

# PAI at SemEval-2023 Task 2: A Universal System for Named Entity Recognition with External Entity Information

Long Ma\*, Kai Lu\*, Tianbo Che, Hailong Huang, Weiguo Gao, Xuan Li  
{malong633, lukai113, chetianbo263, huanghailong590}@pingan.com.cn  
{gaoweiguo801, lixuan208}@pingan.com.cn  
Ping An Life Insurance Company of China, Ltd.

## Abstract

The MultiCoNER II task aims to detect complex, ambiguous, and fine-grained named entities in low-context situations and noisy scenarios like the presence of spelling mistakes and typos for multiple languages. The task poses significant challenges due to the scarcity of contextual information, the high granularity of the entities (up to 33 classes), and the interference of noisy data. To address these issues, our team PAI proposes a universal Named Entity Recognition (NER) system that integrates external entity information to improve performance. Specifically, our system retrieves entities with properties from the knowledge base (i.e. Wikipedia) for a given text, then concatenates entity information with the input sentence and feeds it into Transformer-based models. Finally, our system wins 2 first places, 4 second places, and 1 third place out of 13 tracks. The code is publicly available at <https://github.com/diqiuzhuanzhuan/semEval-2023>.

## 1 Introduction

The objective of the MultiCoNER II shared task (Fetahu et al., 2023b) is to develop a robust named entity recognition (NER) system for multiple languages (including English, Spanish, Hindi, Bangla, Chinese, Swedish, Farsi, French, Italian, Portuguese, Ukrainian, German and aforementioned mixed languages) that can perform well even when the input data has spelling errors, lacks contextual information, or contains out of knowledge base entities. This task entails the recognition of a total of 33 fine-grained entity types within six coarse-grained entity categories. The entity types within each category are highly prone to confusion and ambiguity, and the testing data includes noise such as spelling errors that must be considered. Similar to the previous competition (Malmasi et al., 2022b), the data in this competition contains a lot of short

and low-context sentences, which means the accurate prediction of entity types relies heavily on external knowledge.

Therefore, the key objective of our current work is to achieve effective integration of the model and external knowledge. In this paper, extending our previous work (Ma et al., 2022), which has demonstrated the efficacy of the Dictionary-fused model, we propose a NER system based on Entity Property Knowledge Base. We constructed our knowledge base using all entities and their associated properties (including entity name, alias, sitelink title, description, instanceof, subclassof and occupation) from WikiData. In the retrieval module, we use string matching to retrieve entities with properties for a given sentence. Given the sequence length limitations of our model, we keep longer entities in sentences. Subsequently, we put the entity and property values into the external context in a specific format and feed the external context and original sentence into the NER model. Figure 1 illustrates our motivation, demonstrating how the inclusion of property values can effectively improve performance in practice.

We further propose entity-aware attention mechanism, which can better model the semantic relationships between entities in sentences and external context than conventional attention. In the test phase, we use the voting strategy to form the final predictions.

We make the following observations according to our experiments:

1. The system based on entity property knowledge base has significantly improved the effect. The knowledge base has a high entity coverage, and even more than 95 percent in some low-resource languages such as Bangla and Hindi.
2. Different properties have different effects on different entity types. Specifically, occupation

\*These authors contributed equally to this work.

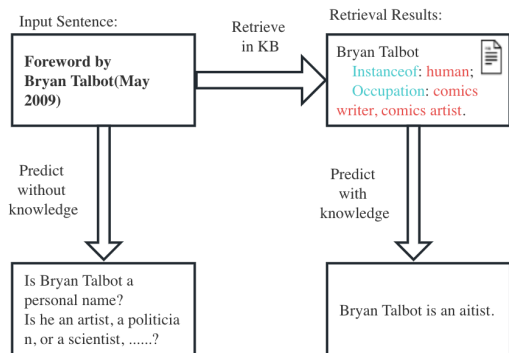


Figure 1: A motivating example from the English dev set. In the retrieval results, "Bryan Talbot" is an entity with two important properties. The blue phases are the property names and the red phases are the property values.

property can improve the recognition of fine-grained person entities.

