

# Improving Few-Shot Cross-Domain Named Entity Recognition by Instruction Tuning a Word-Embedding based Retrieval Augmented Large Language Model

**Subhadip Nandi**

IIT Kanpur, India

subhadipn650@gmail.com

**Neeraj Agrawal**

IISc Bangalore, India

aneeraj@iisc.ac.in

## Abstract

Few-Shot Cross-Domain NER is the process of leveraging knowledge from data-rich source domains to perform entity recognition on data-scarce target domains. Most previous state-of-the-art (SOTA) approaches use pre-trained language models (PLMs) for cross-domain NER. However, these models are often domain specific. To successfully use these models for new target domains, we need to modify either the model architecture or perform model finetuning using data from the new domains. Both of these result in the creation of entirely new NER models for each target domain which is infeasible for practical scenarios. Recently, several works have attempted to use LLMs to solve Few-Shot Cross-Domain NER. However, most of these are either too expensive for practical purposes or struggle to follow LLM prompt instructions. In this paper, we propose IF-WRANER (Instruction Finetuned Word-embedding based Retrieval Augmented large language model for Named Entity Recognition), a retrieval augmented LLM, finetuned for the NER task. By virtue of the regularization techniques used during LLM finetuning and the adoption of word-level embedding over sentence-level embedding during the retrieval of in-prompt examples, IF-WRANER is able to outperform previous SOTA Few-Shot Cross-Domain NER approaches. We have demonstrated the effectiveness of our model by benchmarking its performance on the open source CrossNER dataset, on which it shows more than 2% F1 score improvement over the previous SOTA model. We have deployed the model for multiple customer care domains of an enterprise. Accurate entity prediction through IF-WRANER helps direct customers to automated workflows for the domains, thereby reducing escalations to human agents by almost 15% and leading to millions of dollars in yearly savings for the company.

## 1 Introduction

Named Entity Recognition (NER) (Chinchor and Robinson, 1997) is a key process in information extraction, designed to identify and categorize entities in natural language into predefined entity types. Due to the large variations in entities and the way they are used across domains, NER has been a challenging task in NLP. Most traditional NER models require large volumes of labelled data for training (Wang et al., 2022; Yu et al., 2020; Wang et al., 2020b; Li et al., 2022a). However, collecting large volumes of labeled data is both costly and time consuming. Therefore, we need a model that can perform NER on multiple domains with minimal labeled examples from that domain. To tackle this problem, several solutions have been proposed, which attempt to transfer knowledge from data rich source domain to perform NER on data-scarce target domain. This is referred to as Few-Shot Cross-Domain NER.

The traditional way of solving this involves training PLMs with entity-tagged source domain data, followed by fine-tuning them on target domain data, thereby transferring knowledge from source to target domain. This approach fails to address the semantic gap that may exist between source and target domains. To address this, some previous studies utilize adding auxiliary objects (Liu et al., 2020a; Wang et al., 2020a) or designing new model architectures (Jia et al., 2019; Liu et al., 2020a; Jia and Zhang, 2020) to train with both source and target domain data. Liu et al. leverages continued pretraining with target domain data for a better understanding of domain-specific data. Another line of research (Zheng et al., 2022; Hu et al., 2022) focuses on modeling the label relationship across domains to improve label information transfer. Specifically, LANER (Hu et al., 2022) utilizes an architecture to better leverage semantic relationships among domain labels to increase

cross-domain performance.

Most Few-Shot Cross-Domain NER models however, have one or both of the following weaknesses:

- Models like LNER (Hu et al., 2022), have architectures that are very specific to the source and target domain pairs. Making these models work well for a new target domain requires tweaking the architecture.
- Other approaches require finetuning of model on target domain data. This is not feasible in real life scenarios due to computational resource and time crunch.

Our approach involves a single model architecture, finetuned only using entity tagged source domain data. For adaptation to target domain, our model does not require further fine-tuning.

Recently, with the advent of generative AI, many researchers have tried to use LLMs to solve the Few Shot Cross-Domain NER problem. GPT-NER (Wang et al., 2023) and PromptNER (Ashok and Lipton, 2023) have experimented with different LLM prompting strategies for the task with varying degrees of success. GPT-NER further demonstrates that using the Retrieval Augmented Generation (RAG) (Lewis et al., 2020) framework to select the in-prompt examples further boosts NER performance. One common theme that has emerged from these works is that most of these approaches demonstrate good performance with GPT4 as the backbone LLM. Open source alternatives typically fall well short of SOTA performance as they do not seem to closely follow the prompt instructions. This is a serious problem for real world scenarios as the use of proprietary software like GPT4 can be cost-prohibitive, especially for applications operating at scale. In our approach, we finetune open source LLMs (Touvron et al., 2023), so that they can follow domain specific prompt instructions for the NER task.

Like GPT-NER (Wang et al., 2023), we too utilise the RAG framework for selecting in-prompt examples. We store labelled domain data and their vector embeddings in a vector store and extract relevant domain examples from the store during inference based on similarity between the inference query embedding and the embeddings stored. Most applications using the RAG framework use sentence level embeddings for similarity score calculations. In our work, we show that for the NER task, retrieving examples based on word-level em-

bedding similarity performs much better than that based on sentence-level embedding similarity.

