

NuNER: Entity Recognition Encoder Pre-training via LLM-Annotated Data

Sergei Bogdanov
NuMind

Alexandre Constantin
NuMind

Timothée Bernard
Université Paris Cité

Benoit Crabbé
Université Paris Cité

Etienne Bernard
NuMind
etienne@numind.ai

Abstract

Large Language Models (LLMs) have shown impressive abilities in data annotation, opening the way for new approaches to solve classic NLP problems. In this paper, we show how to use LLMs to create NuNER, a compact language representation model specialized in the Named Entity Recognition (NER) task. NuNER can be fine-tuned to solve downstream NER problems in a data-efficient way, outperforming similar-sized foundation models in the few-shot regime and competing with much larger LLMs. We find that the size and entity-type diversity of the pre-training dataset are key to achieving good performance. We view NuNER as a member of the broader family of *task-specific foundation models*, recently unlocked by LLMs. NuNER and NuNER’s dataset are open-sourced with MIT License¹.

1 Introduction

Named Entity Recognition (NER) — the generic task of extracting and classifying entities from text — is a core component of natural language processing, present in a variety of applications such as medical coding, financial news analysis, or legal documents parsing (Francis et al., 2019; Dozier et al., 2010). Such application typically involves solving a particular NER problem, for a particular set of entity types, thus requiring the creation of a custom model.

For the last five years, the standard procedure for creating such custom model has consisted of using a transformer encoder (Vaswani et al., 2017) pre-trained in a self-supervised way to satisfy a masked language modeling (MLM) objective, such as models of the BERT family (Devlin et al., 2019; Liu et al., 2019). This *foundation model* is then fine-tuned in a supervised way on human-annotated data, either using a simple token-classification, or a more advanced strategy (Zhang et al., 2023).

¹ NuNER and NuNERZero are available on Huggingface

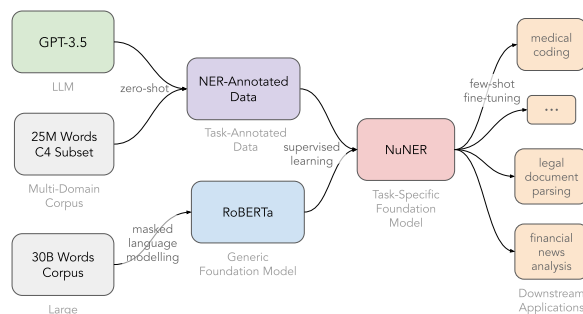


Figure 1: NuNER creation procedure. RoBERTa is further pre-trained on a subset of C4 automatically annotated by GPT-3.5. The resulting model can be fine-tuned on various downstream NER problems.

In the last few years, we have witnessed the emergence of generative large language models (LLMs) such as GPT-3 (Brown et al., 2020) and, more recently, GPT-4 (OpenAI, 2023), which typically have between 100 times and 10,000 times more parameters than BERT. These massive autoregressive transformer models, trained via a next-word prediction objective, are language generators which can be prompted to perform a variety of tasks. For example, these models can be directly used through a well crafted prompt to tackle a particular NER problem with satisfying performance (Wang et al., 2023a). The main issue with this approach is the high inference cost due to the size of LLMs.

A simple solution to this inference-cost issue is to use a correctly-prompted LLM to annotate data for the particular NER problem and then to train a smaller model on this data. LLMs have been shown to outperform crowd workers (Gilardi et al., 2023) on some tasks for a fraction of the cost but this strategy has issues as well. First, crafting a good prompt — like delegating a task to someone else — is not easy; it requires multiple back-and-forth while validating the performance using human-annotated

data. Second, LLMs are not perfect annotators either (Mao et al., 2023). Finally, the best LLMs are mainly hosted on private companies’ opening the door to potential confidentiality and privacy leaks.

We propose an alternative approach that leverages LLMs to reduce the amount of human annotations needed to create custom models. Instead of using an LLM to directly annotate a particular single-domain dataset for a particular NER problem, our idea is to use it to annotate a multi-domain dataset for variety of NER problems. We then further pre-train a small foundation model, such as BERT, on this annotated dataset. The resulting pre-trained model can then be fine-tuned to any downstream NER problem, just like any other foundation model, as depicted in Figure 1.

Because the resulting pre-trained model is specialized to a generic task but is still meant to be fine-tuned to a particular problem, we refer to such a model as a *task-specific foundation model*. Note that compact domain-specific foundation models like SciBERT (Beltagy et al., 2019) or BioBERT (Lee et al., 2019) are common, but task-specific foundation models of this kind are rare, mostly due to the lack of suitable datasets. Generative LLMs are the key to building such models.

In this paper, we apply the above idea to create NuNER, a task-specific foundation model for the generic task of NER.

In Section 3, we describe both the dataset creation and the training procedures. In a nutshell, we use GPT-3.5 to annotate a subset of C4 (Raffel et al., 2020), resulting in a 24.4M words dataset containing 4.38M annotations from 200k unique concepts. We then pre-train a base RoBERTa on this dataset via a contrastive-learning approach (Chen et al., 2020) to obtain NuNER.

In Section 4, we analyze the transfer learning performance of NuNER in an extended few-shot regime. We find that NuNER largely outperforms both its base model and the same base model further pre-trained on NER-BERT data (Liu et al., 2021), which is the largest and most diverse NER dataset we could find. These results demonstrate the validity of our approach.

In Section 5, we investigate the factors influencing NuNER’s abilities. We find that the diversity of the annotations and the size of the pre-training dataset are the most influential factors. Surprisingly, the diversity of the text does not appear to be as influential.

In Section 6, for informational purposes, we

compare fine-tuning NuNER with using GPT-3.5 and GPT-4 via in-context learning. We find that NuNER beats GPT-3.5 and competes with GPT-4 when more than a dozen entities of each type is seen during training. We also compare, via fine-tuning, NuNER with UniversalNER (Zhou et al., 2023), a recent LLM specialized in the NER task. We find that they exhibit similar transfer learning performance when fine-tuned, despite NuNER being 56 times smaller.