3. We have observed that the dictionary can greatly improve the performance on clean data, but may have a negative impact on noisy data. Knowledge retrieved by fuzzy matching such as ElasticSearch can help identify noisy entity.
4. On some tracks, our entity aware attention can better capture the semantic relationships between entities in sentence and external context.

## 2 Related Work

Complex named entities (NE), like the titles of creative works, are not simple nouns and pose challenges for NER systems (Malmasi et al., 2022a; Fetahu et al., 2023a). To mitigate these challenges, NER models utilizing external knowledge have achieved remarkable results. Researchers have integrated gazetteers into models (Bender et al., 2003; Malmasi and Dras, 2016) in earlier studies. Recently, GEMNET (Meng et al., 2021) proposes a novel approach for gazetteer information integration, where a flexible Contextual Gazetteer Representation (CGR) encoder can be fused with any word-level model and a Mixture-of-Experts (MOE) gating network can overcome the feature overuse issue by learning to conditionally combine the context and gazetteer features, instead of assigning them fixed weights. GAIN (Chen et al., 2022)

adapts the representations of gazetteer networks to integrate external entity type information into models. (Ma et al., 2022) uses string matching to retrieve entities with types from the gazetteer and concatenates the entity information with the input text and feeds it to models. However, the prerequisite for these approaches is to construct an entity dictionary with accurate types, which is extremely difficult to fulfill. Knowledge-fused methods mentioned in (Wang et al., 2022; Çark et al., 2022) build an information retrieval (IR) system based on Wikipedia to provide related context information to models. But these methods require retrieving high-quality context, which is difficult to guarantee in the MultiCoNER II shared task (Fetahu et al., 2023b).

Actually, Transformer-based models (Devlin et al., 2019; Vaswani et al., 2017; Yamada et al., 2020) have the ability to automatically learn the mapping between the entity properties and entity types on the basis of input sentence. Therefore, We propose a NER system that fuses entity dictionaries without accurate entity types to overcome these drawbacks of previous work.

## 3 Our System

In this section, we introduce how our system works. Given a sentence  $x = \{x_0, x_1, \dots, x_{n-1}\}$ ,  $x$  consists of  $n$  tokens, and this sentence is input to the knowledge retrieval module. The retrieval module outputs the retrieved entities and property contexts. The system then concatenates the input sentence and outputs from the retrieval module and fed them to a Transformer-based model. To enhance the relevance of retrieved entities to the given information, we devised entity-aware attention mechanism. The overall architecture is shown in Figure 2.

### 3.1 Knowledge Retrieval Module

**Knowledge Base Construction** The entity property KB is constructed based on WikiData. WikiData contains more than 90 million different entities. Two entity examples with related fields are shown in Table 1.

$Qid$  represents an entity, and  $Label_{qid}$ ,  $Alias_{qid}$ ,  $Sitelink\_title_{qid}$ ,  $Subclassof_{qid}$ ,  $Instanceof_{qid}$  and  $Occupation_{qid}$  represents the values of different fields of the entity.  $Entity\_names_{qid}$  represents the entity names, consists of  $Label_{qid}$ ,  $Sitelink\_title_{qid}$  and  $Alias_{qid}$ .  $Property\_qid_{qid}$  represents the property values in qid format,