## 2 Background and Related Work

Traditional approaches to solve the NER problem typically fall into one of the two categories:

- BERT-based (Devlin et al., 2018) models like BERT-Tagger, introduced by Ding et al., 2021 are built by adding a linear classifier on top of BERT such that each token in the sentence is classified into one of the pre-defined entity types. These models are generally trained with a cross-entropy objective.
- Meta-learning approaches like ProtoBERT (Snell et al., 2017), NN-Shot (Yang and Katiyar, 2020) and Struct shot (Yang and Katiyar, 2020) derive a prototype for each entity type by computing the average of the contextual embeddings of the tokens that share the same entity type. Nearest neighbour algorithms are then used to classify each token of an input sentence into one of the entity types based on similarity between the token embedding and the embeddings of the prototypes.

While meta-learning based techniques fare better than BERT based models in the few-shot setting, they still require substantial amounts of data for creating representative prototypes.

Cross-domain Named Entity Recognition (NER) algorithms (Lin and Lu, 2018; Yang et al., 2018) help to address the issue of data scarcity in the target domain by leveraging data from source domains. One commonly used strategy to tackle Few-Shot Cross-Domain NER is multitask learning, which involves the use of auxiliary objects (Liu et al., 2020a; Wang et al., 2020a) or the development of fresh model structures (Jia et al., 2019; Liu et al., 2020a; Jia and Zhang, 2020). These methods aim to enhance the NER performance in the target domain by training on data from both the source and target domains. Another facet of cross-domain NER research concentrates on the transfer of label information across domains. As noted by Zheng et al., 2022, the relationship of labels can be represented as a probability distribution to facilitate the transfer of cross-domain knowledge in NER more effectively. Hu et al., 2022 suggests a method to capitalise on the semantic relationships between domains more efficiently by utilising previous labels (from the source) and the corresponding token. Chen et al., 2023 introduces collaborative prefix

tuning as a solution to the cross-domain NER issue. Although prefix-tuning is considerably quicker than complete fine-tuning, it still requires the addition and modification of new model parameters. Unlike the above approaches, IF-WRANER does not necessitate architectural alterations or model fine-tuning/prefix-tuning to adapt to new target domains, making it more suitable for practical applications.

Large language models (LLMs) (Mann et al., 2020; Hoffmann et al., 2022) have demonstrated remarkable proficiency in in-context learning, where they can generate results for a new test input using only a handful of task-specific examples. Operating under the in-context learning framework, LLMs have yielded encouraging outcomes across a range of NLP tasks, including Machine Translation (MT) (Vilar et al., 2022, Moslem et al., 2023), Question Answering (QA) (Robinson et al., 2022, Li et al., 2022b) etc. For Few-Shot Cross Domain NER, PromptNER (Ashok and Lipton, 2023) and GPT-NER (Wang et al., 2023) have utilized LLMs, achieving a performance level comparable to the industry benchmark through extensive prompt engineering. GPT-NER has redefined the NER problem from a sequence labeling task to a generation task that LLMs can easily adapt to. On the other hand, PromptNER uses the Chain-of-Thought Prompting technique, which offers a precise, adaptable, and user-friendly way to carry out Few-Shot NER and requires prompting the LLM only once. In our approach, we instruct the LLM to provide responses in a structured format and both the extraction of entities and their categorization into entity types happens in a single LLM call.

Another recent innovation that has gained tremendous popularity and adoption post the advent of LLMs is Retrieval-Augmented Generation (RAG) (Lewis et al., 2020) with LLMs. For Few-Shot Cross-Domain NER, we show that using proprietary LLM like GPT4 with the RAG framework we are able to obtain results comparable to SOTA models. However, given the cost implications of using GPT4, we have developed a strategy to fine-tune open-source LLMs. With our approach of finetuning open-source LLM, coupled with regularization techniques and replacing the similarity between sentence level embeddings with similarity between word level embedding as criteria for selecting relevant examples using the retriever, we are able to achieve better performance on most domains compared to previous SOTA models.

## 3 Methodology

### 3.1 Problem Definition

Given a list of predefined entity types for a domain, and a sentence, the Named Entity Recognition (NER) task involves identifying sequences of words in the sentence as entities and categorizing them into correct entity types. With Few-Shot Cross Domain NER, we have the added restriction that the number of labeled examples for the domain in question (target domain) is small. However, we have sufficient labeled examples from another domain (source domain) which we can leverage for model building.

### 3.2 Prompting LLM for NER

NER has historically been viewed as a sequence labeling task that assigns an entity type (partial assignment in-case of multi-word entities) to each word in a given sentence. With the advent of LLMs, many studies have tried to reformulate NER as a text generation task instead of a sequence labelling task. The format of the generated text can vary widely but can be broken into two categories at a high level.

- The output is a dictionary with candidate entities (sequences of words in the sentence) as keys and their corresponding entity types as values.
- The output is a dictionary with all the entity types as keys and their corresponding entities (sequences of words in the sentence) as values.

We saw greater success with the second approach and decided to adopt it for our subsequent experiments. Most approaches that leverage LLMs to solve the NER task follow the general paradigm of in-context learning that includes prompt construction, followed by feeding it to the LLM, which then produces output in the format described in the prompt. The prompts for our experiments have the following format:

- **Task Description:** Explains the task to the LLM which in our case looks like: “You are a smart and intelligent Named Entity Recognition (NER) system. You will be provided with the definition of the entities to extract, the sentence from which to extract the entities and the format in which you are to display the output...”
- **Entity Definitions:** Contains list of entity types for a particular domain and their respective definitions.
- **Input Output Examples:** Top k examples of input and output pairs from corresponding do-

main data such as:

Input: Can i pick this up tomorrow

Output: {product:[], time:[tomorrow] etc.}

- **User Query:** Sentence on which NER is to be performed.

