) or OntoNotes (Hovy et al., 2006) have been the standard benchmarks for a long time. It is not difficult to find systems that perform over 92% of F1 score on those datasets. However, it is known that actual models perform pretty well on relatively easy entity types (such as person names) but have been struggling to extract more complex entities (Luken et al., 2018; Hanselowski et al., 2018). Moreover, the evaluation setup usually involves a large overlap on the entities from the train and test sets, overestimating the systems performance. Removing such overlap resulted on a large performance drop,

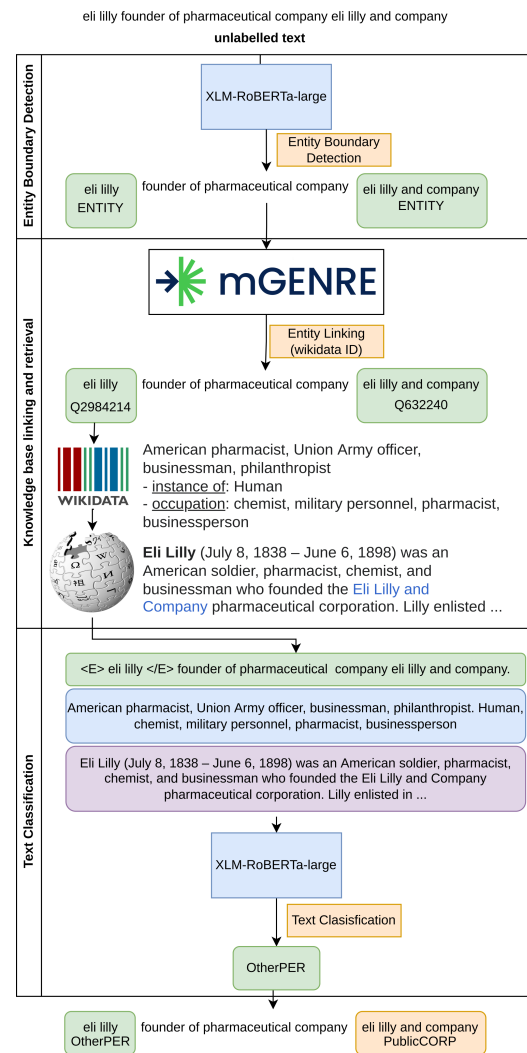


Figure 1: Illustration of our fine-grained named entity recognition system.

consequently showing that models tend to memorize rather than generalize to complex or unseen entities (Meng et al., 2021; Fetahu et al., 2022).

Motivated by these insights, recent works have developed new and more complex datasets like MultiCoNER (Malmasi et al., 2022a) or UFET (Choi et al., 2018) that involve complex entity mentions with higher granularity on the type

definition. To address these new challenges large language models (LLMs) are commonly used, with the hope that after being trained on billions of tokens the models become aware of the relevant entities. Despite the capacity of the models to memorize information about each individual entity, they face a harsh reality: the knowledge that a model can memorize is limited by the date on which the model was pre-trained. So, new emerging entities can be difficult to categorize for models trained years ago. Moreover, there are cases when it is not possible to identify the fine-grained category of an entity just from the context, and thus requires prior knowledge about the entity. For example: *George Bernard Shaw, Douglas Fairbanks, Mary Pickford and Margaret Thatcher are some of the famous guests who stayed at the hotel.*. From the previous sentence, it is impossible to categorize *George Bernard Shaw* and *Margaret Thatcher* as politicians neither *Douglas Fairbanks* and *Mary Pickford* as artist. We hypothesize that a possible solution for this problem is to allow the models access updated Knowledge Bases, and use the updated information to infer the correct type of the entity of interest.

To overcome the mentioned challenges, in this work we present a NER approach (see Figure 1) that (1) identifies possible entity candidates by analyzing the input sentence structure, (2) links the candidate to an existing updated knowledge base if possible, and (3) performs the fine-grained classification using the input sentence plus the retrieved information from the KB about the entity. The approach allowed us to perform fine-grained NER with updated world knowledge when possible, and standard NER otherwise. Our code is publicly available to facilitate the reproducibility of the results and its use in future research¹.

2 Related Work

Named Entity Recognition (NER) (Grishman and Sundheim, 1996) is a core natural language processing task where the goal is to identify entities belonging to predefined semantic types such as organizations, products and locations. Since its inception, different approaches have been developed, from statistical machine learning methods (Zhou and Su, 2002; Konkol and Konopik, 2013; Agerri et al., 2014) to the ones based on neural networks

(Strubell et al., 2017; Xia et al., 2019) and the mix of both (Huang et al., 2015; Chen et al., 2017). Recently, the use of richer contextual embeddings computed via Transformer models (Vaswani et al., 2017; Devlin et al., 2018; He et al., 2020) have considerably improved the state-of-the-art of the task. Nevertheless, these models still have problems detecting and labelling unseen or complex entities (Augenstein et al., 2017; Meng et al., 2021).

