

CLaC at SemEval-2023 Task 2: Comparing Span-Prediction and Sequence-Labeling approaches for NER

Harsh Verma, Sabine Bergler
CLaC Labs, Concordia University
{h_ver, bergler} @cse.concordia.ca

Abstract

This paper summarizes the CLaC submission for the MultiCoNER 2 task which concerns the recognition of complex, fine-grained named entities. We compare two popular approaches for NER, namely Sequence Labeling and Span Prediction. We find that our best Span Prediction system performs slightly better than our best Sequence Labeling system on test data. Moreover, we find that using the larger version of XLM RoBERTa significantly improves performance. Post-competition experiments show that Span Prediction and Sequence Labeling approaches improve when they use special input tokens (<s> and </s>) of XLM-RoBERTa. The code for training all models, preprocessing, and post-processing is available at [this Github repo](#).

1 Introduction

MultiCoNER 2 (Fetahu et al., 2023b,a) is the second version of the MultiCoNER (Malmasi et al., 2022) shared task. It focuses on identifying complex named entities in short sentences that lack context and addresses the challenge of identifying and distinguishing 30 fine-grained entity types and of handling simulated errors, such as typos.

Competition data is tokenized and tagged with the BIO scheme, i.e. the label B for *begin*, I for *inside*, and O for *outside* target entities. For this challenge, the 30 entity types were distinguished using suffixes to the B and I labels, e.g. B-Corp and I-Corp, as illustrated in Example 1.

NER is frequently formulated as a sequence labeling problem (Chiu and Nichols, 2016; Ma and Hovy, 2016; Wang et al., 2022), in which a model learns to label each token individually using the BIO labeling scheme. In contrast, span prediction approaches (Jiang et al., 2020; Li et al., 2020; Fu et al., 2021; Zhang et al., 2023) label entire text spans with entity types. Recently, (Fu et al., 2021) compared the two approaches and found that sequence labeling performed better for long entities

with low label consistency but that span prediction performed better for out of vocabulary words and entities of short to medium length. Because MultiCoNER 2 tests both, the ability to identify out of vocabulary entities and entities of varying length¹, it is interesting to compare the two approaches on this task.

Our submission to the English Track does not use any external knowledge and holds rank 21 of 41 on Codalab² with a macro-F1 of 59.

2 Submitted Systems

2.1 Tokenization

The competition input consisted of tokenized data. The gold labels are given in form of BIO tags, where each token is labeled with a corresponding BIO tag as shown in Example 1. Sample input is given to XLM RoBERTa (Conneau et al., 2020) as the raw input string, without RoBERTa input tags such as <s> and </s> in the submission systems.

(1)	<i>Token</i>	<i>pharma</i>	<i>company</i>	<i>eli</i>
	BIO-Tag	O	O	B-Corp
		<i>lilly</i>	<i>and</i>	<i>company announced</i>
		I-Corp	I-Corp	I-Corp O

2.2 Sequence Labeling Model

Our sequence labeling model uses XLM RoBERTa large (Conneau et al., 2020). Like all BERT-derived models, RoBERTa splits the input tokens further into word piece subtokens, for which the pre-trained embedding vectors are included, and for which the subtoken-level output is generated. These subtokens are classified into the BIO sequence. For submission and scoring, this subtoken sequence has to be decoded into word-level tokens.

¹Creative work titles can be long, e.g. *To Kill A Mocking Bird*

²<https://codalab.lisn.upsaclay.fr/competitions/10025#results>

Token Representation Step Given a sentence $\mathbf{x} = [w_1, w_2, \dots, w_n]$ with n word piece tokens, we generate for each token w_i a pretrained embedding \mathbf{u}_i using the last hidden layer of XLM RoBERTa large (Conneau et al., 2020)

$$\begin{aligned} \text{Embed}(\mathbf{x}) &= \text{Embed}([w_1, w_2, \dots, w_n]) \\ &= [\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_n] \end{aligned}$$

