

Biomedical Nested NER with Large Language Model and UMLS Heuristics

DS@GT CLEF2024 BioNNE Competition Working Note

Wenxin Zhou^{1,*}

¹Georgia Institute of Technology, North Ave NW, Atlanta, GA 30332, United States

Abstract

In this paper, we present our system for the BioNNE English track, which aims to extract 8 types of biomedical nested named entities from biomedical text. We use a large language model (Mixtral 8x7B instruct) and ScispaCy NER model to identify entities in an article and build custom heuristics based on unified medical language system (UMLS) semantic types to categorize the entities. We discuss the results and limitations of our system and propose future improvements. Our system achieved an F1 score of 0.39 on the BioNNE validation set and 0.348 on the test set.

Keywords

large language model, prompt engineering, named entity recognition, nested named entity recognition

1. Introduction

Named entity recognition is an information extraction task that aims to identify named entities in text and classify them into predefined categories. Nested named entity recognition involves detecting both the outer entities and the inner entities. The BioNNE competition [1], being part of CLEF 2024 BioASQ lab [2], focuses on extracting nested named entities from biomedical text. The nested named entity types include anatomy (ANATOMY), chemicals (CHEM), diseases (DISO), physiology (PHYS), scientific findings (FINDING), injury or poisoning damages (INJURY_POISONING), lab procedures (LABPROC), and medical devices (DEVICE) [3].

The challenge offers Russian, English and bilingual tracks. For the English track, the organizers provided a training set with 50 records and a validation set with 50 records. Each record contains a text, which is a PubMed abstract, and a list of entities annotated in BRAT format with the starting and ending locations of the entities in the text. In the testing phase, the organizers released a test set with 154 abstracts and 346 extra files, which resulted in 500 records in total.

Our team focuses on the BioNNE English track. Our system uses a large language model (specifically Mixtral 8x7B instruct model [4]) and a biomedical NER model to find entities in an article. Then, the system uses unified medical language system (UMLS) semantic types to filter and aggregate the entities. The implementation can be found on Github¹.

2. Related Work

2.1. Nested NER

The state-of-the-art nested NER models are the machine reading comprehension (MRC) model [5] and the sequence learning model [6]. Loukachevitch et al. trained the MRC and sequence model on the NEREL-BIO dataset [3], which is the predecessor of the BioNNE dataset. The two datasets are based on PubMed abstracts and with the nested NER annotations. The difference is that the NEREL-BIO dataset

CLEF 2024: Conference and Labs of the Evaluation Forum, September 09-12, 2024, Grenoble, France

*Corresponding author.

✉ wzhou77@gatech.edu (W. Zhou)

ORCID 0009-0002-3325-3357 (W. Zhou)

© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

¹<https://github.com/dsgt-kaggle-clef/bioasq-2024>

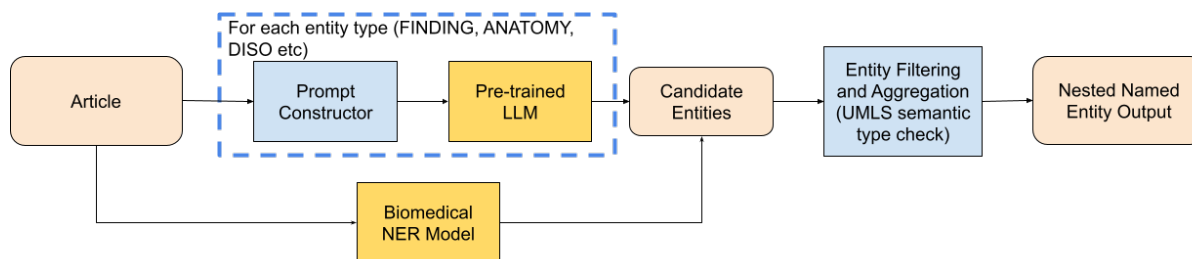


Figure 1: BioNNE System Design

contains more entity types, including the BioNNE entity types and additional entity types such as food, gene and activity. Loukachevitch et al. were able to achieve a macro-F1 score of 0.5968 with MRC model on the NEREL-BIO dataset.