### 3.3 Retrieval Augmented Generation (RAG)

With Retrieval Augmented Generation (RAG), instead of having the same hardcoded domain examples appended for every query, examples are selected dynamically based on their similarity with the input query. To achieve this, embeddings for domain examples are at first computed using a pre-trained universal embedder and the examples along with their embeddings are stored in a vector database. During inference, when a query comes in, its embedding is computed using the same embedder. Then similarity scores of the query embedding with all the embeddings stored in the vector db are computed and the top k most similar examples are selected and appended to the prompt, which is then sent to the LLM for response generation. The RAG framework is largely made up of two components

- **Retriever:** It is responsible for generating query embedding during inference using the embedder and also for extracting top k most similar examples from vector DB using a similarity function.
- **Generator:** The prompt, made up of the user query and the top k examples obtained by the retriever are passed along to the generator component (in our case an LLM) which is responsible for generating a response.

The complete RAG architecture is shown in Figure 1. RAG with LLM yields good results for Few-Shot Cross-Domain NER when using proprietary LLMs like GPT4. However, most open source LLMs struggle to produce output in the format specified as part of prompt instruction. This becomes a challenge, because using GPT4 to perform NER for applications at scale can be extremely costly. Therefore, we need to finetune open source LLMs so that they can follow prompt instructions.

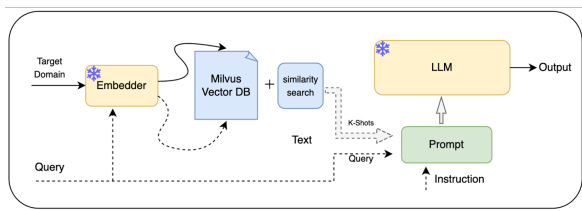


Figure 1: Retrieval Augmented Generation with LLM

### 3.4 Finetuning open-source LLMs for Few-Shot Cross-Domain NER

As per the Cross-Domain NER setting, we have a source domain which has enough entity tagged data. We finetune 7B Meta LLM on this source domain data. The purpose of this finetuning is not to teach the LLM about the source domain. Instead, finetuning accomplishes the task of teaching LLM to perform NER task and generate results in the format specified in the prompt. The prompt content is the same as that in Section 3.2. We utilise the RAG framework during finetuning also, by storing a portion (around 500 examples) of the source domain data in a vector database and finetuning with the rest of the source domain examples. More details on this can be found in Section 4.2. Cross entropy loss, computed from the LLM output and ground truth, is used to finetune the LLM parameters. We employ LoRA (Hu et al., 2021) for this. The detailed finetuning process can be found in Figure 2.

Once finetuned on source domain data, we do not need to make any further changes to the model weights to adapt to different target domains. To evaluate model performance on any target domain, we simply store the labelled examples of that domain in a vector DB, and prompt the finetuned LLM in accordance with the RAG framework to generate outputs.




Figure 2: Finetune 7B Meta LLM on source domain data

### 3.5 Training Regularization

While testing our model on target domain data, we found it to suffer from the problem of overfitting. As an example, for one of our target domains (Politics), our model tagged several politicians as “Person” instead of the entity type “Politician”, despite the prompt instruction explicitly stating that only non-politicians were to be identified as “Person”. This happens because the entity type “Person” was also part of the source domain and during finetuning, the model had memorized how to tag particular entity values as “Person”. During evaluation, the LLM chose to ignore the instruction and contin-



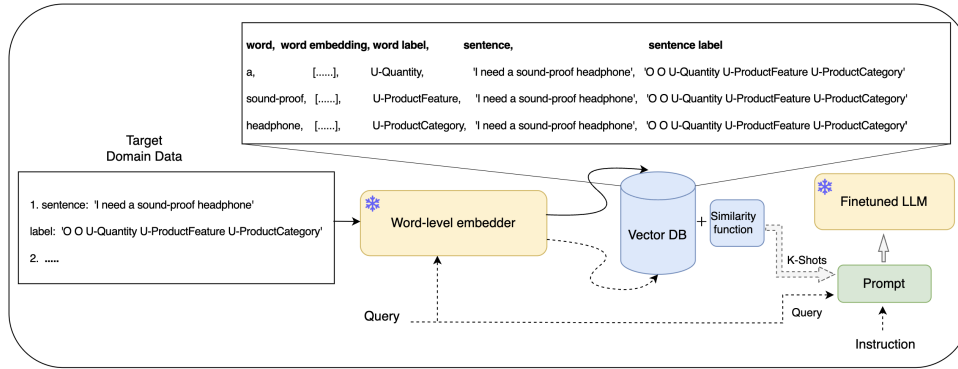


Figure 3: Using word-level embedding instead of sentence-level embedding

ued to identify politicians as “Person” as learned during finetuning. In order to alleviate this problem, we introduce various kinds of noise during model finetuning. This prevents the model from memorizing entities for the different entity types and instead teaches the model to follow prompt instructions. We applied the following regularization techniques:

- We duplicated a percentage of training examples and had some entity types randomly removed from both the input and output of those examples, which were then augmented to our training data. This ensures that the model is penalized when it predicts an entity type which is not part of the prompt, thus forcing it to learn to respect the prompt instruction.
- We randomly shuffled the order of entity types in the prompt for some examples. This prevents the model from memorizing the prompt and helps it achieve robustness against changes in the ordering of entity types in the prompt.