The contributions of our paper are as follows.

1. We introduce and demonstrate the validity of a procedure that consists of annotating raw data with an LLM in order to train a task-specific foundation model for NER.
2. We identify the factors that are likely to improve the performance of the resulting task-specific foundation model.
3. We provide and open-source NuNER², a compact encoder-based language representation model for NER. NuNER outperforms similar-sized models, competes with LLMs, and can be used as a drop-in replacement for RoBERTa.
4. We open-source an LLM-annotated NER dataset², containing 4.38M annotations from 200k entity types, which is suitable for pre-training NER models.

2 Related Work

Early attempts to create NER-specific foundation models for low-resource NER problems focused on leveraging Wikipedia anchors (Mengge et al., 2020; Cao et al., 2019). In particular, Liu et al. (2021) (NER-BERT) combined these anchors and DBpedia Ontology to create NER dataset containing 3.64M entities from 315 entity types. BERT was then pre-trained on this dataset, leading to improved few-shot performance. We demonstrate in Section 4.1 that our LLM-annotation procedure outperforms such an approach.

Subsequent work has focused on pre-training generative LLMs on a large number of existing human-annotated NER datasets to achieve strong zero-shot capabilities (Wang et al., 2023b; Sainz et al., 2023). These works differ from ours in terms of type of data used, model employed, and objective (zero-shot vs. few-shot).

Recently, Zhou et al. (2023) proposed UniversalNER, an LLM with 7B and 13B parameters, also pre-trained on data annotated by GPT-3.5. Our

² NuNER and NuNerZero are available on Huggingface

work primarily differ in the model and training procedure: we found a way to train an encoder model with only 125M parameters on such data, making it substantially cheaper to use. In Section 6, we show that NuNER and UniversalNER have similar transfer-learning abilities when provided with more than a few training examples per entity types.

Even more recently, [Zaratiana et al. \(2023\)](#) proposed GLiNER, which uses UniversalNER’s data to pre-train a small encoder model. Our work was done concurrently and independently from this work. The principal distinction is in the architecture used: GLiNER merges the text and concept encoders, whereas in our approach, they remain independent (see Figure 5). While GLiNER’s integrated approach likely enhances performance, our decision to keep the encoders separate enables NuNER to function as a concept-agnostic language representation model. This gives the possibility to pre-compute text embeddings for a variety of applications such as information retrieval. This also makes NuNER a viable drop-in substitute for BERT or RoBERTa in standard NER methodologies.

3 NuNER

The creation of NuNER is a two-step process: dataset creation and model training.

3.1 Dataset Creation

We begin with a random sample of C4 ([Raffel et al., 2020](#)), an English web crawl corpus that contains text from a wide range of sources, including blog posts, news articles, and social media messages. We selected this dataset for its domain diversity.

We want to annotate this dataset with entities spanning a large and diverse set of types in order for our model to generalize to all kind of NER problems. To achieve this we opt for an unconstrained approach: we allow the LLM to extract any entity it identifies, and give it the freedom to assign any type it deems appropriate for each entity. This includes annotating with concepts more akin to topics than entity types (e.g., “wellness”). We refer to these entity types/topics as *concepts*. To resolve potential ambiguities, we also ask the LLM to provide descriptions for the concepts it identifies. However, we ultimately disregarded these descriptions as we found that they do not improve NuNER’s performance. Our prompt is shown in Figure 2.

Note that we do not ask to return the position

```
The goal is to create a dataset for entity recognition.
Label as many entities, concepts, and ideas as possible
in the input text. Invent new entity types that may not
exist in traditional NER Tasks such as more abstract
concepts and ideas. Make sure the entity concept is
not part of speech but something more meaningful.
Avoid finding meaningless entities.
Output format (separate entities with new line):
entity from the text <> entity concept <> descrip-
tion of entity group/concept
Input:
[INPUT SENTENCE]
```

Figure 2: Prompt used to annotate NuNER’s pre-training data.

of the entity in the text as LLMs are not good at counting. To train NuNER, we have to retrieve this position through an exact string match, which can lead to annotation errors in rare cases.

We use gpt-3.5-turbo-0301 with this prompt to annotate 1.35M sentences, costing less than 500\$. Figure 3 shows one of these annotated sentences. We then apply a simple filter to remove sentences containing an annotation with the concept “concept”, as we consider this too uninformative, and obtain a final dataset of 1M annotated sentences.

“Steven Means has signed a one-year contract extension with the Falcons after making four starts in 2018.”

ENTITY	CONCEPT	DESCRIPTION
Steven Means	NFL player	professional athlete
Falcons	NFL team	professional sports team
2018	year	unit of time

Figure 3: Sentence from C4 annotated with GPT-3.5.

We find that GPT-3.5 performs quite well in this task. From a manual review of 100 examples, we estimate that, in more than 95% of cases, the concept associated with the extracted entity is completely sensible, as shown in Figure 3. In fewer than 5% of cases, the extraction and associated concept are questionable, such as “cosmos seeds” identified as “plant variety” when it is actually a seed. However, many entities are missed. For instance, in the example of Figure 3, the LLM could have also identified “person”. In a sense, this method has high precision but low recall. In Section 3.2, we propose a contrastive learning procedure that mitigates the impact of these false negatives.

The resulting dataset comprises a total of 4.38M entity annotations, distributed across 200k unique concepts. These concepts cover a wide range of

domains, as illustrated in Figure 6. The diversity of concepts in this dataset is far greater than what can be found in human-annotated NER datasets.

We observe a high imbalance in concept frequencies. Common concepts such as “person”, “location”, or “organization” each appears in more than 1% of the extracted entities, while over 100k concepts are seen only once in the dataset. This heavy-tailed distribution of concept frequencies is shown in Figure 4. We further investigate the importance of concept diversity in Section 5.2.