2.2. Large Language Model

Recently, large language models (LLM) have shown great potential in solving NLP tasks. Since they are pre-trained on a large corpus, they can solve various problems in different domains including biomedical question answering and information extraction. Prompt engineering [7] is one of the key techniques to interact with LLM. The model is given a prompt with instructions and examples. The model can then generate the response according to the prompt. This is also known as LLM few-shot in-context learning. The quality of the prompt affects model output and performance. As a result, a lot of research has been done on prompt engineering to optimize the performance of LLM on specific tasks, as an alternative to model finetuning, because finetuning is more expensive and time-consuming.

Chen et al. [8] measured the performance of GPT-3.5 on the Biomedical Language Understanding and Reasoning Benchmark (BLURB). GPT-3.5 achieved 58.5 while the SOTA model had a score of 84.5. Although the existing LLM models are not as good as the SOTA biomedical models, LLMs show great potential in text reasoning and generation. At the same time, LLMs have limitations such as hallucination and inconsistency [9].

Mixtral 8x7B [4], a sparse mixture of experts model (SMoE) is one of the well-known LLM models, which can process five languages including English, French, Italian, German and Spanish with a context of 32k tokens. The basic model and the instruct model of Mixtral 8x7B achieve better performance than the counterparts of GPT-3.5 [10] and Llama2 70B [11] in several benchmarks. Mixtral 8x7B only uses 13B parameters during inference time, which makes it more efficient than other models with the same performance. Therefore, we use Mixtral 8x7B instruct as our LLM model in this study.

3. Methodology

We build a system for the BioNNE English track that uses a general-purpose LLM and a biomedical domain language model to identify entities in an article and then uses custom heuristics based on unified medical language system (UMLS) semantic types to determine the entity types. The system design is shown in Figure 1. Figure 2 shows an example of how DISO and LABPROC entities are extracted from an article by LLM and categorized by UMLS heuristics.

3.1. LLM and Prompt Engineering

Given an article abstract, we use Mixtral 8x7B instruct LLM model to find the entities for each category. We construct a prompt for each category using two of the examples in the training dataset and an instruction that describes the entity type. Then the LLM returns a list of entities separated by semicolon in the response. We then parse the returned entities and store them as candidates for the category.

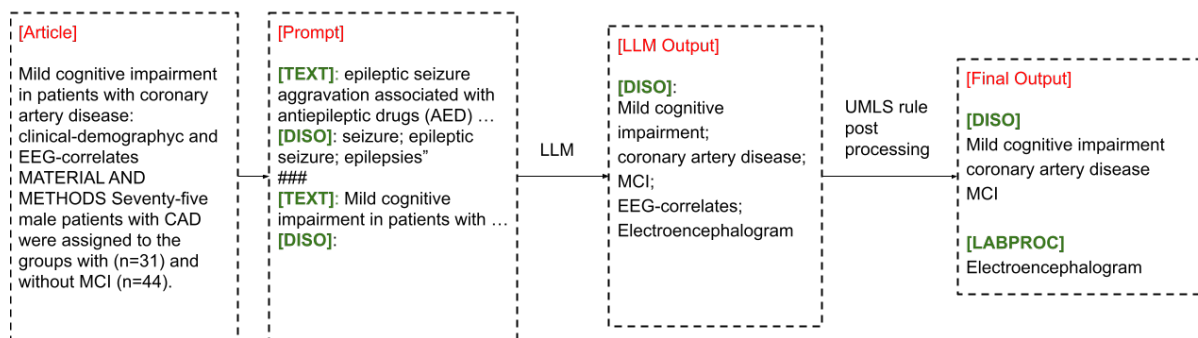


Figure 2: BioNNE Data Flow

We perform post-processing on the response list to remove duplicates as well as entities that are not present in the original text. If the entities in the returned string are not separated by semicolons, we resample the LLM output. For example, Table 1 shows the prompt and response to query ANATOMY entity for record 26271422_en.