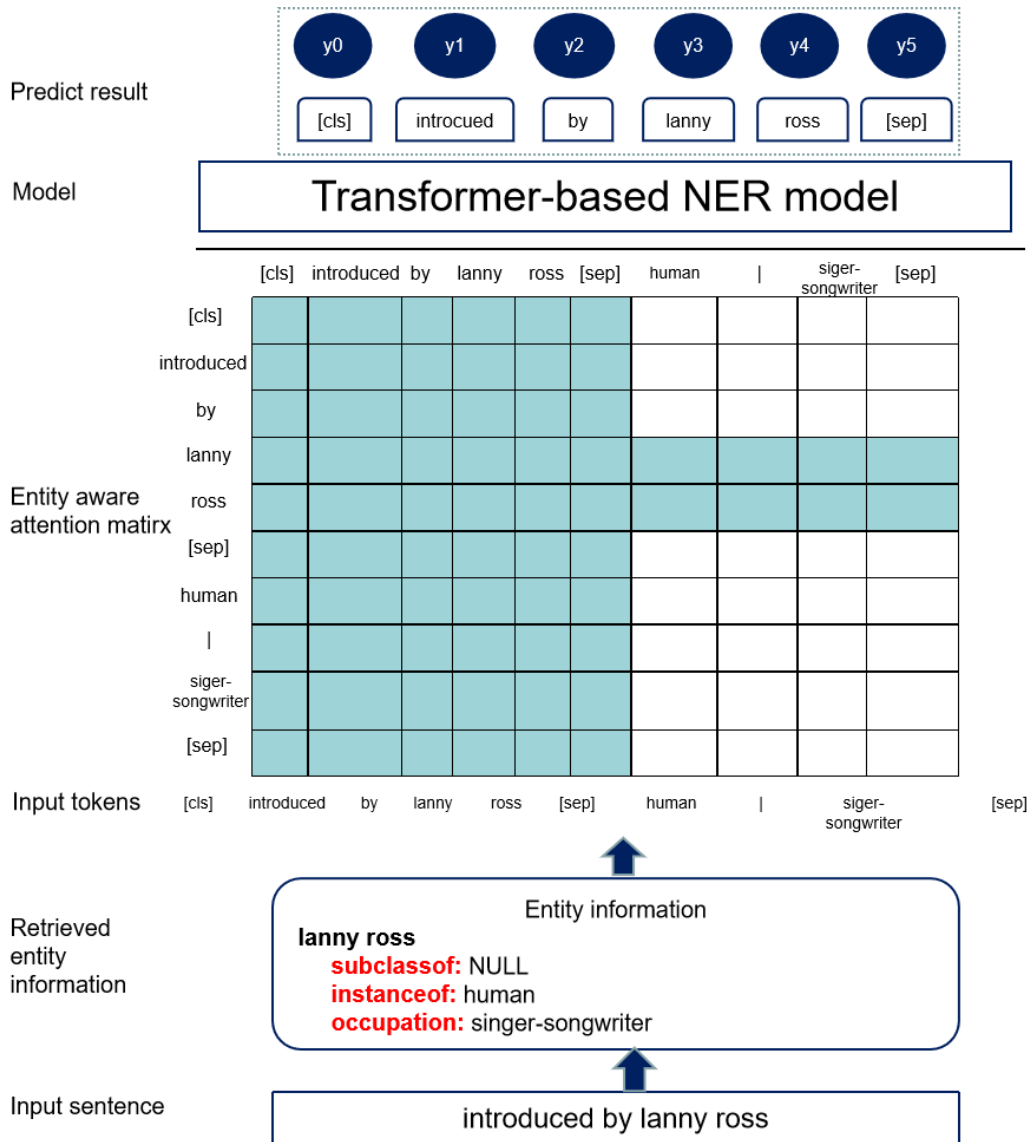


Figure 2: The overall architecture of our proposed system shows a detailed case. "NULL" means that the field in the entity information is empty. In input tokens, different property values are separated by the token "|". The middle of the figure shows the specific entity aware attention matrix, which is described in section 3.2. The blue grid represents a value of 1, and the white grid represents a value of 0.

consists of  $Subclassof_{qid}$ ,  $Instanceof_{qid}$  and  $Occupation_{qid}$ .  $Property\_names_{qid}$  represents the property names in text format, which is the label field values of  $Property\_qid_{qid}$ .  $Context_{qid}$  represent the entity context. In  $Entity\_context_{qid}$ ,  $Property\_names_{qid}$  is separated by a special token "|". Two entity examples with names and context are shown in Table 2. The knowledge base contains all entity names and contexts.

**Entity Retrieval** Given an input sentence, the retrieval module retrieves entity names by string matching. The external context consists of matched entity names and corresponding entity contexts. We prioritize preserving the longer entities when

encountering overlapping matches or the need to truncate information. Finally, retrieval module returns multiple entity context pairs,  $pairs = [(entity_0, context_0), (entity_1, context_1), \dots]$ . The length of the  $pairs$  is  $m$ .