A comparison of the relative contribution of the regularization techniques described has been shown in Section 8.2 of Appendix.

### 3.6 Using word-level embedding instead of sentence-level embedding

NER is a word-level task that focuses more on local evidence rather than a sentence-level task, which is concerned with sentence-level semantics. Let us consider the following query sentence: “I want to buy a 13-inch macbook from store”. We have two candidate sentences for adding as example in prompt: “I want to buy a table from store” and “Show me a 15-inch macbook”. If we consider sentence-level embeddings to compute similarity, candidate 1 is closer to the query sentence. However, for the NER task, candidate 2 would be a

much better example to have in the prompt as it contains a very similar entity to the example sentence. To resolve this, we retrieve examples based on word-level representations rather than sentence-level representations. Implementing this involves the following steps:

- Obtain contextualized word embedding for every entity tagged word across all sentences in domain data. This is done by passing the sentences through an encoder model (bge-base-en (Xiao et al., 2023)). Tokens corresponding to the same word are averaged to obtain embedding for every word in the sentence.
- We store each word embedding, the word itself, the corresponding sentence and sentence label in our vector DB.
- During inference, we obtain embedding for every word in input sentence in the same manner.
- For every word, we find the top k closest matches from vector DB based on cosine similarity of the embeddings and extract the associated labelled examples. We end up with  $k \times N$  examples where  $N$  is the number of words in the input sentence (after removing stop words), along with the word-level similarity scores.
- Among these, we select k unique examples with highest scores.

The process is shown in Figure 3. The impact of using word embeddings instead of sentence embeddings for retrieval is shown in Appendix section 8.3.

## 4 Experimental Setup

### 4.1 Dataset

We have developed our model for the customer care domain of an enterprise. However, given the proprietary nature of the data, we have not shared it here. To demonstrate the generalized nature of

our model, we have also conducted experiments using the open source CrossNER (Liu et al., 2021) dataset. CrossNER contains separate datasets from five diverse domains, namely politics, natural science, music, AI, and literature. We adhere to the official splits for training, validation and test sets, the details of which can be found in Table 1. Models that report their results on the CrossNER dataset typically use a subset of the CoNLL 2003 (Sang and De Meulder, 2003) dataset as source domain, also provided as part of the CrossNER dataset. We follow the same guidelines and use the subset of CoNLL 2003 as our source domain data.

Domain	No. of train examples	No. of dev examples	No. of test examples
Reuters	14987	3466	3684
Politics	200	541	651
Natural Science	200	450	543
Music	100	380	456
Literature	100	400	416
AI	100	350	431

Table 1: Dataset Statistics - Reuters from CoNLL 2003 is used as source domain. The rest of the domains from CrossNER dataset make up the target domains

## 4.2 Experiment Details

We use a 7B Meta LLM as the open source model for our work. We have also experimented with Mistral-7B (Jiang et al., 2023), 13B Meta LLM and others, but based on the trade-off between model performance and model response latency, we found the 7B Meta LLM to be ideal for our use case. Detailed comparison of the models can be found in Appendix section 8.7.

As mentioned in Section 4.1, we use CoNLL 2003 dataset as our source domain and finetune our model using it. We randomly sample 500 examples from the training set, generate embeddings for these examples and store these embeddings along with the labelled examples in a vector DB. The rest of the training examples are used for finetuning the LLM. The validation set is used for the selection of LLM hyperparameters. While evaluating model performance on the target domains, the respective training and validation sets, along with their embeddings are stored in vector DB. Inference is performed on the test set and the data from vector DB serves as potential examples to be used in the prompt.

Since full-finetuning of LLMs is resource intensive we use the Parameter Efficient Fine-Tuning (PEFT) technique LoRA (4-bit) (Hu et al., 2021)

for finetuning our model. AdamW (Loshchilov and Hutter, 2017) optimizer with a learning rate of  $2e-4$  is used during the process. For our RAG framework, we had to choose from a plethora of options for vector DBs, embedding models and similarity metrics. We ended up using bge-base-en (Xiao et al., 2023), Milvus DB (Wang et al., 2021) and cosine similarity respectively, based on empirical results. IVF Flat indexing method is used for indexing data in the Milvus vector DB. We set the value of k (number of in-prompt examples) to 5 for our experiments based on validation dataset results. We used V100 GPU for finetuning our model on CoNLL 2003 data. The whole finetuning process takes around 50 minutes.

## 5 Results

In table 2 we have compared the performance of our model against previous SOTA Cross-Domain NER models on the 5 domains of the CrossNER dataset. Details of the previous SOTA models which we have taken as baselines for our work can be found in Section 8.1 of Appendix. IF-WRANER outperforms most of the models by a significant margin. Only PromptNER with GPT 4 is close in performance to our model. PromptNER with GPT3.5 falls well short of SOTA performance. Among the non-LLM approaches, CP-NER performs the best. We use micro F1-score for performance comparison, the most common metric for evaluating NER models. (Ma and Hovy, 2016; Lample et al., 2016).

## 6 Model Deployment

With IF-WRANER, we have built a model that new domains can use off the shelf, simply by adding entity type definitions and a few labelled examples from the respective domains. We use the tensorrt framework on Triton Inference Server (Tillet et al., 2019) for serving our model. Depending on the traffic and latency requirements for each domain, we create separate instances of our model and serve them with Triton.