Many datasets (Malmasi et al., 2022a; Fetahu et al., 2023a) and methods (Choi et al., 2018; Zhou et al., 2018; Mengge et al., 2020) have been developed as a result of the difficulty to label complex entities. The MultiCoNER (Malmasi et al., 2022a) shared task focuses on detecting semantically ambiguous and complex entities in low-context settings. Almost every participant of the task used Transformer based models as their backbone model, XLM-RoBERTa (Conneau et al., 2020a) being the most popular. The best two models presented for the task (Wang et al., 2022; Chen et al., 2022) show that the use of external KBs, such as Wikipedia and Wikidata, improve their results significantly. Meanwhile, other best performing models (Gan et al., 2022; Pu et al., 2022; Pais, 2022) take advantage of different data augmentation methods without relying on external KBs. The gap between the performance of the first two methods was less noticeable than the gap between the second best and the rest of approaches, implying that exploiting external KBs during training and inference is a promising research line. This difference was greater when more complex entities have to be identified. However, the overall results showed that recognizing complex or infrequent entities is still difficult (Malmasi et al., 2022b).

In this work we focus on the MultiCoNER2 (Fetahu et al., 2023a) dataset, a second version of MultiCoNER (Malmasi et al., 2022b) with higher granularity than its predecessor. Comparing to the state-of-the-art of the first MultiCoNER task, our approach for this task also uses XLM-RoBERTa for entity boundary detection and entity classification, but differs on how we retrieve information from the external KBs. Instead of retrieving relevant text passages to classify a given entity from a query sentence (Wang et al., 2022), or building a Gazetteer network from Wikidata entities and their derived labels to further enrich token representations (Chen et al., 2022), we directly predict Wikidata IDs from entity candidates by integrating mGENRE (De Cao

¹<https://github.com/ikergarcia1996/Context-enriched-NER>

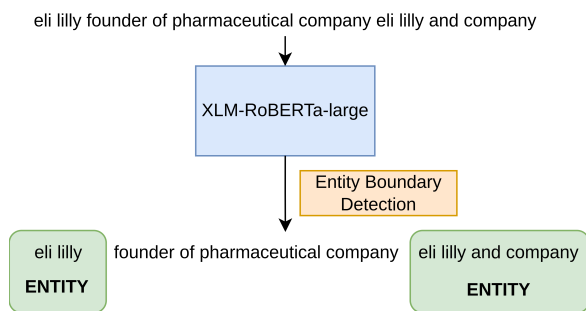


Figure 2: Entity Boundary Detection. We use a XLM-RoBERTa model to predict named entity boundaries.

et al., 2022a) as our entity linking module. Previous models have already tried to link entities with external KBs (Tsai et al., 2016; Zhou et al., 2018; Wang et al., 2021). However, mGENRE allows us to easily access both Wikidata entity descriptors and Wikipedia articles to increase our candidate’s context for entity classification.

3 System description

Our method implements three steps. First we identify entity candidates in the input sentence. Then, we link the candidate to an existing knowledge base. Finally, we predict the fine-grained category for each entity candidate.

3.1 Entity Boundary Detection

Given unlabelled text as input, we predict named entity boundaries by analyzing the input sentence structure (Figure 2). We treat this task as a sequence labelling task in which the model predicts if a given token is part of an entity or not ("B-ENTITY", "I-ENTITY", "O"). We use the multilingual XLM-RoBERTa-large model (Conneau et al., 2020b) with a token classification layer (a linear layer) on top of each token representation. Our implementation is based on the sequence labelling implementation of the Huggingface open-source library (Wolf et al., 2019). We evaluate the model in the development set at the end of each epoch and then select the best performing checkpoint. We train five independent models and then use majority vote as the ensembling strategy at inference time.

3.2 Entity Linking and Information Retrieval

Given a sentence and the boundaries of a named entity between the sentence, we aim to link the entity mention to its corresponding Wikidata/Wikipedia page (Figure 3). To accomplish this task, we employ the mGENRE entity linking system (De Cao

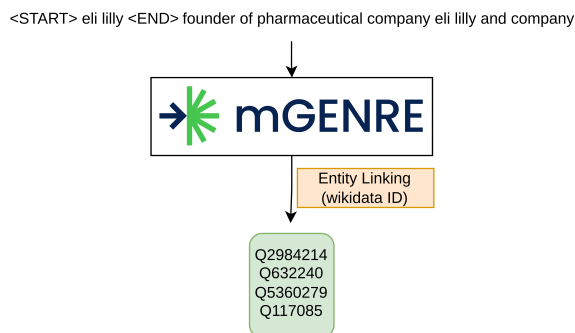


Figure 3: Entity Linking step. Given a sentence and the boundaries of a named entity, we use mGENRE to predict the Wikidata ID that corresponds to the entity

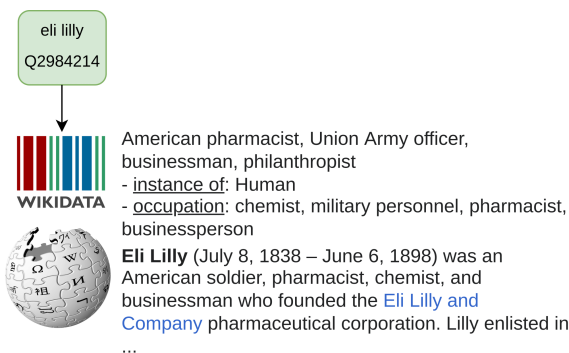


Figure 4: Information Retrieval step. Given a wikidata ID, we retrieve the relevant information from the Wikidata and Wikipedia that can help us determining the correct category of the entity.