### 3.2 Entity Aware Attention

Our system concatenates input sentence  $x$  and entity contexts together, and entity contexts are separated by a special token "\$".

In order to better use the external knowledge in Transformer-based model, entity-aware attention is proposed to link entities in sentence with corresponding contexts. Entity-aware attention

| Field                 | Example "Victor Cousin"          | Example "human"                               |
|-----------------------|----------------------------------|---|
| <b>Qid</b>            | Q434346                          | Q5  |
| <b>Label</b>          | <i>English: Victor Cousin</i>    | <i>English: human</i>                         |
| <b>Alias</b>          | <i>English: NULL</i>             | <i>English: [human being, humankind, ...]</i> |
| <b>Sitelink title</b> | <i>English: Victor Cousin</i>    | <i>English: Human</i>                         |
| <b>Instanceof</b>     | [Q5]                             | [Q55983715]                                   |
| <b>Subclassof</b>     | NULL                             | [Q154954, Q164509, ...]                       |
| <b>Occupation</b>     | [Q4964182, Q82955, Q333634, ...] | NULL  |

Table 1: Entity examples in WikiData. For the label, alias, and sitelink title fields, the examples only show the value in English, but WikiData supports multiple languages. Qid is used to identify a unique entity. The value of the instanceof, subclassof, and occupation properties is qid. The value of label and sitelink title is a string, and other fields contain multiple strings. When the entity lacks a field, we use "NULL" value in the table.

|                       | Example "Victor Cousin"                | Example "human"                          |
|-----------------------|--|--|
| <b>Qid</b>            | Q434346                                | Q5                                       |
| <b>Entity_names</b>   | [Victor Cousin]                        | [human, human being, ...]                |
| <b>Entity_context</b> | human   philosopher   politician ..... | natural person   omnivore   mammal ..... |

Table 2: Entity names and context.

| Language | Manual Dict | Property Dict |
|----------|-------------|---------------|
| BN       | 41.7        | 96.9          |
| DE       | 55.6        | 98.9          |
| EN       | 60.6        | 94.3          |
| ES       | 57.0        | 95.1          |
| FA       | 34.6        | 84.8          |
| FR       | 53.5        | 95.2          |
| HI       | 37.0        | 99.6          |
| IT       | 60.4        | 96.2          |
| PT       | 56.3        | 93.7          |
| SV       | 55.9        | 93.9          |
| UK       | 26.9        | 65.2          |
| ZH       | 42.9        | 99.6          |
| AVG      | 48.5        | 92.8          |

Table 3: Coverage rate on dev data. The definition of coverage rate is the same as (Chen et al., 2022). "Coverage rate" refers to the number of entities in the data also found in our dictionary / the number of entities in the data.

matrix  $M$  is a binary matrix with values 0 or 1, which makes context only attended by the corresponding entity. Under both conditions  $M_{i,j} = 1$ , other times  $M_{i,j} = 0$ . In the first condition,  $i, j < n + 2$ .  $n + 2$  takes into account [CLS] token at the front of the sentence and [SEP] token at the end of the sentence. In the second condition,  $i$  in  $position_{entity_k}$ , and  $j$  in  $position_{context_k}$ ,  $0 \leq k < m$ .  $position_{entity_k}$  and  $position_{context_k}$  are position subscript sets of  $entity_k$  and  $context_k$

respectively. Entity-aware attention matrix  $M$  is input into Transformer-based NER model as attention mask. An entity-aware attention matrix example is shown in Figure 2.

## 4 Experimental Setup

### 4.1 Data and Evaluation Metrics

MultiCoNER II contains 6 coarse-grained labels, which can be further subdivided into 33 fine-grained labels. The specific labels are shown in table 8. Fine-grained named entity recognition (NER) is a particularly challenging task, especially in low-context and ambiguous situations. For instance, identifying names of individuals who are scientists or musicians can be super difficult without sufficient context.