Using IF-WRANER we are able to achieve reasonable latency and throughput numbers on A100 GPUs. Some domains however, have very low latency requirements and as per our experiments, a 7B-parameter IF-WRANER cannot meet these latency requirements. For such domains, we create a new model with Tinyllama (Zhang et al., 2024) as the base LLM. Tinyllama is a 1.1B model with the same architecture and tokenizer as 7B Meta

Model	Politics	Natural Science	Music	Literature	AI	Average
BiLSTM-CRF (Lample et al., 2016)	56.60	49.97	44.79	43.03	43.56	47.59
Coach (Liu et al., 2020b)	61.50	52.09	51.66	48.35	45.15	51.75
CROSS-DOMAIN LM (Jia et al., 2019)	68.44	64.31	63.56	59.59	53.70	61.92
FLAIR (Akbik et al., 2018)	69.54	64.71	65.60	61.35	52.48	62.73
BARTNER (Yan et al., 2021)	69.90	65.14	65.35	58.93	53.00	62.46
LIGHTNER (Liu et al., 2020b)	72.78	66.74	72.28	65.17	35.82	62.56
LST-NER (Zheng et al., 2022)	73.25	70.07	76.83	70.76	63.28	70.84
LANER (Hu et al., 2022)	74.06	71.83	78.78	71.11	65.79	72.31
CP-NER (Chen et al., 2023)	74.25	75.82	79.10	72.17	67.95	73.86
GPT-NER (Wang et al., 2023)	74.71	70.77	78.30	62.18	66.07	70.41
PromptNER (GPT3.5) (Ashok and Lipton, 2023)	71.74	64.83	77.78	64.15	59.35	67.57
PromptNER (GPT4) (Ashok and Lipton, 2023)	78.61	72.59	84.26	74.44	64.83	74.95
RAG + GPT4 using sentence embeddings	78.2	73.52	83.61	71.32	66.91	74.71
RAG + GPT4 using word embeddings	78.63	73.95	84.25	74.68	68.19	75.94
<b>IF-WRANER (ours)</b>	<b>79.8</b>	<b>75.31</b>	<b>85.43</b>	<b>75.52</b>	<b>68.81</b>	<b>76.97</b>

Table 2: Comparisons of previous SOTA models for Cross-Domain NER and IF-WRANER in terms of F1 scores(%) are provided. The Average indicates the average F1 score across five domains in the CrossNER benchmark

Model	Performance			Latency (s)	QPS	Cost/month (\$)
	CrossNER	Domain A	Domain B			
IF-WRANER (ours)	76.97	83.72	79.95	2X	1X	1X
Tiny-IF-WRANER (ours)	73.62	79.64	76.46	1X	1X	1X
RAG with GPT 4 (our implementation)	74.71	80.95	78.04	4X	1X	120X
PromptNER with GPT4(Ashok and Lipton, 2023)	74.95	81.15	78.12	4.2X	1X	120X

Table 3: Comparison of model performance, latency, throughput and cost for IF-WRANER, Tiny-IF-WRANER, RAG with GPT4 and PromptNER. F1 score(%) is used as model performance metric. Latency, throughput and cost are expressed in seconds(s), queries per second(QPS) and USD respectively. The cost for IF-WRANER and Tiny-IF-WRANER is the cost of using A100 GPUs while the cost of PromptNER is the cost of calling GPT4 openai endpoint

Domain	Data size in vector DB	Test Data Size	Number of entity types
Domain A	200	400	8
Domain B	230	500	15

Table 4: Characteristics of proprietary datasets

LLM, pretrained on 3 trillion tokens. We finetune Tinyllama in exactly the same way as before. This finetuned Tinyllama, Tiny-IF-WRANER, is able to serve the domains with very low latency requirement. As expected, due to its smaller model size, Tiny-IF-WRANER suffers from a drop in F1-score. In Table 3, we have compared our models (IF-WRANER and Tiny-IF-WRANER) with GPT4 based models in terms of performance, latency, throughput and cost. We see that our models are able to serve the same throughput at much lower latencies and cost. The domain with low latency requirement is represented as domain A. We use tiny-IF-WRANER to serve its users. Domain B, without such a requirement, uses IF-WRANER. Both domain A and domain B are customer care domains of an e-commerce enterprise. Due to the

proprietary nature of the domains, we cannot make their datasets available. However, we have shared some characteristics of the datasets in Table 4.

## 7 Conclusion

In this work, we have introduced IF-WRANER, a retrieval augmented instruction following LLM, that outperforms SOTA models for Few-Shot Cross-Domain NER. Unlike many of the models developed for Cross-Domain NER, we do not need to finetune or make structural modifications to our model to adapt to new domains. Also, IF-WRANER manages to attain SOTA performance using non-proprietary LLM, making it much more cost effective compared to proprietary LLM based Cross-Domain NER models. Our model is flexible and can be easily used by end users with no technical expertise. All they have to do is provide definitions for their domain’s entity types and a few labelled examples. For serving domains with very low latency requirements, we have proposed tiny-IF-WRANER which uses Tinyllama instead of 7B Meta LLM as its base LLM.

