

Evaluating Named Entity Recognition Using Few-Shot Prompting with Large Language Models

Hédi Zeghidi and Ludovic Moncla

INSA Lyon, CNRS, Université Claude Bernard Lyon 1,
LIRIS, UMR5205, 69621 Villeurbanne, France

Correspondence: ludovic.moncla@insa-lyon.fr

Abstract

This paper evaluates Few-Shot Prompting with Large Language Models for Named Entity Recognition (NER). Traditional NER systems rely on extensive labeled datasets, which are costly and time-consuming to obtain. Few-Shot Prompting or in-context learning enables models to recognize entities with minimal examples. We assess state-of-the-art models like GPT-4 in NER tasks, comparing their few-shot performance to fully supervised benchmarks. Results show that while there is a performance gap, large models excel in adapting to new entity types and domains with very limited data. We also explore the effects of prompt engineering, guided output format and context length on performance. This study underscores Few-Shot Learning’s potential to reduce the need for large labeled datasets, enhancing NER scalability and accessibility.

1 Introduction

Named Entity Recognition (NER) is a fundamental task in Natural Language Processing (NLP), involving the identification and classification of entities such as names, locations, and dates within text. Traditional NER systems typically require large, annotated datasets for effective training, which can be expensive and time-consuming to curate. This limitation hinders the rapid deployment and scalability of NER applications across diverse domains and languages. Few-Shot Learning (FSL) offers a promising alternative by enabling models to perform NER with minimal labeled examples. Leveraging the capabilities of Large Language Models (LLMs), FSL can significantly reduce the dependency on extensive annotated data. This paper investigates the performance of LLMs in NER tasks using FSL, comparing their efficacy to conventional fully supervised methods. Through this evaluation, we aim to highlight the potential of FSL to transform the landscape of NER, making it more accessible and scalable.

Despite their impressive performance in many NLP tasks, NER remains challenging for LLMs (Lu et al., 2024; González-Gallardo et al., 2024). Mainly because NER is inherently a sequence labeling task, whereas LLMs are primarily designed for text generation. Existing methods utilizing LLMs typically extract only a list of entities from text along with their respective classes. We aim to improve the process by identifying each entity’s position in the text and outputting it in JSON format, including both token and character positions.

This study aims to answer the following question: How effective are FSL techniques in LLMs for performing NER tasks? Our experiments are performed using the GeoEDdA dataset¹ (Moncla et al., 2024) and several LLMs are evaluated and compared to BERT-like model (Devlin et al., 2018) baselines.

2 Related works

Some works study the zero-shot capabilities of GPT for NER. For instance, Xie et al. (2023) proposed the NER task decomposition into a set of simpler subproblems by labels and perform a decomposed-question-answering task in addition to a syntactic augmentation. However, most of works are investigating FSL. Li and Zhang (2023) investigate the performances of LLMs and smaller pre-trained language models (e.g., BERT, T5) for domain-specific NER. Ashok and Lipton (2023) propose to prompts an LLM to produce a list of potential entities along with corresponding explanations justifying their compatibility with the provided entity type definitions. Wang et al. (2023) propose adding special tokens for marking the entities to extract. The evaluation shows that this method achieve comparable results as supervised methods.

For newly models with expanded context windows, (Agarwal et al., 2024) propose an experimen-

¹<https://huggingface.co/datasets/GEODE/GeoEDdA>

tal study comparing FSL versus many-shot learning (hundreds or thousands of examples). Results indicate significant performance improvements when increasing the number of provided examples from very few to many. Results also indicate that the order of examples in the prompt affects many-shot performance and adding more examples than optimal can sometimes degrade performance for certain tasks. Additionally, providing only questions without answers in the examples works well in some configurations. These observations are model-dependent, based on tests conducted with Gemini 1.5, GPT-4, and Claude 3.5.

3 Methodology

The objective of this preliminary work is to evaluate the capabilities of LLMs on the NER task at the token and span (or entity) levels. For these two tasks, the output must be a JSON format with information for each detected token or span. A span refer to an entity and can be composed of several tokens.

Our methodology relies on a prompt containing descriptions of the task, tagset and one example (with both input and output). A preliminary experiment reveals that LLMs struggle to accurately retrieve or calculate token or span positions from raw input text. Although the output format was correct, the generated position values were inaccurate, resembling hallucinations or random numbers rather than referring to actual token positions, despite appearing as numerical values. To address this issue, we propose a solution that includes tokenization information—specifically, token position details for span detection—within the input data.