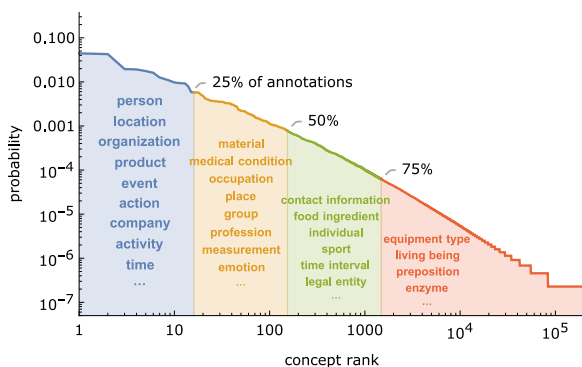


Figure 4: Frequency of each concept assigned by GPT-3.5, sorted from most to least common. We observe a heavy-tailed distribution.

3.2 Model Training

We want to pre-train a language representation model using our annotated dataset, which contains 200k concepts and exhibits a high concept-imbalance. Additionally, some concepts are similar to each other, such as “company” and “company name”. Moreover, many potential entities in a sentence are not extracted. Due to these factors, training a conventional token classifier that only takes the text as input and returns a distribution over the entire set of concepts is not practical. We instead propose a training method based on the contrastive learning framework (Chen et al., 2020).

Our training network, depicted in Figure 5, consists of two separate sub-networks: The first is NuNER — which encodes the input text as a sequence of vectors. The second encodes a concept name as a unique vector. The text vectors are matrix-multiplied with the concept vector to obtain logits, then divided by temperature and passed through a logistic sigmoid to yield probabilities. During training, this setup encourages each token embedding to align with the concept embedding if the token instantiates the concept, and to become opposite otherwise.

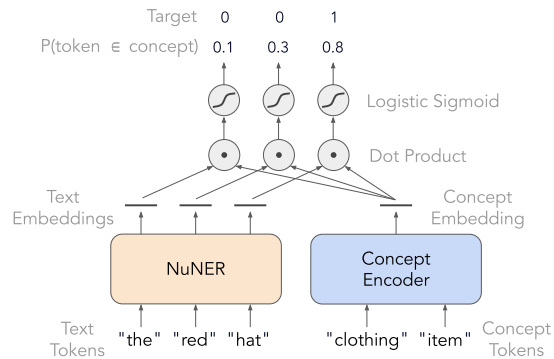


Figure 5: NuNER’s pre-training procedure. The text and concept encoder are separated. Their embeddings are compared to obtain probabilities.

For each training batch, we collect all the concepts exemplified in the batch and construct a binary array of dimensions $\#sentences \times \#tokens \times \#concepts_in_batch$. This array indicates the presence (1) or absence (0) of each concept present in the batch for every token in each sentence. We then use the binary cross-entropy loss to fit the network’s probabilities to these target arrays.

Since we do not account for concepts not present in the batch, and because our dataset misses some concepts in some sentences, the probabilities generated by such a training network would need calibration for use in a zero-shot setting. However, in our case, this miscalibration is not a significant issue as we are only interested in using the text encoder NuNER.

3.3 Training Details

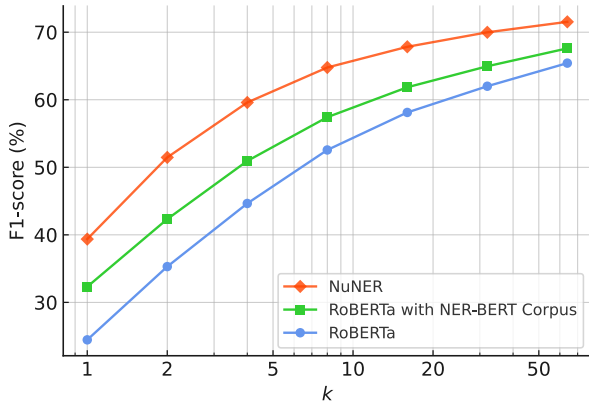
We use RoBERTa-base (Liu et al., 2019) for both the text encoder and the concept encoder. Our model is trained for 10 epochs on the full 1M sentence dataset where 90% of sentences are used for training and 10% for validation. We choose a learning rate $lr = 0.00003$, batch size= 48, and temperature before the sigmoid $\tau = 5$. We use AdamW optimizer (Kingma and Ba, 2015) with $\beta_1 = 0.9, \beta_2 = 0.999, \epsilon = 10^{-6}$, weight decay=0.01, and a linear scheduler with a warm-up for the first 10% of the training steps. The bottom 6 layers of the text encoder are frozen as we found it leads to better training stability. After training, we discard the concept encoder and keep the text encoder NuNER.

4 Transfer Learning Performance

NuNER is designed to be fine-tuned to downstream NER problems in a data-efficient way. We are

MODEL	5-WAY 1 ~ 2	10-WAY 1 ~ 2	5-WAY 5 ~ 10	10-WAY 5 ~ 10	AVG
TADNER INTRA	60.78±0.32	55.44±0.08	67.94±0.17	60.87±0.22	61.26
NUNER-BERT INTRA	62.48±0.28	57.63±0.38	69.16±0.28	62.99±0.27	63.07
TADNER INTER	64.83±0.14	64.06±0.19	72.12±0.12	69.94±0.15	67.74
NUNER-BERT INTER	67.37±0.31	66.54±0.40	73.50±0.09	71.04±0.14	69.61

Table 1: Few-NERD performance using TadNER (Li et al., 2023) and a modified TadNER using NuNER-BERT as the backbone.