## References

- Introducing mpt-7b: A new standard for open-source, commercially usable llms. <https://www.databricks.com/blog/mpt-7b>.
- Alan Akbik, Duncan Blythe, and Roland Vollgraf. 2018. Contextual string embeddings for sequence labeling. In *Proceedings of the 27th international conference on computational linguistics*, pages 1638–1649.
- Dhananjay Ashok and Zachary C Lipton. 2023. Promptner: Prompting for named entity recognition. *arXiv preprint arXiv:2305.15444*.
- Xiang Chen, Lei Li, Shumin Deng, Chuanqi Tan, Changliang Xu, Fei Huang, Luo Si, Huajun Chen, and Ningyu Zhang. 2021. Lightner: A lightweight tuning paradigm for low-resource ner via pluggable prompting. *arXiv preprint arXiv:2109.00720*.
- Xiang Chen, Lei Li, Shuofei Qiao, Ningyu Zhang, Chuanqi Tan, Yong Jiang, Fei Huang, and Huajun Chen. 2023. One model for all domains: collaborative domain-prefix tuning for cross-domain ner. *arXiv preprint arXiv:2301.10410*.
- Nancy Chinchor and Patricia Robinson. 1997. Muc-7 named entity task definition. In *Proceedings of the 7th Conference on Message Understanding*, volume 29, pages 1–21.
- 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*.
- Ning Ding, Guangwei Xu, Yulin Chen, Xiaobin Wang, Xu Han, Pengjun Xie, Hai-Tao Zheng, and Zhiyuan Liu. 2021. Few-nerd: A few-shot named entity recognition dataset. *arXiv preprint arXiv:2105.07464*.
- Matthijs Douze, Alexandr Guzhva, Chengqi Deng, Jeff Johnson, Gergely Szilvasy, Pierre-Emmanuel Mazaré, Maria Lomeli, Lucas Hosseini, and Hervé Jégou. 2024. The faiss library. *arXiv preprint arXiv:2401.08281*.
- Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. 2022. Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Jinpeng Hu, He Zhao, Dan Guo, Xiang Wan, and Tsung-Hui Chang. 2022. A label-aware autoregressive framework for cross-domain ner. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 2222–2232.
- Chen Jia, Xiaobo Liang, and Yue Zhang. 2019. Cross-domain ner using cross-domain language modeling. In *Proceedings of the 57th annual meeting of the association for computational linguistics*, pages 2464–2474.
- Chen Jia and Yue Zhang. 2020. Multi-cell compositional lstm for ner domain adaptation. In *Proceedings of the 58th annual meeting of the association for computational linguistics*, pages 5906–5917.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.
- Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. 2016. Neural architectures for named entity recognition. *arXiv preprint arXiv:1603.01360*.
- Teven Le Scao, Angela Fan, Christopher Akiki, Elie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, Matthias Gallé, et al. 2023. Bloom: A 176b-parameter open-access multilingual language model.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474.
- Jingye Li, Hao Fei, Jiang Liu, Shengqiong Wu, Meishan Zhang, Chong Teng, Donghong Ji, and Fei Li. 2022a. Unified named entity recognition as word-word relation classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 10965–10973.
- Junlong Li, Zhuosheng Zhang, and Hai Zhao. 2022b. Self-prompting large language models for open-domain qa. *arXiv preprint arXiv:2212.08635*.
- Bill Yuchen Lin and Wei Lu. 2018. Neural adaptation layers for cross-domain named entity recognition. *arXiv preprint arXiv:1810.06368*.
- Z Liu, F Jiang, Y Hu, C Shi, and P Fung. Ner-bert: A pre-trained model for low-resource entity tagging. arxiv 2021. *arXiv preprint arXiv:2112.00405*.
- Zihan Liu, Genta Indra Winata, and Pascale Fung. 2020a. Zero-resource cross-domain named entity recognition. *arXiv preprint arXiv:2002.05923*.
- Zihan Liu, Genta Indra Winata, Peng Xu, and Pascale Fung. 2020b. Coach: A coarse-to-fine approach for cross-domain slot filling. *arXiv preprint arXiv:2004.11727*.
- Zihan Liu, Yan Xu, Tiezheng Yu, Wenliang Dai, Ziwei Ji, Samuel Cahyawijaya, Andrea Madotto, and