4 Experiments and Results

The experiments² are performed using the GeoEDdA dataset (Moncla et al., 2024) which contains semantic annotations (at the token and span levels) for named entities (i.e., Spatial, Person, and Misc), nominal entities, spatial relations, and geographic coordinates. Nested named entities (Gaio and Moncla, 2017) also present in this dataset were not considered in this experiment. One example³ from the training set containing at least one entity of each class is used for FSL (included in the

²Source code publicly available at <https://github.com/GEODE-project/ner-llm>

³<https://huggingface.co/datasets/GEODE/GeoEDdA/viewer/default/train?row=53>

Model	Precision	Recall	F1
GPT-3.5	0.81	0.36	0.50
GPT-4	0.75	0.62	0.67
GPT-4o	0.68	0.72	0.70
Fine-tuned BERT	0.93	0.94	0.93

Table 1: Token classification scores

prompt) and the entire test set (200 articles from the Diderot Encyclopedie) is used for evaluation.

GPT models from the OpenAI⁴ API (gpt-3.5-turbo-0125, gpt-4-0613 and gpt-4o-2024-05-13) were evaluated through the LangChain Python framework. Token-level scores are shown in Table 1. Micro average precision, recall and F1-score are used as evaluation metrics in a strict matching set (exact boundary matching). While the performance is below that of a traditionally fine-tuned BERT model⁵ (fine-tuned using the fully supervised training set containing 1,800 encyclopaedia entries), experiments have shown significant variation across the different classes. At the span level, some unusual issues (compared to traditional NER) arise as the output format is more complex than at the token level. The experiment confirmed the performance improvement between GPT versions 4 and 3.5. Using GPT-3.5, 28% of the predicted spans only are correct (both boundaries and labels) compared to 49% with GPT-4o, while 13% are partially correct (partial boundaries and correct labels) with GPT-3.5 and 9% with GPT-4o Also, using GPT-3.5, some answers do not refer to the input document but rather to the example from the prompt. In very few cases, some answers do not strictly follow the predefined JSON output format and some token attributes may be missing (such as start, end, text, or label).

We also experimented smaller and local LLMs using LM Studio⁶ such as Phi3 (Microsoft), Gemma (Google), Mistral (MistralAI), Qwen (Qwen Team, affiliated with Alibaba Group), and Llama (Meta). The models were executed on an Nvidia RTX 3500 ADA GPU and obtained very varying results. Some LLMs provide the correct JSON output syntax but don't really understand the defined set of labels. Others completely don't understand the task and do anything or simply repeat the input sentence.

⁴<https://openai.com>

⁵<https://huggingface.co/GEODE/bert-base-french-cased-edda-ner>

⁶<https://lmstudio.ai>

Acknowledgments

The authors are grateful to the ASLAN project (ANR-10-LABX-0081) of the Université de Lyon, for its financial support within the French program "Investments for the Future" operated by the National Research Agency (ANR).

2023 Conference on Empirical Methods in Natural Language Processing, pages 7935–7956, Singapore.

References

- Rishabh Agarwal, Avi Singh, Lei M Zhang, Bernd Bohnet, Stephanie Chan, Ankesh Anand, Zaheer Abbas, Azade Nova, John D Co-Reyes, Eric Chu, et al. 2024. [Many-shot in-context learning](#). *arXiv preprint arXiv:2404.11018*.
- Dhananjay Ashok and Zachary C Lipton. 2023. [Promptner: Prompting for named entity recognition](#). *arXiv preprint arXiv:2305.15444*.
- 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*.
- Mauro Gaio and Ludovic Moncla. 2017. Extended named entity recognition using finite-state transducers: An application to place names. In *The ninth international conference on advanced geographic information systems, applications, and services*, Nice, France.
- Carlos-Emiliano González-Gallardo, Tran Thi Hong Hanh, Ahmed Hamdi, and Antoine Doucet. 2024. Leveraging open large language models for historical named entity recognition. In *The 28th International Conference on Theory and Practice of Digital Libraries*, Ljubljana, Slovenia.
- Mingchen Li and Rui Zhang. 2023. [How far is language model from 100% few-shot named entity recognition in medical domain](#). *arXiv preprint arXiv:2307.00186*.
- Qiu hao Lu, Rui Li, Andrew Wen, Jinlian Wang, Liwei Wang, and Hongfang Liu. 2024. [Large language models struggle in token-level clinical named entity recognition](#). *arXiv preprint arXiv:2407.00731*.
- Ludovic Moncla, Denis Vigier, and Katherine McDonough. 2024. GeoEDdA: A Gold Standard Dataset for Geo-semantic Annotation of Diderot & d’Alembert’s Encyclopédie. In *Second International Workshop on Geographic Information Extraction from Texts (GeoExT)*, Glasgow, Scotland.
- Shuhe Wang, Xiaofei Sun, Xiaoya Li, Rongbin Ouyang, Fei Wu, Tianwei Zhang, Jiwei Li, and Guoyin Wang. 2023. [Gpt-ner: Named entity recognition via large language models](#). *arXiv preprint arXiv:2304.10428*.
- Tingyu Xie, Qi Li, Jian Zhang, Yan Zhang, Zuozhu Liu, and Hongwei Wang. 2023. [Empirical study of zero-shot ner with chatgpt](#). In *Proceedings of the*