Table 1

Example Prompt and Response for ANATOMY entity

Prompt:

Instruction: Return phrases or entities that comprise organs, body part, cells and cell components, body substances in TEXT, in the ANATOMY concatenated by ";"

[TEXT]: Impact of bosentan therapy on stress-induced pulmonary hypertension in patients with systemic sclerosis. AIM To describe hemodynamic and clinical changes in patients with elevated mean pulmonary artery pressure (MPAP) ...

[ANATOMY]: pulmonary; artery; pulmonary artery; lung;heart; left heart; atrial; right atrial; cardiac; arterial; vascular; pulmonary arterial; pulmonary vascular

###

[TEXT]: The authors present the material of their study of the morphological and molecular biological features of damage to the stem cell ...

[ANATOMY]: lung biopsies; respiratory acinus; lung tissue; mesenchymal cell; myofibroblast; mesenchymal stem cell; SCN; stem cell; cell; lung;pulmonary; acinus; stem cell niches; tissue;mesenchymal; SCN areas; respiratory acini; biopsies; sections; acini; cells

###

[TEXT]: Endoprosthetic replacement with lifting of abdominal wall in treatment of umbilical and postoperative ventral hernias. The results of ...

[ANATOMY]:

Response:

abdominal wall defect; abdomen ptosis; polypropylene prosthesis; main flap; roundish edges; additional flap; super lightweight; hypogastric area; musculoaponeurotic tissues lifting; physical health component; psychic health component; standard; original

3.2. Biomedical NER model

In addition to LLM, we also use the ScispaCy biomedical NER model *en_ner_bc5cdr_md* [12] to find the named entities. The ScispaCy model is trained on the BC5CDR corpus, which detects DISEASE and

Table 2

BioNNE type and UMLS type mapping

BioNNE Category	UMLS semantic types
DISO	B2.2.1.2 Pathologic Function
CHEM	A1.4.1 Chemical
ANATOMY	A1.2 Anatomical Structure A2.1.5.2 Body Location or Region
LABPROC	B1.3.1.1 Laboratory Procedure B1.3.1.2 Diagnostic Procedure
INJURY_POISONING	B2.3 Injury or Poisoning
DEVICE	A1.3.1 Medical Device
PHYS	B2.2.1.1 Physiologic Function
FINDING	A2.2 Finding

CHEMICAL entities in the document. We store the entities found by the ScispaCy model as candidates for the DISO and CHEM categories.

3.3. UMLS Heuristics

Finally, we query UMLS [13] to find the semantic types for the candidate entities. The UMLS release version used in this study is 2023AB. Specifically, we use UMLS [14] `"/search/current"` API and set the search term as our entity to retrieve the top 5 Concept Unique Identifiers (CUI) and their names associated with the entity. For each CUI, we use the `"/content/current/CUI/cui"` API to retrieve the semantic type name for the CUI and then use the `semantic_uri` to find the treeid of the semantic type. Finally, the treeid is mapped to the BioNNE category to finalize the category of an entity.

An entity is finalized as a specific category only when both the language model (either LLM or NER model) and the UMLS heuristics agree on the category. If UMLS heuristics show that the entity fits multiple categories in the top 5 CUI results, we will use the category type that comes first as the final category. The mapping between BioNNE category and UMLS semantic type is shown in Table 2. The mapping is based on the NEREL-BIO paper [3], with additional UMLS semantic types added for ANATOMY and LABPROC categories.

3.4. Acronym Detection

We use the ScispaCy abbreviation detector [12] to detect acronyms in the abstract. If the long form of the acronym is identified by the language model and UMLS heuristics as a named entity, the acronym will be assigned the same category as the long form entity.

4. Results and Discussion

Our model achieved 0.348 F1 score on the leaderboard for the BioNNE English track test set. The F1 score we achieved on the validation set is 0.39, which is close to the test F1 score. Since the golden answers for the test set have been not released, we will discuss the validation set results here.

The English validation set contains 50 records. The F1 scores for each category sorted in descending order are shown in Table 3. Our model performs well on DISO, CHEM, and ANATOMY entities, with F1 scores above 0.5, but the F1 scores for FINDING and PHYS are below 0.3.

We examine the model predictions for FINDING and PHYS entities and have the following observations.