MODEL	1	4	16	64
ROBERTA	24.5	44.7	58.1	65.4
ROBERTA w. NER-BERT	32.3	50.9	61.9	67.6
NUNER	39.4	59.6	67.8	71.5

Figure 7: Transfer learning performance of NuNER, RoBERTa, and RoBERTa pre-trained on NER-BERT data as function of k . NuNER substantially outperforms both models for all training sizes. Full table and dataset-wise results are shown in Table 4 and Figure 13 in the appendix.

pre-training on a large subset of Few-NERD before being evaluated using 5,000 few-shot train-test splits. In the INTRA setting, the entities seen during pre-training and evaluation belong to different coarse-grained types, while in the INTER setting, the entities share the same coarse-grained types.

The current state-of-the-art on Few-NERD is TadNER (Li et al., 2023), an advanced framework that employs a span-detection network and a type-classification network, using BERT for both networks. We adapt TadNER by replacing BERT with NuNER based on the same BERT. Furthermore, since NuNER is designed to have all its entity-related knowledge in its last layer, we modify TadNER to only use this last layer instead of averaging over the last four layers.

Results are presented in Table 1. We observe that NuNER outperforms TadNER in all settings and training sizes, establishing it as the new SOTA for this benchmark. This further hints at the benefits of using NuNER’s pre-training procedure.

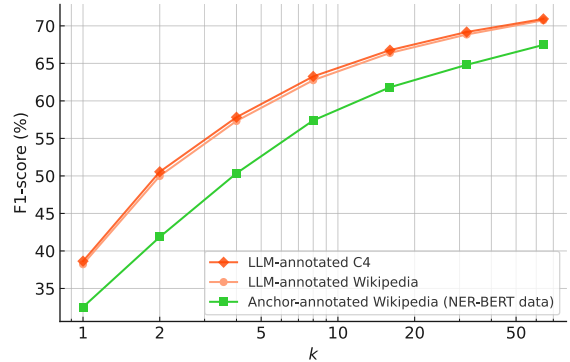


Figure 8: Effect of text diversity on NuNER’s performance. Wikipedia and C4 lead to similar performance when they are both annotated by the LLM. Results table is shown in Table 5 in the appendix.

5 Ablation Studies

We aim to understand which factors in the pre-training procedure most significantly affect the performance of NuNER. To this end, we investigate the impact of text diversity, concept diversity, pre-training dataset size, and model size Appendix A.1, using the same benchmark as in Section 4.1. We also study NuNER’s multilingual performance in Appendix A.2.

5.1 Effect of Text Diversity

In Section 4.1, we saw that NuNER’s pre-training data, based on C4, leads to better performance than NER-BERT’s pre-training data, based on Wikipedia. This might simply be because C4 is more diverse than Wikipedia.

To investigate this, we downsample both datasets to 50k sentences each, and annotate the Wikipedia subset with our LLM-annotation procedure. This way, the only distinguishing factor between these datasets is their original corpus: C4 vs. Wikipedia. We then pre-train RoBERTa on these two datasets and measure the transfer learning performance of the resulting models.

We see in Figure 8 that the LLM-annotated C4 subset and the LLM-annotated Wikipedia subset result in very similar model performance. This

shows that the main reason for the performance gap is the LLM-annotation procedure rather than the underlying corpus.

5.2 Effect of Concept Diversity

To understand the effect of concept diversity, we first take a random sample of 100k annotated examples from NuNER’s dataset, which includes approximately 80k unique concepts. We then retain only the annotations from the top- n most frequent concepts, simulating an annotation procedure that excludes rare concepts. We use $n = 4, 16, 154, 1.5k,$ and $80k$ concepts, corresponding to 12.5%, 25%, 50%, 75%, and 100% of all annotations, respectively (see Figure 4). We pre-train NuNER on the resulting datasets and measure the transfer learning performance for $k = 8$.

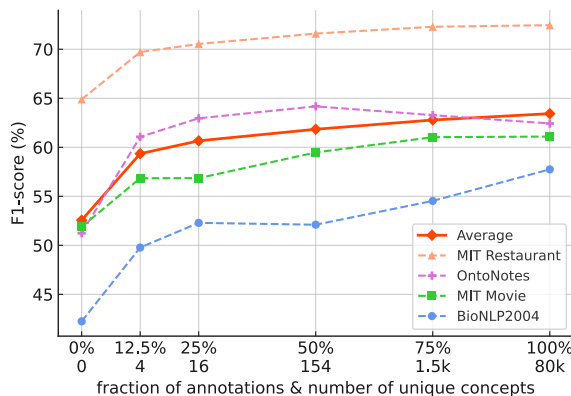


Figure 9: Effect of concept diversity on NuNER’s performance. Results table is shown in Table 6 in the appendix.

Results are shown in Figure 9. As expected, overall performance increases with concept diversity. However, there are variations across datasets. BioNLP appears to benefit the most from concept diversity while it seems to harm OntoNotes past 154 concepts. This difference is likely because BioNLP contains rarer concepts than OntoNotes. The performance degradation on OntoNotes may indicate the difficulty of encoding a large number of concepts into the 768-dimensional embedding vector.

Note that, in this experiment, the number of annotations grows with concept diversity, which might bring an additional effect. Results of Section 5.3 shows that such effect would account here for less than 1% of F1-score.

5.3 Effect of Dataset Size

We next investigate the effect of the pre-training dataset size, ranging from 1k examples to 1M examples. We again pre-train NuNER on each dataset and measure its transfer learning performance for $k = 8$. Results are shown in Figure 10.

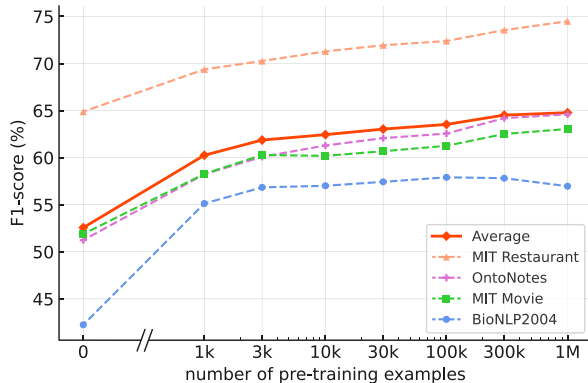


Figure 10: Effect of pre-training dataset size on NuNER’s performance. Results table is shown in Table 7 in the appendix.