- Pascale Fung. 2021. Crossner: Evaluating cross-domain named entity recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 13452–13460.
- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Xuezhe Ma and Eduard Hovy. 2016. End-to-end sequence labeling via bi-directional lstm-cnns-crf. *arXiv preprint arXiv:1603.01354*.
- Ben Mann, N Ryder, M Subbiah, J Kaplan, P Dhariwal, A Neelakantan, P Shyam, G Sastry, A Askell, S Agarwal, et al. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.
- Yasmin Moslem, Rejwanul Haque, John D Kelleher, and Andy Way. 2023. Adaptive machine translation with large language models. *arXiv preprint arXiv:2301.13294*.
- Joshua Robinson, Christopher Michael Rytting, and David Wingate. 2022. Leveraging large language models for multiple choice question answering. *arXiv preprint arXiv:2210.12353*.
- Oscar Sainz, Iker García-Ferrero, Rodrigo Agerri, Oier Lopez de Lacalle, German Rigau, and Eneko Agirre. 2023. Gollie: Annotation guidelines improve zero-shot information-extraction. *arXiv preprint arXiv:2310.03668*.
- Erik F Sang and Fien De Meulder. 2003. Introduction to the conll-2003 shared task: Language-independent named entity recognition. *arXiv preprint cs/0306050*.
- Jake Snell, Kevin Swersky, and Richard Zemel. 2017. Prototypical networks for few-shot learning. *Advances in neural information processing systems*, 30.
- Philippe Tillet, Hsiang-Tsung Kung, and David Cox. 2019. Triton: an intermediate language and compiler for tiled neural network computations. In *Proceedings of the 3rd ACM SIGPLAN International Workshop on Machine Learning and Programming Languages*, pages 10–19.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- David Vilar, Markus Freitag, Colin Cherry, Jiaming Luo, Viresh Ratnakar, and George Foster. 2022. Prompting palm for translation: Assessing strategies and performance. *arXiv preprint arXiv:2211.09102*.
- Chenguang Wang, Xiao Liu, Zui Chen, Haoyun Hong, Jie Tang, and Dawn Song. 2022. Deepstruct: Pre-training of language models for structure prediction. *arXiv preprint arXiv:2205.10475*.
- Jianguo Wang, Xiaomeng Yi, Rentong Guo, Hai Jin, Peng Xu, Shengjun Li, Xiangyu Wang, Xiangzhou Guo, Chengming Li, Xiaohai Xu, et al. 2021. Milvus: A purpose-built vector data management system. In *Proceedings of the 2021 International Conference on Management of Data*, pages 2614–2627.
- Jing Wang, Mayank Kulkarni, and Daniel Preotiuc-Pietro. 2020a. Multi-domain named entity recognition with genre-aware and agnostic inference. In *Proceedings of the 58th annual meeting of the association for computational linguistics*, pages 8476–8488.
- 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*.
- Xinyu Wang, Yong Jiang, Nguyen Bach, Tao Wang, Zhongqiang Huang, Fei Huang, and Kewei Tu. 2020b. Automated concatenation of embeddings for structured prediction. *arXiv preprint arXiv:2010.05006*.
- Shitao Xiao, Zheng Liu, Peitian Zhang, and Niklas Muennighoff. 2023. C-pack: Packaged resources to advance general chinese embedding.
- Hang Yan, Tao Gui, Junqi Dai, Qipeng Guo, Zheng Zhang, and Xipeng Qiu. 2021. A unified generative framework for various ner subtasks. *arXiv preprint arXiv:2106.01223*.
- Jie Yang, Shuailong Liang, and Yue Zhang. 2018. Design challenges and misconceptions in neural sequence labeling. *arXiv preprint arXiv:1806.04470*.
- Yi Yang and Arzoo Katiyar. 2020. Simple and effective few-shot named entity recognition with structured nearest neighbor learning. *arXiv preprint arXiv:2010.02405*.
- Juntao Yu, Bernd Bohnet, and Massimo Poesio. 2020. Named entity recognition as dependency parsing. *arXiv preprint arXiv:2005.07150*.
- Urchade Zaratiana, Nadi Tomeh, Pierre Holat, and Thierry Charnois. 2023. Gliner: Generalist model for named entity recognition using bidirectional transformer. *arXiv preprint arXiv:2311.08526*.
- Peiyuan Zhang, Guangtao Zeng, Tianduo Wang, and Wei Lu. 2024. Tinyllama: An open-source small language model. *arXiv preprint arXiv:2401.02385*.
- Junhao Zheng, Haibin Chen, and Qianli Ma. 2022. Cross-domain named entity recognition via graph matching. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2670–2680.

## 8 Appendix

### 8.1 Baselines

To evaluate the effectiveness of the proposed method, we compare it with several baselines, including:

- COACH (Liu et al., 2020b): Utilizes patterns of slot entities and combines the features for each slot entity in order to improve entity type predictions.
- CROSS-DOMAIN LM (Jia et al., 2019): Utilizes a parameter generation network to merge crossdomain language modeling with NER.
- FLAIR (Akbik et al., 2018): Utilizes the internal states of a character-level language model to generate contextual string embeddings, which are integrated into the NER model.
- BARTNER (Yan et al., 2021): Uses the pre-trained BART model to generate entity spans, treating the NER task as a sequence generation problem.
- LST-NER (Zheng et al., 2022): Models the relationship between labels as a probability distribution and builds label graphs in both the source and target label spaces for cross-domain NER tasks.
- LANER (Hu et al., 2022): Uses a new approach for cross-domain named entity recognition by utilizing an autoregressive framework to strengthen the connection between labels and tokens.
- LIGHTNER (Chen et al., 2021): Utilizes a plug-gable prompting method to improve NER performance in low-resource settings.
- CP-NER (Chen et al., 2023): Utilizes collaborative prefix tuning to learn domain-specific prefixes for flexible NER execution
- GPT-NER (Wang et al., 2023): Redefines NER from a sequence labeling task to a generation task that LLMs can perform easily.
- PromptNER (Ashok and Lipton, 2023): Uses the Chain-of-Thought Prompting to perform NER.

Recently GoLLIE (Sainz et al., 2023) and GLINER (Zaratiana et al., 2023) have demonstrated SOTA performance on CrossNER dataset for the zero shot setting. However, they have not neither evaluated their models on the few-shot setting for CrossNER nor provided an easy way to do so. Therefore, we have not included them as baselines for our work.

## 8.2 Ablation Study

We have performed an ablation study to compare the contributions of the different regularization techniques towards our model’s performance. As shown in Table 5, regularization by augmenting examples with randomly removed entity types makes the most significant contribution to our model’s

performance.