Token Classification Step In this layer, we classify every token representation into a set of named entity types corresponding to the BIO(beginning, inside, outside) tagging scheme. Assuming Θ is the set of all named entity types, the set of all BIO tags \mathbf{B} is of size $(2 \times |\Theta|) + 1$. We use a linear layer to map each subtoken $\mathbf{u}_i \in \mathbb{R}^d$ to a prediction $\mathbf{p}_i \in \mathbb{R}^{|\mathbf{B}|}$, where d is the length of the token embedding. The predictions are used to calculate loss of sample \mathbf{x} with n tokens as follows:

$$\text{Loss}(\mathbf{x}) = \frac{-1}{n} \sum_{i=1}^n \log(\text{Softmax}(\mathbf{p}_i)_{y_i}) \quad (1)$$

Here y_i represents the index of the gold BIO label of the i^{th} token.

Decoding Step For this task, the boundaries of each predicted span must align with actual word boundaries, which poses an issue due to the word piece tokenization. We align every predicted span with the nearest enveloping words. Concretely, let (b, e) represent the beginning and end offsets of a predicted span s . If b and e are contained in words w_b and w_e respectively, they are remapped to the beginning and end offsets of the containing words.

2.3 Span Prediction Model

Token Representation Layer The token representation layer is identical to that of the Sequence Labeling model.

Span Representation Layer Let a span s be a tuple $s = (b, e)$ where b and e are the begin and end word piece token indices, and s represents the text segment $[w_b, w_{b+1}, \dots, w_e]$ where w_i is the i^{th} word piece token. In this layer, we enumerate all possible spans, represented by the tuple (b, e) . Because $b \leq e$, there are $\frac{n^2}{2}$ possible spans. We follow (Fu et al., 2021) and encode each span s_i as the concatenation of their begin and end word piece token embeddings $\mathbf{v}_i = [\mathbf{u}_{b_i}; \mathbf{u}_{e_i}]$. The output of

the decoding layer is $\mathbf{V} \in \mathbb{R}^{k \times (2 \times d)}$ where $k = \frac{n^2}{2}$ and d is length of the token embedding vector.

Span Classification Layer In this layer, we classify each span representation with a named entity type. We introduce an additional label Neg_Span which represents the absence of a named entity. In particular, a linear layer maps each span representation $\mathbf{v}_i \in \mathbb{R}^{(2 \times d)}$ to a prediction $\mathbf{p}_i \in \mathbb{R}^{|\Omega|}$, where Ω is the set of all named entity types (including Neg_Span) and d is the size of the token embedding. The predictions are used to calculate the loss for sentence \mathbf{x} with l possible spans as follows:

$$\text{Loss}(\mathbf{x}) = \frac{-1}{l} \sum_{i=1}^l \log(\text{Softmax}(\mathbf{p}_i)_{y_i}) \quad (2)$$

Here y_i represents the index of the gold label of the i^{th} span.

Decoding Similarly to the Sequence Labeling model, we align all spans with word boundaries. Because the Span Prediction model predicts overlapping spans, we remove overlaps with the following procedure:

1. For each span s , if s is completely contained within another span S , we remove s . We keep removing fully contained spans until none are left.
2. For each span s_1 which partially overlaps another span s_2 , we randomly select one of the two spans and remove it. We keep removing until no overlapping spans are left.

The remaining spans are then mapped to BIO tags.

2.4 Training

XLM RoBERTa large is fine-tuned on the training data using the Adam optimizer (Kingma and Ba, 2015) with a learning rate of 1e-5 and a batch size of 1. The best model is selected using early stopping. Training for 10 epochs takes around 6 hours on one Nvidia RTX 3090 gpu.

3 Results and Discussion

We evaluate all of our systems on the Codalab Competition Website³, where submissions are evaluated using the entity level macro F1 metric. All models

³<https://codalab.lisn.upsaclay.fr/competitions/10025>

use the XLM-RoBERTa large LM unless otherwise indicated.