et al., 2021; De Cao et al., 2022b). mGENRE is a sequence-to-sequence system for Multilingual Entity Linking, which can generate entity names in over 100 languages from left to right, token-by-token in an autoregressive manner, conditioned by the context. In order to ensure that only valid entity identifiers are generated mGENRE employs a prefix tree (trie) to enable constrained beam search. mGENRE predicts both, Wikipedia page title that corresponds to the entity and the language in which the entity mention is written, and subsequently both identifiers to the Wikidata ID. In our work, we utilize the pre-trained mGENRE model provided by its authors. The authors fine-tuned an mBART (Liu et al., 2020; Lewis et al., 2020) model that had been pre-trained on 125 languages using Wikipedia hyperlinks in 105 languages. For each entity, we use constrained beam search to predict the five most probable Wikidata IDs for each predicted entity from the Entity Boundary Detection step.

As illustrated in Figure 4, we leverage the predicted Wikidata ID to obtain the Wikidata descrip-

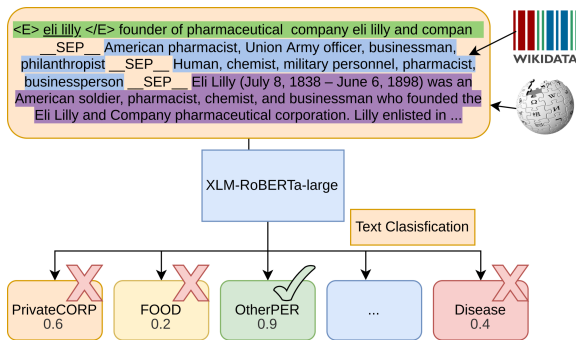


Figure 5: Given the boundaries of a named entity and the information retrieved from Wikipedia/Wikidata, we use a XLM-RoBERTa model to predict the correct category for the entity.

tion for the entity, as well as the contents of the *instance_off* and *occupation* arguments. In the multilingual track, we also retrieve the *subclass_off* argument. Due to time constraints, we were unable to retrieve this argument in the other tracks. As Wikidata pages contain links to their corresponding Wikipedia webpages, we also retrieve the Wikipedia summary. These summaries are generally more detailed than the Wikidata descriptions.

We initiate the retrieval process by querying the Wikidata ID with the highest probability, as predicted by mGENRE. If the predicted ID corresponds to a page that has been deleted, is empty, or is a list/disambiguation page, we discard the ID and proceed to the next most probable predicted ID. We found that, while Wikidata descriptions and arguments are mostly available in English for the entities in MultiCoNER, this is not the case for other languages such as Bangla or Farsi. Similarly, a large number of entities have a Wikipedia page available in English but not in other languages. For this reason, we always retrieve the Wikidata description and arguments, as well as the Wikipedia summary, in English.

3.3 Entity Category Classifier

In our final step, we aim to classify the entity candidates into fine-grained categories by combining all the knowledge we have. We create a text input by concatenating the sentence where the entity is annotated with an HTML-style markup, the wikidata description, a list of all retrieved arguments, and the Wikipedia summary. We use the special token `__SEP__` to delimit each piece of information. If we fail to retrieve a Wikidata or Wikipedia page, we use the text "No Wikidata/Wikipedia summary found," and the model predicts the category solely

based on the context of the entity.

We use the multilingual XLM-RoBERTa-large model (Conneau et al., 2020b) with a text classification layer (a linear layer) on top of the first token. We implement this using the Huggingface open-source library’s text classification implementation (Wolf et al., 2019). During training, we generate the train and development datasets using the gold entity boundaries. During inference, we predict the categories for the entity boundaries computed in the entity boundary detection step. We evaluate the model at the end of each epoch and choose the model with the highest performance on the development set. We train five independent models and merge the predictions at inference using majority voting.

4 Experimental Setup and Dataset

We use the official MultiCoNER2 dataset (Fetahu et al., 2023a) in all tracks to train our entity boundary and entity classification models. To query external knowledge, we utilize the MediaWiki API to retrieve data from the live version of Wikipedia/Wikidata. MultiCoNER2 consists of two substacks, the monolingual task with 12 monolingual datasets and a multilingual task that contains data from all languages. For 7 languages (EN, ZH, IT, ES, FR, PT, SV) some sentences were corrupted with noise, either on context tokens or entity tokens. The results of the shared task are evaluated based on entity-level macro F1 scores, where all label categories contribute equally to the final score, regardless of the number of entities labeled with each category. The dataset consists of 36 pre-defined fine-grained categories which are grouped into 6 coarse-grained categories: Medical entities, Locations, Creative Works, Groups, Persons and Products. A description of the hyperparameters used, dataset stats and hardware used is available in the Appendix.