As expected, the overall performance increases with data size, and continues to improve slightly from 300k to 1M examples. Again, we see variability across datasets, with the performance for BioNLP decreasing after 100k examples, while all other datasets experience a monotonic increase. The reason for this discrepancy is unclear.

6 Comparison with LLMs

Although the comparisons are indicative, we finally provide for information a comparison of NuNER with modern generative LLMs. We choose GPT-3.5 (gpt-3.5-turbo-16k-0613) for its popularity and longer context, GPT-4 (gpt-4-0613) for its high performance (Zheng et al., 2023), and UniversalNER (UniversalNER-7B-type) for its specialization in the NER task. UniversalNER has 56 times more parameters than NuNER, and GPT-3.5 and GPT-4 are likely to have around 1,000 and 10,000 times more parameters than NuNER, respectively.

GPT-4 and GPT-3.5 are used via in-context learning using Spacy’s [NER V3 prompt](#). This advanced prompt template allows to create a prompt for a particular NER problem by providing it the set of entity types and some training examples.

UniversalNER (Zhou et al., 2023) is trained to be used for zero-shot inferences, conducted through a conversation in which one sequentially prompt the model to identify each entity type. To adapt this model for a few-shot setting, we need to fine-tune

it. We use the original training settings of UniversalNER but modify them to enhance few-shot learning performance, as detailed in Appendix A.4. For NuNER, we simply attach a two-layer fully-connected network regularized via dropout, and fine-tune the entire network for 30 epochs, costing less than 5\$ in total.

Because of the high costs of GPT-4, we deviate from the extended few-shot learning protocol of Section 4.1. We only use the MIT Restaurant and BioNLP datasets, and downsample test sets to 1,000 examples. Also, we create training sets with a specific number of words belonging to a given entity type, that we call k_w , instead of using the $k \sim 2k$ entity-based mining method. We conduct several runs for each training-size (using the same training sets for all models) and average results, except for GPT-4 where we only perform one run. Results are presented in Figure 11.

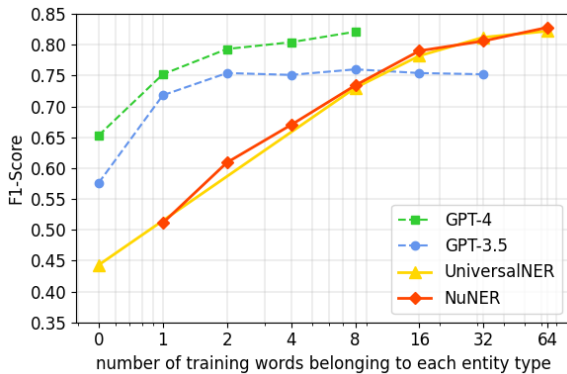


Figure 11: Comparison of NuNER with LLMs. Dashed curves indicate in-context learning and solid curves indicate fine-tuning. Results tables are shown in Table 11 and Table 12 in the appendix.

We find that GPT-3.5 and GPT-4 already perform well in the zero-shot regime, and show rapid improvement when examples are added to the prompt. However, for larger training sets, the performance of GPT-3.5 quickly plateaus, and it ends up being outperformed by both UniversalNER and NuNER when $k_w > 8$. The same thing might happen with GPT-4 but we cannot conclude with such noisy and incomplete results. UniversalNER starts lower than both GPT-3.5 and GPT-4, but steadily catches on to eventually surpass GPT-3.5.

NuNER begins at a lower performance level than GPTs as it is not intended to be a zero-shot model. However, it matches UniversalNER’s performance and eventually surpasses GPT-3.5. It remains unclear whether it would also surpass GPT-4. Nevertheless, when using GPT-4 with $k = 8$, inferring

one test example costs around 0.1\$, while NuNER achieves similar performance while costing less than 0.0001\$ per example, allowing for industrial use.

To compare NuNER and UniversalNER in a more standard and reproducible manner, we conduct an additional experiment using the $k \sim 2k$ setting on all four datasets mentioned in Section 4.1. In this experiment, we use the full test sets and measure entity-level micro F1-score, averaged over these datasets. We perform 3 runs for each k . We see in Table 2 that NuNER and UniversalNER have similar performance.

MODEL	8 ~ 16	64 ~ 128
UNIVERSALNER	57.89 ± 4.34	71.02 ± 1.53
NUNER	58.75 ± 0.93	70.30 ± 0.35

Table 2: NuNER vs. UniversalNER few-shot entity-level F1-score in the $k \sim 2k$ setting showing similar performance. Dataset-wise tables can be found in Table 9 and in Table 10 in the appendix.

The fact that NuNER surpasses GPT-3.5 and possibly also GPT-4 is likely attributable to the limitations of in-context learning. The situation might differ with proper, albeit challenging, few-shot fine-tuning for these LLMs. More surprising to us is that NuNER achieves performance comparable to UniversalNER despite being 56 times smaller and trained on similar data. This could be due to an inherent advantage of encoders over generative models for this task. Alternatively, it might be related to NuNER’s pre-training procedure, which encourages human concepts to emerge in the last layers of the network, being easily accessible during few-shot training. Further experiments would be needed to explore this aspect.

7 Conclusion

Modern large language models are opening new possibilities for addressing traditional NLP tasks. We have introduced a procedure that uses these LLMs to create a compact, yet data-efficient, NER-specific foundation model. We foresee an increasing trend in the development of such task-specific foundation models, which will facilitate the creation of high-quality custom NLP models without requiring intensive human or computational resources.

Limitations

This paper presents the NuNER model which was created to be an efficient feature extractor for the Named Entity Recognition problems. In this work, we only studied the Few-shot learning capabilities of the model, without experimenting with the potential Zero-shot modifications of NuNER.