The task has 12 monolingual tracks, including multiple low-resource language tracks that are difficult to obtain external knowledge, such as Hindi and Bangla. In this paper, all results use entity level micro F1 as evaluation metric except table 5. Table 5 use macro F1 which is official evaluation metric.

### 4.2 Training detail

We use Chinese BERT with Whole Word Masking (Cui et al., 2019) as our pre-trained model in Chinese and utilize BERT multilingual base model (uncased)(Wolf et al., 2020) in other tracks. An 8-fold cross-validation training strategy is applied

| System               | DE          | BN          | ZH          | EN          | ES          | HI          | SV          | FA          | FR          | IT          | PT          | UK          | AVG         |
|----------------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
| <b>Baseline</b>      | 72.1        | 75.0        | 68.3        | 69.7        | 73.4        | 71.1        | 75.2        | 60.2        | 74.0        | 78.5        | 71.1        | 67.6        | 71.3        |
| <b>Manual Dict</b>   | 78.1        | 78.5        | 75.2        | 74.5        | 73.8        | 79.2        | 78.8        | 66.7        | 77.9        | 78.7        | 77.5        | 69.2        | 75.7        |
| <b>Property Dict</b> | <b>88.5</b> | <b>89.4</b> | <b>91.8</b> | <b>81.2</b> | <b>81.9</b> | <b>88.1</b> | <b>84.2</b> | <b>73.5</b> | <b>83.2</b> | <b>85.3</b> | <b>80.6</b> | <b>74.8</b> | <b>83.5</b> |

Table 4: Dev micro F1 score with different dictionary. **Manual Dict** means the model with type dict which needs manual relational mapping, **Property Dict** means the model with entity property knowledge.

| Data Type         | DE    | BN    | ZH    | EN    | ES    | HI    | SV    | FA    | FR    | IT    | PT    | UK    |
|-------------------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|
| <b>Clean Data</b> | 88.09 | 84.39 | 86.23 | 86.16 | 79.35 | 80.96 | 81.53 | 68.46 | 89.5  | 88.94 | 84.56 | 71.28 |
| <b>Noisy Data</b> | -     | -     | 41.9  | 65.41 | 55.25 | -     | 55.22 | -     | 78.71 | 76.53 | 76.12 | -     |
| <b>Overall</b>    | 88.09 | 84.39 | 74.87 | 80    | 71.67 | 80.96 | 72.38 | 68.46 | 86.17 | 84.88 | 81.64 | 71.28 |

Table 5: Macro F1 score on clean data and noisy data in the evaluation phase. "-" means that there is no noisy data.

|                           | Average F1  |
|---------------------------|-------------|
| All property              | <b>83.5</b> |
| All property - subclassof | 82.5        |
| All property - instanceof | 80.4        |
| All property - occupation | 79          |

Table 6: Average micro F1 score of all monolingual tracks in the dev set when using different properties to construct the dictionary.

|                                     | BN   | IT   |
|-------------------------------------|------|------|
| <b>Model with default attention</b> | 78.5 | 78.7 |
| <b>Model with entity</b>            | 79.5 | 79   |

Table 7: Dev micro F1 score with entity aware attention on Bangla and Italian.

in the evaluation except for MULTI task. In the evaluation phase, all the best models vote on the prediction results, and the voting weights are determined by the F1 score of each model on the validation set.

## 5 Results and Analysis

### 5.1 Main Results

To show the effectiveness of our system, we evaluate the results of a baseline system without knowledge. We reproduce the dictionary construction process of (Chen et al., 2022) on MultiCoNER II, and obtain a type dictionary through manual relational mapping. The results of the baseline, Manual Type dictionaries, and our system are shown in Table 4. The automatic property dictionary exceeds the baseline by 12.2 F1 on average, demonstrating the effectiveness of our system. Without manual relational mapping, our proposed property dictionary outperforms the manually constructed type

dictionary by 7.8 F1 on average, demonstrating the capability of the model to learn the correlation between property values and entity types automatically.

### 5.2 Coverage Rate Trial

Previous studies (Rijhwani et al., 2020; Meng et al., 2021) have found that systems with higher entity coverage have higher performance. Table 3 shows the entity coverage of Manual Type Dictionary and Property Dictionary. The coverage rate of the property dictionary is 44 percent higher than that of the artificial dictionary and even reaches 99 percent in some languages, such as Chinese and Hindi. The track with higher coverage rate has bigger improvement in Table 4.