5 Results

Our team participated in all the 12 monolingual tracks and the multilingual track. More than 35 team participated in the shared task. We achieved the **3rd** positions in four tracks. The shared task results computed by the organizers are shown in Table 1. We first compare our system with an XLM-RoBERTa-large (Conneau et al., 2020b) baseline. For this baseline, we added a token classification layer (linear layer) on top of each token represen-

Team Name	EN	ES	SV	UK	PT	FR	FA	DE	ZH	HI	BN	IT	MULTI	AVG
XLM-RoBERTa-large	59.85	62.60	65.76	65.67	61.09	62.98	57.17	62.86	51.89	68.06	65.97	65.17	69.12	62.93
NetEase	-	-	-	-	-	-	-	-	84.05	-	-	-	-	-
USTC-NELSLIP	72.15	74.44	75.47	74.37	71.26	74.25	68.85	78.71	66.57	82.14	80.59	75.70	75.62	74.72
PAI	80.00	71.67	72.38	71.28	81.61	86.17	68.46	88.09	74.87	80.96	84.39	84.88	77.00	78.59
DAMO-NLP	85.53	89.78	89.57	89.02	85.97	89.59	87.93	84.97	75.98	78.56	81.60	89.79	84.48	85.59
Our System	72.82	73.81	76.54	75.25	72.28	74.52	69.49	80.35	64.86	79.56	78.95	74.67	78.17	74.71

Table 1: Our system macro-F1 score for all the tracks compared with our baseline and the systems that achieved the best results.

tation of an XLM-RoBERTa-large model and fine-tuned the model to directly predict the fine-category for each token. The results were much lower than the ones achieved with our system, which confirms our hypothesis that external knowledge is needed to predict the entity categories.

We also compare our system with the performance of systems from others teams that got the first position in any of the tracks. Our system was most competitive in Hindi and Bangla, as well as the multilingual track. Hindi and Bangla are by far the languages with the smallest Wikipedia’s, ranked 59th and 63rd by the number of articles respectively.² We attribute this to our decision to retrieve the external knowledge from the English Wikipedia/Wikidata for all the tracks. Because we predict a wikidata ID, we are able to retrieve knowledge in any language in which it is available. This enables our system to be used for low-resource languages for which no large, up-to-date knowledge base is available, thanks to being able to link the entities found to a knowledge base in high-resource languages. On the other hand, our system is not as competitive in very high-resource languages, such as English and Chinese.

5.1 Clean vs Noisy Performance

In seven languages (EN, ZH, IT, ES, FR, PT, SV), some sentences contained noise, either on context tokens or entity tokens. Table 2 shows the results of our system on the set of sentences that were not corrupted and the set of corrupted sentences. We observed a significant drop in performance when noise was introduced in the sentences, especially for the Chinese track. We believe that this is because our entity linking step depends too much on the entity itself and does not sufficiently take into account the context in which the entity appears. Improving our system’s ability to consider contextual

²https://meta.wikimedia.org/wiki/List_of_Wikipedias

Track	Clean	Noisy
EN	76.64	64.36
ES	77.65	66.09
SV	80.75	68.69
PT	76.00	65.54
FR	78.60	65.81
ZH	70.35	48.37
IT	78.16	67.66

Table 2: Our system macro-F1 score in the clean and noisy data from the test sets as computed by the organizers

information is something we need to work on.

5.2 Ablation Study

In this section, we evaluate the different steps in our system and check how different design decisions affect performance.

How important is external knowledge to classify entities? We aim to measure the importance of retrieving external knowledge for accurate entity classification in the MultiCoNER dataset. To achieve this, we train our text classification system on the training split and evaluate its performance on the development data. We use the gold labeled entity spans for both splits. We assess the relevance of external knowledge by varying the amount of external knowledge used in the evaluation. Table 3 shows the results obtained using different combinations of external knowledge, including using only the entity in its context (no external knowledge). As we observed in Table 1, the results without external knowledge are very poor. The most relevant external knowledge is the content of the Wikidata arguments. However, as we provide more external knowledge to the model, the performance improves. These results also demonstrate that our system achieves remarkable results in classifying entities into fine-grained categories when the cor-

Initial Context	Wikidata Description	Wikidata Arguments	Wikipedia Summary	EN	ZH
✓				79	81
✓	✓			82	84
✓		✓		88	88
✓			✓	86	86
✓	✓	✓		89	86
✓	✓	✓	✓	90	89

Table 3: Our system macro-F1 score in the development split when using gold entity boundaries and different amount of external knowledge

	EN	ZH	Multi
Entity Boundary Detection	90.05	90.49	90.52

Table 4: F1-score of the Entity Boundary Classification Step in the development split

rect entity spans are provided, indicating that the entity boundary detection step hinders the performance of our system. The entity classification confusion matrix for the English development set is available in Appendix D.