Model	F1 score (%)
base finetuned model	73.65
base finetuned model + entity types removed	76.34
base finetuned model + entity types shuffled	74.41
base finetuned model + both	76.97

Table 5: Contribution of different regularization techniques towards model performance improvement

## 8.3 Effect of replacing sentence-level embedding with word-level embedding

We also studied the effect of using word-level embedding instead of sentence-level embedding in the retrieval of top k examples on both GPT 4 as well as on our model. Results of this can be found in Table 6.

Model	F1 score (%)
IF-WRANER using sentence-level embedding	75.72
IF-WRANER using word-level embedding	76.97
RAG with GPT 4 using sentence-level embedding	74.71
RAG with GPT 4 using word-level embedding	75.94

Table 6: Effect of replacing sentence-level embedding with word-level embedding on model performance

## 8.4 Effect of changing the number of retrieved examples in LLM prompt

We also studied the effect of changing the number of retrieved examples included in the prompt to IF-WRANER. Results of this can be found in Table 7. We did not create separate deployments for testing the effect of varying number of retrieved examples. The latency numbers indicate the average time taken to complete one request to IF-WRANER on V100 GPU. Based on the trade-off between F1 score and latency we settled on 5 as the optimal number of retrieved examples to add to our LLM prompt.

## 8.5 Selecting the optimal vector DB and indexing scheme

As mentioned in the paper, for the retrieval component, we have used Milvus DB as the vector DB for storing the word-level embeddings, IVF Flat as the indexing scheme and cosine similarity search for retrieving similar examples. We compared its performance against other vector DBs/vector similarity search libraries such as FAISS (Douze et al.,

Number of retrieved examples	F1 score on CrossNER(%)	F1 score on Domain A(%)	F1 score on Domain B(%)	Median Latency (s)
1	65.32	68.14	64.98	1.6
3	74.28	76.25	72.49	1.8
<b>5</b>	<b>76.97</b>	<b>83.72</b>	<b>79.95</b>	<b>1.8</b>
10	77.01	83.68	79.42	2.2
20	76.88	84.29	79.27	2.8
30	76.92	82.16	78.49	3.2

Table 7: Effect of changing the number of retrieved examples included in LLM prompt

2024) with different indexing schemes. A more detailed view of this can be found in Table 8.

With flat indexing a direct comparison is made between embeddings whereas with IVF Flat indexing, embeddings of examples are clustered first and then the query embedding is compared with cluster embeddings. It is therefore faster. HSNW performs additional optimizations and is therefore even faster than IVF, but also shows a drop in F1 score. Performance of FAISS and Milvus are roughly similar, with Milvus showing slightly better numbers. Milvus with IVF Flat indexing provides a good tradeoff for our usecase and was therefore adopted.

Vector DB	Index	Similarity Function	F1 score on CrossNER(%)	Search latency on 200 examples (s)	Search latency on 10 <sup>5</sup> examples (s)
Milvus	Flat	cosine	77.01	0.03	0.3
<b>Milvus</b>	<b>IVF Flat</b>	<b>cosine</b>	<b>76.97</b>	<b>0.024</b>	<b>0.1</b>
FAISS	FlatL2	L2	76.99	0.04	0.3
FAISS	IVF Flat	L2	76.95	0.026	0.1
FAISS	HSNW	L2	76.22	0.015	0.09

Table 8: Comparison of different vector DBs and indexing schemes for use in the retrieval component of IF-WRANER. Search latency corresponds to median search latency.

### 8.6 Selecting the optimal embedder model

While deciding on the retriever component, we also considered different embedder models. The table below compares different open-source embedder models, each using less than 2GB memory, for our use case. Models like text-ada-embedding, which only provide sentence-level embedding and abstract away the token level embedding vectors are excluded due to our focus being only on word-level embeddings. Based on the experiments bge-base-en seems to work well across all domains.

Embedder Model	Memory requirement (GB)	CrossNER F1 score(%)	Domain A F1 score(%)	Domain B F1 score(%)
<b>bge-base-en</b>	<b>1.63</b>	<b>76.97</b>	<b>83.72</b>	<b>79.95</b>
gte-large-en	1.62	76.88	83.75	78.86
uae-large-v1	1.25	76.42	83.66	77.12

Table 9: Comparison of different embedder models for the retrieval component of IF-WRANER

### 8.7 Selecting the optimal open-source LLM

We experimented with different open-source LLMs such as 7B Meta LLM, 13B Meta LLM, Mistral-7B and so on. Based on performance and latency scores, we decided to use 7B Meta LLM for our experiments. We did not experiment with LLMs larger than 13B owing to latency and infra constraints. Details of this can be found in Table 10.

Backbone LLM	F1 score (%)	Latency on single V100 GPU (s)
<b>7B Meta LLM (Touvron et al., 2023)</b>	<b>76.97</b>	<b>1.8</b>
13B Meta LLM (Touvron et al., 2023)	77.42	2.6
Mistral-7B (Jiang et al., 2023)	77.14	1.9
Bloom-7b1 (Le Scao et al., 2023)	74.42	1.9
MPT-7B (MPT)	76.12	1.8

Table 10: Comparison of different open-source LLMs

### 8.8 Prompt Guidelines

Based on our experimentation with different kinds of prompts and analysis of the model responses, we found that IF-WRANER demonstrates its best performance when the following prompting guidelines are followed:

- Entity definitions should be clear. Having atleast one example in the definition itself helps.
- When there is ambiguity between two entity types, such as “Organization” and “Political Party” and one is a subgroup of the other, then the subgroup entity type, in this case “Political Party”, should appear first in the prompt. The model displays this behavior despite the regularization techniques applied to the model.