1. The LLM model generates a lot of false positives for the FINDING entities. The UMLS heuristics can filter out some of the false positives. However, since our UMLS concept search does not search for exact terms, some of the LLM-recognized phrases that are part of some FINDING entity word are not filtered out, even though those LLM-recognized phrases (when looked at independently)

Table 3

F1 Score for each category on the validation set

Category	Precision	Recall	F1 Score
DISO	0.7565	0.5613	0.6444
CHEM	0.7857	0.4695	0.5878
ANATOMY	0.8101	0.4082	0.5429
LABPROC	0.5357	0.2632	0.3529
INJURY_POISONING	0.2778	0.4545	0.3448
DEVICE	0.6250	0.2	0.3030
PHYS	0.4875	0.1444	0.2229
FINDING	0.1288	0.2537	0.1709

do not belong to the FINDING type. For example, for the sentence (in record 27029443_en), "The authors suggest the algorithm for choosing the order of priority of surgical interventions on coronary and brachiocephalic arteries", the LLM recognizes "suggest" as a FINDING entity, but the UMLS heuristics do not filter it out because it corresponds to "Abnormal/suggest Ca", which is a FINDING concept.

2. PHYS entity recognition has a low recall. On one hand, PHYS entities sometimes contain generic one-word terms such as "healthy", "lifetime", "size", "shape" and "adults", which are not recognized by LLM or UMLS as PHYS entities. On the other hand, some long phrases (which are outer nested entities) recognized by LLM as PHYS entities are excluded by UMLS heuristics because the UMLS search cannot find such terms. For example, "spirometric indicators" and "peripheral blood oxygen saturation" in record 27030325_en are excluded by UMLS heuristics.

The general takeaways from the results are:

1. Context is important in NER task. The same phrases have different entity types in different documents. For example, the words "inflammatory" and "albuminuria" can be PHYS or DISO in different contexts. Our UMLS heuristics only check the semantic meaning of phrases without looking at the context, which limits the performance of our model.
2. Our model is often unable to detect the outer nested named entities. For example, it recognizes "cardiac contractility" as PHYS, but fails to recognize "low cardiac contractility" as FINDING. Sometimes those outer nested named entities are identified by LLM but rejected by UMLS heuristics because UMLS search cannot find matching terms in the UMLS system.

4.1. Impact of UMLS Heuristics

To understand the impact of UMLS heuristics on the system performance, we remove the UMLS heuristics from the system pipeline. When the LLM model identifies multiple entity types for a phrase, we assign the final entity type according to the INJURY_POISONING, ANATOMY, PHYS, DISO, CHEM, LABPROC, DEVICE and FINDING order. The macro-F1 score on the validation set without UMLS heuristics is 0.2151, which is significantly lower than the F1 score (0.348) when UMLS heuristics are used. The F1 scores for every category are shown in Table 4. We observe that without UMLS heuristics, the precision score for each category becomes lower. This is because the LLM model generates many false positives. In terms of recall, only the INJURY_POISONING and ANATOMY types have better recall scores when UMLS heuristics are not used. This is because we prioritize assigning the final entity type as INJURY_POISONING and ANATOMY when the LLM detects multiple matching entity types. This result shows that UMLS heuristics play an important role in eliminating the false positives of LLM predictions.

4.2. Training Dataset and Results

The Mixtral 8x7B model officially supports five languages, which do not include Russian. Therefore, we only attempt the English track. It is possible that other LLM models that support Russian can be used

Table 4

F1 Score for each category on the validation set when UMLS heuristics are not used

Category	Precision	Recall	F1 Score
DISO	0.4551	0.4676	0.4613
CHEM	0.3731	0.3049	0.3356
ANATOMY	0.2036	0.4629	0.2828
LABPROC	0.1964	0.0973	0.1302
INJURY_POISONING	0.1304	0.5455	0.2105
DEVICE	0.069	0.08	0.0741
PHYS	0.0742	0.10	0.0852
FINDING	0.1357	0.1463	0.1408

to accomplish the Russian/bilingual track using the same modeling pipeline. The method we use in this study is few-shot prompt engineering and we only use two examples from the training dataset to construct the LLM prompt. We did not include more examples in the prompt, as the input token size of LLM is limited and the cost of processing a long prompt is high. We also do not expect adding more examples in the prompt would improve the performance significantly, as the prompt is mostly used to guide the model to generate outputs in the desired format. The entities recognized by the model are mostly based on the original knowledge of the model. However, we believe the remaining training examples will be useful if we want to finetune the LLM model on the BioNNE task in the future.