Entity boundary detection performance Our initial hypothesis was that external knowledge is required to accurately classify entities into fine-grained categories. However, we also hypothesized that detecting the boundaries of named entities does not require external knowledge. For instance, it is possible to identify that a sequence of words corresponds to the name of a person without needing to know anything about that person. Table 4 shows the performance the entity boundary detection step in the development dataset, which does not utilize any external knowledge. We were able to identify a significant proportion of the entities in the dataset. However, the results were lower than those of the entity classification step in Table 3. This indicates that the entity boundary detection step is the weakest point in our system. We believe that incorporating external knowledge into this step would enhance the results.

6 Conclusion

We have developed a system that identifies potential entity candidates and leverages external knowledge from an up-to-date knowledge base to classify them into a set of predefined fine-grained categories, addressing the challenges posed by temporal knowledge and unknown entities. Our results demonstrate exceptional performance in classifying entities into fine-grained categories, un-

derscoring the need for external knowledge in accurately classifying entities in the MultiCoNER dataset. However, our current entity boundary detection step does not incorporate external knowledge, which we plan to improve in the future.

Furthermore, our system exhibits promising results for low-resource languages where a comprehensive Wikipedia is unavailable. By linking the identified entities to a knowledge base in a high-resource language, it can be used to process low-resource languages. In future work, we intend to integrate all the steps and insights gained from our current system into a single end-to-end model. Although each step performs well when evaluated independently, the current pipeline multiplies errors at each step, which we aim to address in the integrated model.

Acknowledgements

This work has been partially supported by the HiTZ center and the Basque Government (Research group funding IT-1805-22). We also acknowledge the funding from the following projects: (i) DeepKnowledge (PID2021-127777OB-C21) MCIN/AEI/10.13039/501100011033 and ERDF A way of making Europe; (ii) DeepR3 (TED2021-130295B-C31) by MCIN/AEI/10.13039/501100011033 and EU NextGeneration programme EU/PRTR. (iii) IARPA BETTER Program contract No. 2019-19051600006 (ODNI, IARPA).

Iker García-Ferrero, Oscar Sainz and Ander Salaberria are supported by doctoral grants from the Basque Government (PRE_2022_2_0208, RE_2022_2_0110 and PRE_2022_2_0219, respectively). Jon Ander Campos enjoys a doctoral grant from the Spanish MEC (FPU18/01271).

References

- Rodrigo Agerri, Josu Bermudez, and German Rigau. 2014. *IXA pipeline: Efficient and ready to use multilingual NLP tools*. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation, LREC 2014, Reykjavik, Iceland, May 26-31, 2014*, pages 3823–3828. European Language Resources Association (ELRA).
- Isabelle Augenstein, Leon Derczynski, and Kalina Bontcheva. 2017. Generalisation in named entity recognition: A quantitative analysis. *Computer Speech & Language*, 44:61–83.
- Beiduo Chen, Jun-Yu Ma, Jiajun Qi, Wu Guo, Zhen-Hua Ling, and Quan Liu. 2022. *USTC-NELSLIP*

- at SemEval-2022 task 11: Gazetteer-adapted integration network for multilingual complex named entity recognition. In *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*, pages 1613–1622, Seattle, United States. Association for Computational Linguistics.
- Tao Chen, Ruifeng Xu, Yulan He, and Xuan Wang. 2017. Improving sentiment analysis via sentence type classification using bilstm-crf and cnn. *Expert Systems with Applications*, 72:221–230.
- Eunsol Choi, Omer Levy, Yejin Choi, and Luke Zettlemoyer. 2018. **Ultra-fine entity typing**. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 87–96, Melbourne, Australia. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020a. **Unsupervised cross-lingual representation learning at scale**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020b. **Unsupervised cross-lingual representation learning at scale**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Nicola De Cao, Gautier Izacard, Sebastian Riedel, and Fabio Petroni. 2021. **Autoregressive entity retrieval**. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Nicola De Cao, Ledell Wu, Kashyap Papat, Mikel Artetxe, Naman Goyal, Mikhail Plekhanov, Luke Zettlemoyer, Nicola Cancedda, Sebastian Riedel, and Fabio Petroni. 2022a. **Multilingual autoregressive entity linking**. *Transactions of the Association for Computational Linguistics*, 10:274–290.
- Nicola De Cao, Ledell Wu, Kashyap Papat, Mikel Artetxe, Naman Goyal, Mikhail Plekhanov, Luke Zettlemoyer, Nicola Cancedda, Sebastian Riedel, and Fabio Petroni. 2022b. **Multilingual autoregressive entity linking**. *Transactions of the Association for Computational Linguistics*, 10:274–290.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Besnik Fetahu, Zhiyu Chen, Sudipta Kar, Oleg Rokhlenko, and Shervin Malmasi. 2023a. **Multi-CoNER v2: a Large Multilingual dataset for Fine-grained and Noisy Named Entity Recognition**.
- Besnik Fetahu, Anjie Fang, Oleg Rokhlenko, and Shervin Malmasi. 2022. **Dynamic gazetteer integration in multilingual models for cross-lingual and cross-domain named entity recognition**. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2777–2790, Seattle, United States. Association for Computational Linguistics.
- Besnik Fetahu, Sudipta Kar, Zhiyu Chen, Oleg Rokhlenko, and Shervin Malmasi. 2023b. **SemEval-2023 Task 2: Fine-grained Multilingual Named Entity Recognition (MultiCoNER 2)**. In *Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023)*. Association for Computational Linguistics.
- Weichao Gan, Yuanping Lin, Guangbo Yu, Guimin Chen, and Qian Ye. 2022. **Qtrade AI at SemEval-2022 task 11: An unified framework for multilingual NER task**. In *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*, pages 1654–1664, Seattle, United States. Association for Computational Linguistics.
- Ralph Grishman and Beth Sundheim. 1996. **Message understanding conference-6: A brief history**. In *Proceedings of the 16th Conference on Computational Linguistics - Volume 1, COLING '96*, page 466–471, USA. Association for Computational Linguistics.
- Andreas Hanselowski, Hao Zhang, Zile Li, Daniil Sorokin, Benjamin Schiller, Claudia Schulz, and Iryna Gurevych. 2018. **UKP-athene: Multi-sentence textual entailment for claim verification**. In *Proceedings of the First Workshop on Fact Extraction and VERification (FEVER)*, pages 103–108, Brussels, Belgium. Association for Computational Linguistics.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2020. **Deberta: Decoding-enhanced bert with disentangled attention**. *arXiv preprint arXiv:2006.03654*.
- Eduard Hovy, Mitchell Marcus, Martha Palmer, Lance Ramshaw, and Ralph Weischedel. 2006. **OntoNotes: The 90% solution**. In *Proceedings of the Human Language Technology Conference of the NAACL, Companion Volume: Short Papers*, pages 57–60, New York City, USA. Association for Computational Linguistics.
- Zhiheng Huang, Wei Xu, and Kai Yu. 2015. **Bidirectional lstm-crf models for sequence tagging**. *arXiv preprint arXiv:1508.01991*.
- Michal Konkol and Miloslav Konopík. 2013. **Crf-based czech named entity recognizer and consolidation of czech ner research**. In *Text, Speech, and Dialogue*:

- 16th International Conference, TSD 2013, Pilsen, Czech Republic, September 1-5, 2013. *Proceedings 16*, pages 153–160. Springer.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. **BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020. **Multilingual denoising pre-training for neural machine translation**. *CoRR*, abs/2001.08210.
- Ilya Loshchilov and Frank Hutter. 2019. **Decoupled weight decay regularization**. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Jackson Luken, Nanjiang Jiang, and Marie-Catherine de Marneffe. 2018. **QED: A fact verification system for the FEVER shared task**. In *Proceedings of the First Workshop on Fact Extraction and VERification (FEVER)*, pages 156–160, Brussels, Belgium. Association for Computational Linguistics.
- Shervin Malmasi, Anjie Fang, Besnik Fetahu, Sudipta Kar, and Oleg Rokhlenko. 2022a. **MultiCoNER: A large-scale multilingual dataset for complex named entity recognition**. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 3798–3809, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Shervin Malmasi, Anjie Fang, Besnik Fetahu, Sudipta Kar, and Oleg Rokhlenko. 2022b. **SemEval-2022 Task 11: Multilingual Complex Named Entity Recognition (MultiCoNER)**. In *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*. Association for Computational Linguistics.
- Tao Meng, Anjie Fang, Oleg Rokhlenko, and Shervin Malmasi. 2021. **GEMNET: Effective gated gazetteer representations for recognizing complex entities in low-context input**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1499–1512, Online. Association for Computational Linguistics.
- Xue Mengge, Bowen Yu, Zhenyu Zhang, Tingwen Liu, Yue Zhang, and Bin Wang. 2020. **Coarse-to-Fine Pre-training for Named Entity Recognition**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6345–6354, Online. Association for Computational Linguistics.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. **fairseq: A fast, extensible toolkit for sequence modeling**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Demonstrations*, pages 48–53. Association for Computational Linguistics.
- Vasile Pais. 2022. **RACAI at SemEval-2022 task 11: Complex named entity recognition using a lateral inhibition mechanism**. In *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*, pages 1562–1569, Seattle, United States. Association for Computational Linguistics.
- Keyu Pu, Hongyi Liu, Yixiao Yang, Jiangzhou Ji, Wenyi Lv, and Yaohan He. 2022. **CMB AI lab at SemEval-2022 task 11: A two-stage approach for complex named entity recognition via span boundary detection and span classification**. In *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*, pages 1603–1607, Seattle, United States. Association for Computational Linguistics.
- Emma Strubell, Patrick Verga, David Belanger, and Andrew McCallum. 2017. **Fast and accurate entity recognition with iterated dilated convolutions**. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2670–2680, Copenhagen, Denmark. Association for Computational Linguistics.
- Erik F. Tjong Kim Sang and Fien De Meulder. 2003. **Introduction to the conll-2003 shared task: Language-independent named entity recognition**. In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003 - Volume 4, CONLL '03*, page 142–147, USA. Association for Computational Linguistics.
- Chen-Tse Tsai, Stephen Mayhew, and Dan Roth. 2016. **Cross-lingual named entity recognition via wikification**. In *Proceedings of the 20th SIGNLL Conference on Computational Natural Language Learning*, pages 219–228, Berlin, Germany. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. **Attention is all you need**. *Advances in neural information processing systems*, 30.
- Xinyu Wang, Yong Jiang, Nguyen Bach, Tao Wang, Zhongqiang Huang, Fei Huang, and Kewei Tu. 2021. **Improving named entity recognition by external context retrieving and cooperative learning**. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1800–1812.
- Xinyu Wang, Yongliang Shen, Jiong Cai, Tao Wang, Xiaobin Wang, Pengjun Xie, Fei Huang, Weiming Lu,