### 5.3 Effect of Property

In order to analyze the impact of different properties on the results, we conduct ablation experiments of properties. Table 6 shows the results of using different properties. We observe that the *Occupation* property has a greater impact on the performance. Without *Occupation* property, the average results drop by 4.5 F1. Further, we analyze the relationship between properties and taxonomies. Table 8 shows the taxonomy F1 score when using different property combinations. The *subclassof* property is strongly related to product class. The *instanceof* property is strongly correlated with Creative Works, Group, Location, and Medical classes. *Occupation* is a property about people. Experiments show that occupation property plays a crucial role in identifying Person class.

### 5.4 Effect of Entity Aware Attention

We propose entity-aware attention, which can represent the relationship between entities and context in

| Coarse-grained taxonomy | Fine-grained taxonomy | All property | All property - subclassof | All property - instanceof | All property - occupation |
|-------------------------|-----------------------|--------------|---------------------------|---------------------------|---------------------------|
| Creative Works          | ArtWork               | 76.8         | 76.2                      | <b>69.8</b>               | 78.7                      |
|                         | MusicalWork           | 82.4         | 81.4                      | <b>77.3</b>               | 82.6                      |
|                         | Software              | 82.0         | 81.5                      | <b>79.2</b>               | 84.0                      |
|                         | VisualWork            | 85.3         | 83.9                      | <b>81.2</b>               | 87.6                      |
|                         | WrittenWork           | 79.7         | 78.3                      | <b>73.8</b>               | 77.6                      |
| Group                   | PublicCorp            | 80.7         | 81.6                      | <b>66.5</b>               | 78.3                      |
|                         | AerospaceManufacturer | 92.0         | 91.1                      | <b>83.5</b>               | 87.4                      |
|                         | CarManufacturer       | 84.0         | 82.8                      | <b>80.7</b>               | 82.4                      |
|                         | MusicalGRP            | 90.5         | 89.3                      | <b>85.8</b>               | 90.5                      |
|                         | ORG                   | 79.1         | 78.6                      | <b>68.6</b>               | 76.9                      |
|                         | PrivateCorp           | 83.3         | 74.8                      | <b>70.3</b>               | 75.5                      |
| Location                | SportsGRP             | 89.8         | 87.5                      | <b>86.9</b>               | 89.8                      |
|                         | Facility              | 76.6         | 75.5                      | <b>73.2</b>               | 79.0                      |
|                         | HumanSettlement       | 90.7         | 90.2                      | <b>87.7</b>               | 90.3                      |
|                         | OtherLOC              | 74.3         | <b>66.5</b>               | 71.1                      | 73.3                      |
| Medical                 | Station               | 82.5         | 82.9                      | <b>77.4</b>               | 84.0                      |
|                         | AnatomicalStructure   | 75.7         | 78.2                      | <b>72.0</b>               | 74.8                      |
|                         | Disease               | 77.2         | <b>72.5</b>               | <b>72.5</b>               | 74.8                      |
|                         | MedicalProcedure      | 73.4         | <b>68.0</b>               | 68.6                      | 76.8                      |
|                         | Medication/Vaccine    | 81.9         | 79.4                      | <b>78.4</b>               | 80.8                      |
| Person                  | Symptom               | 82.1         | 74.9                      | <b>70.0</b>               | 76.6                      |
|                         | Artist                | 90.3         | 89.9                      | 89.9                      | <b>82.7</b>               |
|                         | Athlete               | 87.1         | 88.6                      | 87.6                      | <b>74.1</b>               |
|                         | Cleric                | 81.5         | 78.9                      | 81.6                      | <b>65.2</b>               |
|                         | OtherPER              | 75.6         | 75.0                      | 73.4                      | <b>57.5</b>               |
|                         | Politician            | 83.6         | 83.7                      | 84.1                      | <b>62.8</b>               |
|                         | Scientist             | 77.7         | 78.9                      | 77.0                      | <b>52.7</b>               |
| Product                 | SportsManager         | 89.0         | 90.5                      | 89.9                      | <b>67.9</b>               |
|                         | Clothing              | 65.1         | <b>59.2</b>               | 65.5                      | 62.3                      |
|                         | Drink                 | 77.8         | <b>71.1</b>               | 75.3                      | 78.6                      |
|                         | Food                  | 66.0         | <b>61.4</b>               | 63.4                      | 66.0                      |
|                         | OtherPROD             | 73.0         | <b>68.0</b>               | 70.2                      | 72.8                      |
|                         | Vehicle               | 69.0         | <b>65.5</b>               | 70.4                      | 67.6                      |