Table 1 shows the performance of the two submitted systems S(submitted) in bold.

Post-competition experimentation showed that increasing the batch size from 1 to 4 improved results by at least 3.2, this improvement is shown in Table 1 for systems B.

We also found that adding the embeddings for RoBERTa special tokens <s> and </s> further improved performance, especially raising performance of the Sequence model Seq by 2.1, nearly erasing the performance advantage of the span-based model. The results are shown in Table 1 for systems B+E.

System Type	Seq			Span		
	F1	P	R	F1	P	R
S(submitted)	53.2	52.8	54.5	55.0	56.5	55.5
B	56.9	57.3	57.5	59.0	61.2	58.4
B+E	59.0	60.7	58.9	59.7	61.3	59.5

Table 1: Performance on the test set in bold. The submitted systems improve when the batch size is increased from 1 to 4 (systems B), and they improve further when special RoBERTa token embeddings are added (systems B+E).

Table 2 shows the performance of our systems on the validation set.

The Span Prediction models showed significantly higher performance on the validation set than the test set. The post-competition improvements affect the sequence and span-based models differently and nearly erase the performance difference.

Model	Span _S	Seq _S	Span _B	Seq _B
Macro F1	58.5	59.3	63.6	60.2

Table 2: Performance on validation set for systems S and B

Table 3 compares the performance of large and base versions of XLM RoBERTa. Span_{base} and Seq_{base} are identical to Span_B and Seq_B except that the base models use XLM RoBERTa *base* instead of large. The larger models seem to be performing significantly better.

Model	Span _B	Seq _B	Span _{base}	Seq _{base}
Macro F1	59.0	56.9	52.5	51.1

Table 3: Performance of large and base pretrained models on test set for systems B

4 Conclusion

We submitted two systems to MultiCoNER 2, one inspired by a Sequence Labeling approach and another inspired by a Span Prediction approach. We find that our best Span Prediction system performs slightly better than our best Sequence Labeling system on test data. We showed significant increases in our systems’ performance when using a larger pretrained language model, batch size 4, and special tokens <s> and </s>.

References

- Jason P.C. Chiu and Eric Nichols. 2016. [Named entity recognition with bidirectional LSTM-CNNs](#). *Transactions of the Association for Computational Linguistics*, 4:357–370.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [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.
- Besnik Fetahu, Zhiyu Chen, Sudipta Kar, Oleg Rokhlenko, and Shervin Malmasi. 2023a. MultiCoNER v2: a Large Multilingual dataset for Fine-grained and Noisy Named Entity Recognition.
- 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.
- Jinlan Fu, Xuanjing Huang, and Pengfei Liu. 2021. [SpanNER: Named entity re-recognition as span prediction](#). 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 7183–7195, Online. Association for Computational Linguistics.
- Zhengbao Jiang, Wei Xu, Jun Araki, and Graham Neubig. 2020. [Generalizing natural language analysis through span-representation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2120–2133, Online. Association for Computational Linguistics.
- Diederik P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *Proceedings of the 3rd International Conference on Learning Representations, ICLR’15*.

- Xiaoya Li, Jingrong Feng, Yuxian Meng, Qinghong Han, Fei Wu, and Jiwei Li. 2020. [A unified MRC framework for named entity recognition](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5849–5859, Online. Association for Computational Linguistics.
- Xuezhe Ma and Eduard Hovy. 2016. [End-to-end sequence labeling via bi-directional LSTM-CNNs-CRF](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1064–1074, Berlin, Germany. Association for Computational Linguistics.
- Shervin Malmasi, Anjie Fang, Besnik Fetahu, Sudipta Kar, and Oleg Rokhlenko. 2022. 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.
- 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.
- Sheng Zhang, Hao Cheng, Jianfeng Gao, and Hoifung Poon. 2023. [Optimizing bi-encoder for named entity recognition via contrastive learning](#). In *The Eleventh International Conference on Learning Representations*.