Yueting Zhuang, Kewei Tu, Wei Lu, and Yong Jiang. 2022. [DAMO-NLP at SemEval-2022 task 11: A knowledge-based system for multilingual named entity recognition](#). In *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*, pages 1457–1468, Seattle, United States. Association for Computational Linguistics.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, and Jamie Brew. 2019. [Huggingface’s transformers: State-of-the-art natural language processing](#). *CoRR*, abs/1910.03771.

Congying Xia, Chenwei Zhang, Tao Yang, Yaliang Li, Nan Du, Xian Wu, Wei Fan, Fenglong Ma, and Philip Yu. 2019. [Multi-grained named entity recognition](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1430–1440, Florence, Italy. Association for Computational Linguistics.

Ben Zhou, Daniel Khashabi, Chen-Tse Tsai, and Dan Roth. 2018. [Zero-shot open entity typing as type-compatible grounding](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2065–2076, Brussels, Belgium. Association for Computational Linguistics.

GuoDong Zhou and Jian Su. 2002. Named entity recognition using an hmm-based chunk tagger. In *Proceedings of the 40th annual meeting of the association for computational linguistics*, pages 473–480.

A Dataset details

We use the MultiCoNER2 dataset in all our experiments (Fetahu et al., 2023a). The number of sentences for each language and dataset split is available in Table 5. The dataset Taxonomy is visualized in Figure 6. The label distribution is visualized in Figure 7.

B Hyper parameters

B.1 Entity Boundary Detection

We use the XLM-RoBERTa-large model in all our experiments. We train the model for 8 epochs, with a batch size of 16, AdamW optimizer (Loshchilov and Hutter, 2019) with 2e-5 learning rate. We use a max sequence size of 256 tokens, larger sequences are truncated to the first 256 tokens. The rest of hyper parameters are the default values for the sequence labelling implementation of the Huggingface library³. As described in 3.1 we evaluate the model in the development set at the end of each

³<https://github.com/huggingface/transformers/tree/main/examples/pytorch/token-classification>

epoch and then select the best performing checkpoint. We train five independent models and then use majority vote as the ensembling strategy at inference time.

B.2 Entity Linking and Information Retrieval

We use the mGENRE implementation⁴ in the fairseq library (Ott et al., 2019). We use the prefix tree provided by the authors and marginalization. This constrains the model to only generate valid identifiers. We retrieve data from Wikipedia using the pymediawiki⁵ wrapper and parser for the MediaWiki API. We use Wikidata client library for Python⁶ to retrieve knowledge from Wikidata.

B.3 Entity Category Classifier

We use the XLM-RoBERTa-large model in all our experiments. We train the model for 8 epochs, with a batch size of 16, AdamW optimizer (Loshchilov and Hutter, 2019) with 2e-5 learning rate. We use a max sequence size of 256 tokens. The rest of hyper parameters are the default values for the text classification implementation of the Huggingface library⁷. As described in 3.3 we evaluate the model in the development set at the end of each epoch and then select the best performing checkpoint. We train five independent models and then use majority vote as the ensembling strategy at inference time.

C Hardware used

We perform all our experiments using a single NVIDIA A100 GPU with 80GB memory. The machine used has two AMD EPYC 7513 32-Core Processors and 512GB of RAM. Although our system can also run on GPUs with 24GB of VRAM.

D Text Classification Confusion Matrix

We show the text classification confusion matrix (Figure 8) for the English development test using the gold labelled spans to better understand which type of entities are the most difficult ones to classify for our model. As we can see in the matrix, the entities related to persons are the most ambiguous and difficult to classify.

⁴https://github.com/facebookresearch/GENRE/tree/main/examples_mgenre

⁵<https://github.com/barrust/mediawiki>

⁶<https://github.com/dahlia/wikidata>

⁷<https://github.com/huggingface/transformers/tree/main/examples/pytorch/text-classification>