References

- Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. [Scibert: A pretrained language model for scientific text](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 3613–3618. Association for Computational Linguistics.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Yixin Cao, Zikun Hu, Tat-seng Chua, Zhiyuan Liu, and Heng Ji. 2019. [Low-resource name tagging learned with weakly labeled data](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 261–270, Hong Kong, China. Association for Computational Linguistics.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. [A Simple Framework for Contrastive Learning of Visual Representations](#).
- Nigel Collier, Tomoko Ohta, Yoshimasa Tsuruoka, Yuka Tateisi, and Jin-Dong Kim. 2004. [Introduction to the bio-entity recognition task at JNLPBA](#). In *Proceedings of the International Joint Workshop on Natural Language Processing in Biomedicine and its Applications (NLPBA/BioNLP)*, pages 73–78, Geneva, Switzerland. COLING.
- 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, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.
- Ning Ding, Guangwei Xu, Yulin Chen, Xiaobin Wang, Xu Han, Pengjun Xie, Haitao Zheng, and Zhiyuan Liu. 2021. [Few-nerd: A few-shot named entity recognition dataset](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 3198–3213. Association for Computational Linguistics.
- Christopher Dozier, Ravikumar Kondadadi, Marc Light, Arun Vachher, Sriharsha Veeramachaneni, and Ramdev Wudali. 2010. [Named entity recognition and resolution in legal text](#). In *Semantic Processing of Legal Texts: Where the Language of Law Meets the Law of Language*, volume 6036 of *Lecture Notes in Computer Science*, pages 27–43. Springer.
- Sumam Francis, Jordy Van Landeghem, and Marie-Francine Moens. 2019. [Transfer learning for named entity recognition in financial and biomedical documents](#). *Inf.*, 10(8):248.
- Fabrizio Gilardi, Meysam Alizadeh, and Maël Kubli. 2023. [ChatGPT outperforms crowd workers for text-annotation tasks](#). *Proceedings of the National Academy of Sciences*, 120(30):e2305016120.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [Lora: Low-rank adaptation of large language models](#). In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A Method for Stochastic Optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2019. [BioBERT: a pre-trained biomedical language representation model for biomedical text mining](#). *Bioinformatics*, 36(4):1234–1240.
- Yongqi Li, Yu Yu, and Tiejun Qian. 2023. [Type-aware decomposed framework for few-shot named entity recognition](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 8911–8927, Singapore. Association for Computational Linguistics.
- Jingjing Liu, Panupong Pasupat, Scott Cyphers, and James R. Glass. 2013. [Asgard: A portable architecture for multilingual dialogue systems](#). In *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2013, Vancouver, BC, Canada, May 26-31, 2013*, pages 8386–8390. IEEE.

- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized BERT pretraining approach](#). *CoRR*, abs/1907.11692.
- Zihan Liu, Feijun Jiang, Yuxiang Hu, Chen Shi, and Pascale Fung. 2021. [NER-BERT: A pre-trained model for low-resource entity tagging](#). *CoRR*, abs/2112.00405.
- Rui Mao, Guanyi Chen, Xulang Zhang, Frank Guerin, and Erik Cambria. 2023. [GpEval: A survey on assessments of chatgpt and GPT-4](#). *CoRR*, abs/2308.12488.
- Leland McInnes and John Healy. 2018. [UMAP: uniform manifold approximation and projection for dimension reduction](#). *CoRR*, abs/1802.03426.
- Xue Mengge, Bowen Yu, Zhenyu Zhang, Tingwen Liu, Yue Zhang, and Bin Wang. 2020. [Coarse-to-Fine Pre-training for Named Entity Recognition](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6345–6354, Online. Association for Computational Linguistics.
- OpenAI. 2023. [GPT-4 technical report](#). *CoRR*, abs/2303.08774.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- 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](#). *CoRR*, abs/2310.03668.
- Simone Tedeschi and Roberto Navigli. 2022. [MultiNERD: A multilingual, multi-genre and fine-grained dataset for named entity recognition \(and disambiguation\)](#). In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 801–812, Seattle, United States. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is All you Need](#). In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Shuhe Wang, Xiaofei Sun, Xiaoya Li, Rongbin Ouyang, Fei Wu, Tianwei Zhang, Jiwei Li, and Guoyin Wang. 2023a. [GPT-NER: named entity recognition via large language models](#). *CoRR*, abs/2304.10428.
- Xiao Wang, Weikang Zhou, Can Zu, Han Xia, Tianze Chen, Yuansen Zhang, Rui Zheng, Junjie Ye, Qi Zhang, Tao Gui, Jihua Kang, Jingsheng Yang, Siyuan Li, and Chunsai Du. 2023b. [Instructuie: Multi-task instruction tuning for unified information extraction](#). *CoRR*, abs/2304.08085.
- Ralph Weischedel, Martha Palmer, Mitchell Marcus, Eduard Hovy, Sameer Pradhan, Lance Ramshaw, Ni-anwen Xue, Ann Taylor, Jeff Kaufman, Michelle Franchini, Mohammed El-Bachouti, Robert Belvin, and Ann Houston. 2013. [OntoNotes Release 5.0](#).
- Urchade Zaratiana, Nadi Tomeh, Pierre Holat, and Thierry Charnois. 2023. [Gliner: Generalist model for named entity recognition using bidirectional transformer](#). *CoRR*, abs/2311.08526.
- Sheng Zhang, Hao Cheng, Jianfeng Gao, and Hoifung Poon. 2023. [Optimizing bi-encoder for named entity recognition via contrastive learning](#). In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. [Judging llm-as-a-judge with mt-bench and chatbot arena](#). *CoRR*, abs/2306.05685.
- Wenxuan Zhou, Sheng Zhang, Yu Gu, Muhao Chen, and Hoifung Poon. 2023. [Universalner: Targeted distillation from large language models for open named entity recognition](#). *CoRR*, abs/2308.03279.