Table 8: The F1 score on taxonomy with different properties combination. In the table, the bold numbers are the lowest F1 on taxonomies.

a sentence. Table 7 shows the effect of entity-aware attention on several languages. The entity-aware attention improves by 1 F1 and 0.3 F1 over the baseline on Bangla and Italian, demonstrating its effectiveness.

### 5.5 Clean Data and Noisy Data

Table 5 shows the results of our system on clean data, noisy data, and overall data in evaluation. The results demonstrate that our system achieves higher performance on clean data than on noisy data, which poses a challenge due to the noisy entities that can not be retrieved through string matching. The Chinese track exhibits the largest discrepancy between clean and noisy data since the Chinese training and validation sets only contain clean data, while the test set includes both clean and noisy data.

## 6 Conclusion

In this paper, we describe our NER system based on entity properties, which wins two tracks in Multi-CoNER II shared task. We construct a KB based on entity properties, which is used to retrieve the relevant entity names and contexts for a given sentence. Our property dictionary is built without the need for manual relational mapping and achieves high coverage on the test set. We have found that different entity types require different properties. We propose the entity-aware attention mechanism to better learn the relationship between entities and contexts. In the future, we plan to adopt fuzzy matching to improve the performance on noisy data and explore our system for other low-resource tasks.

## References

- Oliver Bender, Franz Josef Och, and Hermann Ney. 2003. Maximum entropy models for named entity recognition. In *Proceedings of the seventh conference on Natural language learning at HLT-NAACL 2003*, pages 148–151.
- Buse Çarık, Fatih Beyhan, and Reyyan Yeniterzi. 2022. [SU-NLP at SemEval-2022 task 11: Complex named entity recognition with entity linking](#). In *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*, pages 1648–1653, Seattle, United States. Association for Computational Linguistics.
- 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.
- Yiming Cui, Wanxiang Che, Ting Liu, Bing Qin, Ziqing Yang, Shijin Wang, and Guoping Hu. 2019. Pre-training with whole word masking for chinese bert. *arXiv preprint arXiv:1906.08101*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. 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.
- Long Ma, Xiaorong Jian, and Xuan Li. 2022. [PAI at SemEval-2022 task 11: Name entity recognition with contextualized entity representations and robust loss functions](#). In *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*, pages 1665–1670, Seattle, United States. Association for Computational Linguistics.
- Shervin Malmasi and Mark Dras. 2016. Location mention detection in tweets and microblogs. In *Computational Linguistics: 14th International Conference of the Pacific Association for Computational Linguistics, PACLING 2015, Bali, Indonesia, May 19-21, 2015, Revised Selected Papers 14*, pages 123–134. Springer.
- Shervin Malmasi, Anjie Fang, Besnik Fetahu, Sudipta Kar, and Oleg Rokhlenko. 2022a. MultiCoNER: a Large-scale Multilingual dataset for Complex Named Entity Recognition.
- 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.
- Shruti Rijhwani, Shuyan Zhou, Graham Neubig, and Jaime Carbonell. 2020. [Soft gazetteers for low-resource named entity recognition](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8118–8123, Online. 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](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- 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, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Ikuya Yamada, Akari Asai, Hiroyuki Shindo, Hideaki Takeda, and Yuji Matsumoto. 2020. Luke: Deep contextualized entity representations with entity-aware self-attention. *arXiv preprint arXiv:2010.01057*.