5. Future Work

There are two directions for future improvement.

First, we can improve the heuristics for determining the category of an entity. Our current algorithm is weak at detecting the outer nested entities. Even though the LLM model recognizes them, we fail to assign the correct category for those entities using the existing UMLS heuristics. We can add a new set of heuristics for detecting the outer nested entities. Specifically, if a long phrase contains an inner entity that is a named entity, we can assign the outer entity with the same category as the inner entity or assign the outer entity with a category detected by the LLM.

Second, we can improve the performance of the LLM model. Currently LLM model generates many false positives and it may assign several entity types to a single phrase. For example, it recognizes "neuronal dysfunction" as DISO, ANATOMY, PHYS, FINDING and CHEM in the sentence "the severity of coronary artery lesions and low cardiac contractility affect the degree of cerebral ischemia and neuronal dysfunction detected by spectral EEG power" (in record 25726786_en). This is probably because the general-purpose LLM model (Mixtral 8x7B instruct) trained on the general text does not have the domain knowledge of the biomedical field. Therefore it cannot distinguish the subtle differences among the BioNNE entity types. We can fine-tune the LLM model on the BioNNE training data as well as other biomedical NER datasets to improve the performance of the model in this task.

Although the Mixtral 8x7B model does not officially support the Russian language, a recent study by Fenogenova et al. [15] shows that Mistral 7B (a smaller version of Mixtral 8x7B) performs well on Russian language benchmarks compared to other LLMs. Future work can explore using Mixtral or other LLM models to solve the BioNNE bilingual NER problems using prompt engineering techniques in the paper or fine-tuning the model with bilingual datasets.

6. Conclusions

We build a system that uses a general-purpose LLM, a biomedical domain NER model and UMLS-based heuristics to extract nested named entities from biomedical text. Our model achieved an F1 score of 0.39 and 0.348 on the BioNNE English validation and test sets. Although the results are not comparable to those of the state-of-the-art models [3], that were trained on the domain-specific dataset, our study

demonstrates the potential of using general-purpose LLM and prompt engineering with domain-specific rules to solve biomedical NER tasks.

Acknowledgements

We want to thank the Data Science @ Georgia Tech (DS@GT) CLEF team for their support and acknowledge the use of Grammarly [16] to proofread this paper.