A Appendix

A.1 Effect of Model Size

We investigated the influence of model size. We pre-train a version of NuNER using RoBERTa large (355M parameters) and compare it to the original NuNER (155M parameters). We see in Figure 12 an overall increase of the F1-score of a few percent. This increase is more pronounced for smaller training sets ($k < 10$) than for larger ones. Combined with the positive impacts of concept diversity and dataset size, this result suggests that scaling up both models and data would lead to further performance improvements.

A.2 Multilingual NuNER

We also decided to see the transferability of our approach to multilingual NER problems. The multilingual version of NuNER is based on the Multilingual BERT, pre-trained on 104 different languages. To further train multilingual NuNER, we created a new multilingual dataset. Specifically, for each selected language, we sampled 16000 text

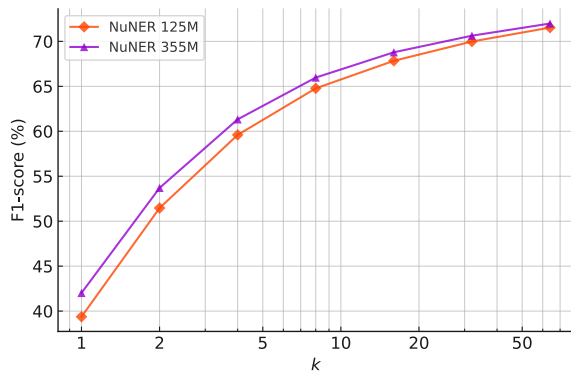


Figure 12: Effect of model size on NuNER’s performance. Results table shown in Table 8 in the appendix.

k	mBERT	NuNER
1	10.84	27.71
2	19.19	39.55
4	31.60	49.77
8	43.84	56.58
16	52.69	61.16
32	58.65	65.38
64	62.20	67.68

Table 3: NuNER vs multilingual BERT on MultiNERD dataset

fragments from the multilingual C4 dataset (Rafel et al., 2020) and labeled them, using the same prompt in English but explicitly mentioning to output the results in the language of the current fragment. We chose to label texts in English, French, Spanish, Portuguese, Dutch, Italian, Polish, and Russian.

For the model evaluation, we used the same protocol as described in Section 4.1, while evaluating the MultiNERD dataset (Tedeschi and Navigli, 2022). For the given k , the results are averaged over all evaluated languages. The results are found in Table 3.

As we can see, multilingual NuNER confidently outperforms multilingual BERT for all k , even for the bigger k the gap between models stays respectable, despite mBERT also being pre-trained on the evaluated languages.

A.3 Comparison with LLMs - Details

In the experiment of Figure 11 we only conduct one training run for GPT-4 for cost reasons. For all the other models we conduct 16 training runs for $k_w = 1$, 8 training runs for $k_w = 2$, 4 training runs for $k_w = 4$, 2 training runs for $k_w = 8$, and a unique training run for higher training sizes. These choices were made to reduce GPT-3.5 costs and to for the training data to be kept consistent across all

models for each training size.

A.4 UniversalNER Training Details

We create UniversalNER’s training data following Zhou et al. (2023)’s procedure: We use the dataset-specific instruction tuning template and populate it to extract all entity types. We don’t need to add any negative sampling here since all possible entity types are queried, therefore having an empty list when no entities of a given type are present in the example.

We fine-tune UniversalNER using authors’ fine-tuning code, only changing the number of training epochs and the number of gradient accumulation steps to obtain better few-shot results. After experimenting with external datasets, we found that using 20 epochs for $k_w = 8$ and $k = 8$, 15 epochs for $k_w = 16$ and $k_w = 32$, and 10 epoch for $k_w = 64$ and $k = 64$ worked well. We used these values and, for $k_w = 8$ and $k = 8$, also reduced the number of gradient accumulation steps to 4. We didn’t train UniversalNER on $k_w < 8$ as regularization was difficult to tune.

We also tried adopting LORA (Hu et al., 2022) for fine-tuning UniversalNER. Although, this methodology was more stable thanks to its implied regularization, it consistently led to worse results than full fine-tuning and required more time to converge.

We believe our heuristics allow to train UniversalNER pretty well in a few-shot setting, and are enough to make a rough comparison with NuNER. However, we should note that fine-tuning such large model on such small amount of data is not easy, and there are certainly better automatic machine learning procedures for this model.

A.5 Extra Tables and Figures

In this section we present the tables that were partially cut or just illustrated as plots in the body of the paper. We also show dataset-wise results that were shown as averages in the body of the paper.