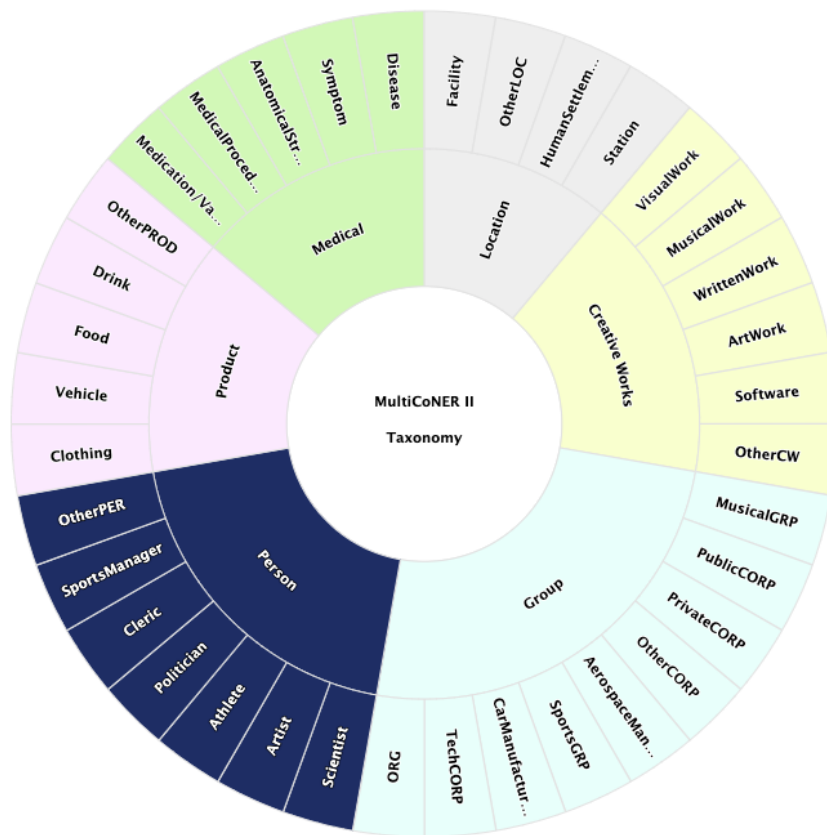


Figure 6: MultiCoNER2 Taxonomy

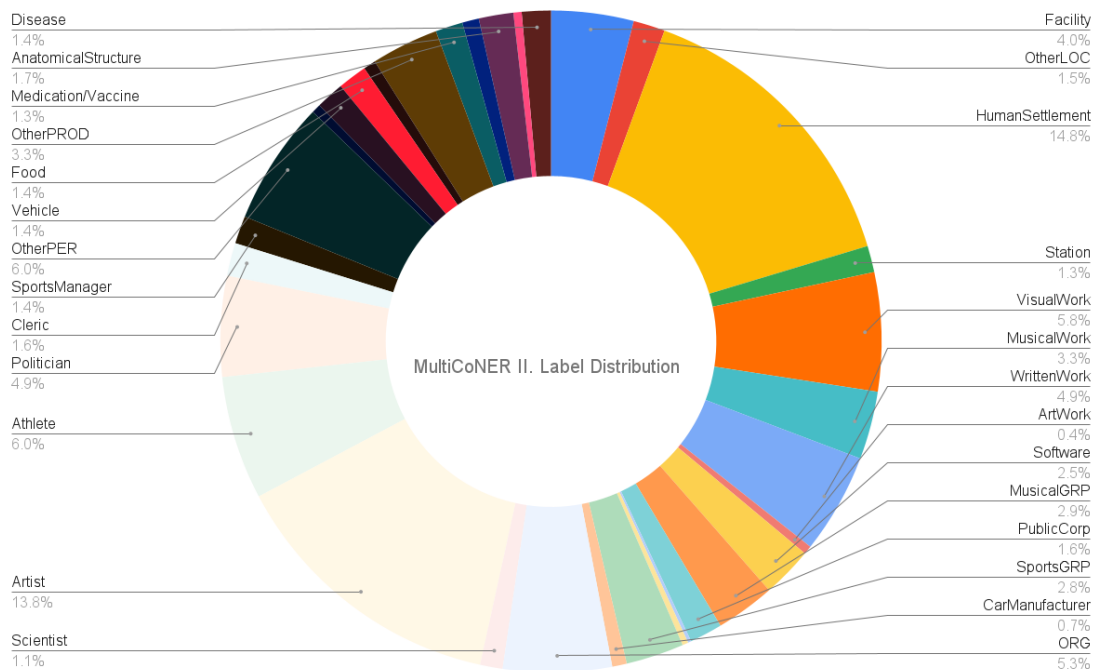


Figure 7: MultiCoNER2 label distribution computed as the sum of the label count for every language and split




Figure 8: text classification confusion matrix for the English development test using gold labelled spans

Lang	Split	Sentences
BN	train	9,708
	dev	507
	test	19,859
DE	train	9,785
	dev	512
	test	20,145
EN	train	16,778
	dev	871
	test	249,980
ES	train	16,453
	dev	854
	test	246,900
FA	train	16,321
	dev	855
	test	219,168
FR	train	16,548
	dev	857
	test	249,786
HI	train	9,632
	dev	514
	test	18,399
IT	train	16,579
	dev	858
	test	247,881
PT	train	16,469
	dev	854
	test	229,490
SV	train	16,363
	dev	856
	test	231,190
UK	train	16,429
	dev	851
	test	238,296
ZH	train	9,759
	dev	506
	test	20,265
MULTI	train	170,824
	dev	8,895
	test	358,668

Table 5: Number of sentences for each file in the dataset