References

- [1] V. Davydova, N. Loukachevitch, E. Tutubalina, Overview of bionne task on biomedical nested named entity recognition at bioasq 2024, 2024.
- [2] A. Nentidis, G. Katsimpras, A. Krithara, S. Lima-López, E. Farré-Maduell, M. Krallinger, N. Loukachevitch, V. Davydova, E. Tutubalina, G. Paliouras, Overview of BioASQ 2024: The twelfth BioASQ challenge on Large-Scale Biomedical Semantic Indexing and Question Answering, in: L. Goeriot, P. Mulhem, G. Quénot, D. Schwab, L. Soulier, G. Maria Di Nunzio, P. Galuščáková, A. García Seco de Herrera, G. Faggioli, N. Ferro (Eds.), *Experimental IR Meets Multilinguality, Multimodality, and Interaction. Proceedings of the Fifteenth International Conference of the CLEF Association (CLEF 2024)*, 2024.
- [3] N. Loukachevitch, S. Manandhar, E. Baral, I. Rozhkov, P. Braslavski, V. Ivanov, T. Batura, E. Tutubalina, NEREL-BIO: A Dataset of Biomedical Abstracts Annotated with Nested Named Entities, *Bioinformatics* (2023). URL: <https://doi.org/10.1093/bioinformatics/btad161>. doi:10.1093/bioinformatics/btad161, btad161.
- [4] A. Q. Jiang, A. Sablayrolles, A. Roux, A. Mensch, B. Savary, C. Bamford, D. S. Chaplot, D. de Las Casas, E. B. Hanna, F. Bressand, G. Lengyel, G. Bour, G. Lample, L. R. Lavaud, L. Saulnier, M.-A. Lachaux, P. Stock, S. Subramanian, S. Yang, S. Antoniak, T. L. Scao, T. Gervet, T. Lavril, T. Wang, T. Lacroix, W. E. Sayed, Mixtral of experts, *ArXiv abs/2401.04088* (2024). URL: <https://api.semanticscholar.org/CorpusID:266844877>.
- [5] X. Li, J. Feng, Y. Meng, Q. Han, F. Wu, J. Li, A unified MRC framework for named entity recognition, in: D. Jurafsky, J. Chai, N. Schlueter, J. Tetreault (Eds.), *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, Association for Computational Linguistics, Online, 2020, pp. 5849–5859. URL: <https://aclanthology.org/2020.acl-main.519>. doi:10.18653/v1/2020.acl-main.519.
- [6] T. Shibuya, E. Hovy, Nested named entity recognition via second-best sequence learning and decoding, *Transactions of the Association for Computational Linguistics* 8 (2020) 605–620. URL: <https://aclanthology.org/2020.tacl-1.39>. doi:10.1162/tacl_a_00334.
- [7] J. White, Q. Fu, S. Hays, M. Sandborn, C. Olea, H. Gilbert, A. Elnashar, J. Spencer-Smith, D. Schmidt, A prompt pattern catalog to enhance prompt engineering with chatgpt, 2023. doi:10.48550/arXiv.2302.11382.
- [8] Q. Chen, H. Sun, H. Liu, Y. Jiang, T. Ran, X. Jin, X. Xiao, Z. Lin, H. Chen, Z. Niu, An extensive benchmark study on biomedical text generation and mining with ChatGPT, *Bioinformatics* 39 (2023) btad557. URL: <https://doi.org/10.1093/bioinformatics/btad557>. doi:10.1093/bioinformatics/btad557.
- [9] S. Tian, Q. Jin, L. Yeganova, P.-T. Lai, Q. Zhu, X. Chen, Y. Yang, Q. Chen, W. Kim, D. C. Comeau, R. Islamaj, A. Kapoor, X. Gao, Z. Lu, Opportunities and challenges for ChatGPT and large language models in biomedicine and health, *Briefings in Bioinformatics* 25 (2024) bbad493. URL: <https://doi.org/10.1093/bib/bbad493>. doi:10.1093/bib/bbad493.
- [10] OpenAI, Gpt-3.5 turbo, OpenAI (2023). URL: <https://platform.openai.com/docs/models/gpt-3-5-turbo>.
- [11] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal,

- Llama open and efficient foundation language models, arXiv preprint arXiv:2302.13971 (2023). URL: <https://arxiv.org/abs/2302.13971>.
- [12] M. Neumann, D. King, I. Beltagy, W. Ammar, ScispaCy: Fast and Robust Models for Biomedical Natural Language Processing (2019) 319–327. URL: <https://www.aclweb.org/anthology/W19-5034>. doi:10.18653/v1/W19-5034. arXiv:arXiv:1902.07669.
- [13] O. Bodenreider, The unified medical language system (umls): integrating biomedical terminology, *Nucleic acids research* 32 (2004) D267–D270. URL: <https://doi.org/10.1093/nar/gkh061>. doi:10.1093/nar/gkh061.
- [14] Umls terminology services (2023). URL: <https://uts.nlm.nih.gov/uts/umls/home>.
- [15] A. Fenogenova, A. Chervyakov, N. Martynov, A. Kozlova, M. Tikhonova, A. Akhmetgareeva, A. Emelyanov, D. Shevelev, P. Lebedev, L. Sinev, U. Isaeva, K. Kolomeytseva, D. Moskovskiy, E. Goncharova, N. Savushkin, P. Mikhailova, D. Dimitrov, A. Panchenko, S. Markov, Mera: A comprehensive llm evaluation in russian, 2024. URL: <https://arxiv.org/abs/2401.04531>. arXiv:2401.04531.
- [16] Grammarly, Grammarly handbook (2024). URL: <https://www.grammarly.com/handbook/>.