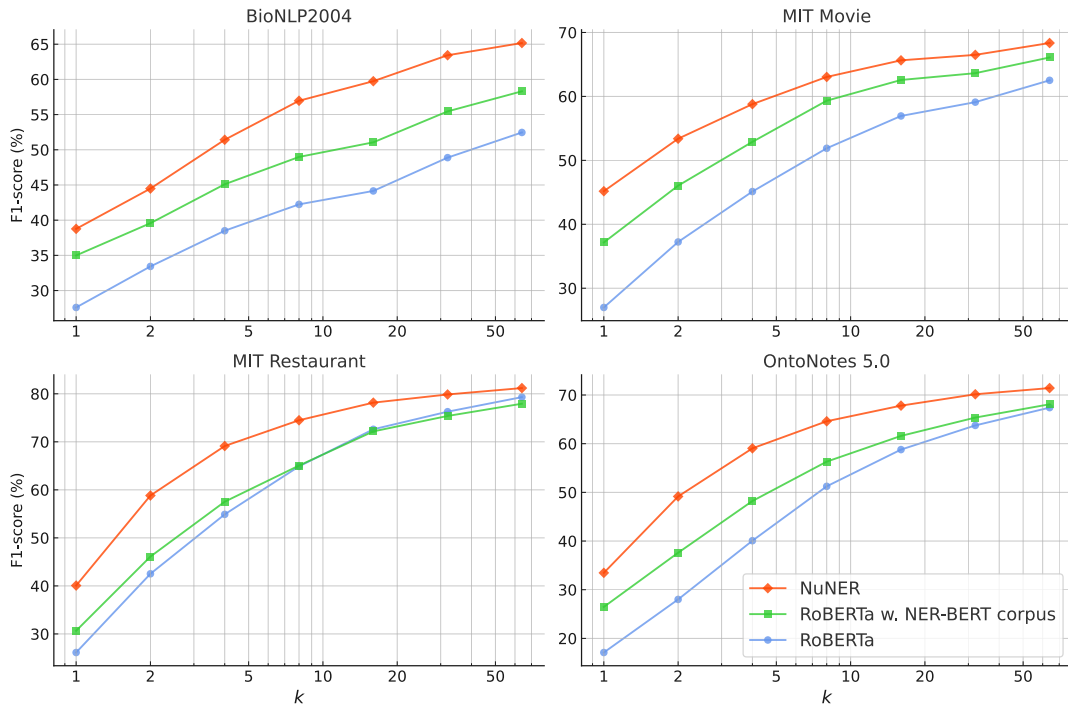


Figure 13: Dataset-wise results of Figure 7. Transfer learning performance of NuNER, RoBERTa, and RoBERTa pre-trained on NER-BERT data as function of k . NuNER substantially outperforms both models for all training sizes. We see similar behaviors for every dataset.

k	ROBERTA	ROBERTA w. NER-BERT	NuNER
1	24.46	32.32	39.37
2	35.30	42.33	51.46
4	44.65	50.94	59.60
8	52.56	57.42	64.78
16	58.12	61.85	67.84
32	62.00	64.96	69.98
64	65.42	67.61	71.53

Table 4: Results table of Figure 7. Transfer learning performance of NuNER, RoBERTa, and RoBERTa pre-trained on NER-BERT data as function of k . NuNER substantially outperforms both models for all training sizes.

k	4	16	154	1500	80k
	12.5%	25%	50%	75%	100%
1	32.78	36.95	39.56	39.45	38.83
2	44.40	47.33	49.75	50.53	50.69
4	52.57	54.75	56.60	57.58	57.97
8	59.35	60.66	61.84	62.78	63.43
16	63.74	65.05	65.49	66.38	66.83
32	66.76	67.67	67.96	68.92	69.27
64	69.31	70.02	70.03	70.76	71.04

Table 6: Results table of Figure 9. Effect of concept diversity on NuNER’s performance.

k	ANCH-AN. WIKI	LLM-AN. WIKI	LLM-AN. C4
1	32.52	38.23	38.64
2	41.86	49.99	50.55
4	50.35	57.32	57.82
8	57.40	62.76	63.26
16	61.82	66.39	66.76
32	64.81	68.85	69.18
64	67.46	70.74	70.93

Table 5: Results table of Figure 8. Effect of text diversity on NuNER’s performance. Wikipedia and C4 lead to similar performance when they are both annotated by the LLM.

k	1K	3K	30K	100K	300K	1M
1	32.8	35.6	38.3	38.9	39.3	39.4
2	45.5	48.0	50.2	50.8	51.6	51.5
4	54.0	56.2	57.6	58.1	59.1	59.6
8	60.3	61.9	63.0	63.5	64.5	64.8
16	64.3	65.4	66.6	66.9	67.7	67.8
32	66.5	67.8	69.0	69.3	69.8	70.0
64	68.3	69.4	70.7	71.1	71.6	71.5

Table 7: Full results table of Figure 10 for all training sizes. Effect of pre-training dataset size on NuNER’s performance.

k	NUNER	NUNER-LARGE
1	39.37	42.02
2	51.46	53.70
4	59.60	61.33
8	64.78	65.97
16	67.84	68.79
32	69.98	70.63
64	71.53	71.99

Table 8: Full results table of Figure 12 for all training sizes. Effect of model size on NuNER’s performance.

DATASET	NUNER	UNIVERSALNER
BIONLP	43.84 ± 2.18	41.79 ± 17.08
MIT MOVIE	58.69 ± 1.16	61.57 ± 2.26
MIT RESTAURANT	62.66 ± 2.28	63.4 ± 1.69
ONTONOTES	69.82 ± 1.64	64.8 ± 1.17
AVERAGE	58.75 ± 0.93	57.89 ± 4.34

Table 9: Dataset-wise results of Table 2 for $k = 8$.

DATASET	NUNER	UNIVERSALNER
BIONLP	61.18 ± 1.15	59.94 ± 6.03
MIT MOVIE	65.56 ± 0.29	69.50 ± 0.57
MIT RESTAURANT	73.66 ± 0.65	75.84 ± 0.97
ONTONOTES	80.81 ± 0.41	78.78 ± 0.26
AVERAGE	70.30 ± 0.35	71.02 ± 1.53

Table 10: Dataset-wise results of Table 2 for $k = 64$

k_w	GPT-3.5	GPT-4	UNIVERSALNER	NUNER
0	51.70	69.54	28.35	-
1	77.26	79.66	-	55.88
2	78.83	81.48	-	65.76
4	80.01	84.54	-	75.76
8	81.00	85.35	76.56	78.13
16	82.00	-	80.30	82.07
32	81.34	-	82.36	82.27
64	-	-	83.40	84.02

Table 11: Results table of Figure 11 for MIT Restaurant. Comparison of NuNER with LLMs. F1-macro token classification metric.

k_w	GPT-3.5	GPT-4	UNIVERSALNER	NUNER
0	63.48	61.09	60.23	-
1	67.62	70.82	-	46.37
2	70.70	77.10	-	55.95
4	70.22	76.32	-	58.16
8	70.86	78.80	69.52	69.09
16	68.88	-	76.09	76.93
32	69.12	-	80.09	78.96
64	-	-	80.96	81.48

Table 12: Results table of Figure 11 for BioNLP. Comparison of NuNER with LLMs. F1-macro token classification metric.