

Small Language Model Makes an Effective Long Text Extractor

Yelin Chen^{1*†}, Fanjin Zhang^{2*‡}, Jie Tang^{2‡}

¹School of Computer Science and Technology, Xinjiang University, Urumqi 830049, China

²Department of Computer Science and Technology, Tsinghua University, Beijing 100084, China
ylin@stu.xju.edu.cn, {fanjinz, jietang}@tsinghua.edu.cn

Abstract

Named Entity Recognition (NER) is a fundamental problem in natural language processing (NLP). However, the task of extracting longer entity spans (e.g., awards) from extended texts (e.g., homepages) is barely explored. Current NER methods predominantly fall into two categories: span-based methods and generation-based methods. Span-based methods require the enumeration of all possible token-pair spans, followed by classification on each span, resulting in substantial redundant computations and excessive GPU memory usage. In contrast, generation-based methods involve prompting or fine-tuning large language models (LLMs) to adapt to downstream NER tasks. However, these methods struggle with the accurate generation of longer spans and often incur significant time costs for effective fine-tuning. To address these challenges, this paper introduces a lightweight span-based NER method called SeNER, which incorporates a bidirectional arrow attention mechanism coupled with LogN-Scaling on the [CLS] token to embed long texts effectively, and comprises a novel bidirectional sliding-window plus-shaped attention (BiSPA) mechanism to reduce redundant candidate token-pair spans significantly and model interactions between token-pair spans simultaneously. Extensive experiments demonstrate that our method achieves state-of-the-art extraction accuracy on three long NER datasets and is capable of extracting entities from long texts in a GPU-memory-friendly manner.

Code — <https://github.com/THUDM/scholar-profiling/tree/main/sener>

Introduction

Named entity recognition (NER), a fundamental task in information extraction (IE), aims to identify spans indicating specific types of entities. It serves as the foundation for numerous downstream tasks, including relation extraction (Miwa and Bansal 2016), knowledge graph construction (Xu et al. 2017), and question answering (Mollá, Van Zaanen, and Smith 2006).

*These authors contributed equally.

†Part of the work was done when Yelin interned at Zhipu AI.

‡Corresponding authors.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Dr. Ware's profile snippet:

Prior to joining the UMMS faculty, Dr. Ware founded QualityMetric Incorporated and served as its CEO and Chairman for more than 10 years where he led the development of health status measures used in the Health Insurance Experiment . He has published more than 400 peer-reviewed articles , was the first recipient of the International Society for Pharmacoeconomics and Outcomes Research (ISPOR) Avedis Donabedian Outcomes Research Lifetime Achievement Award .


Work Experience Gender Award

Figure 1: An example of entity/attribute extraction from an author’s homepage, where “work experience” is a long entity block and “award” is a long entity.


Despite extensive studies, existing NER research rarely focuses on extracting named entities from long texts, a common real-world scenario such as extracting author attributes from homepages and identifying “methods” and “problems” in academic papers. For example, in Figure 1, “work experience” is a long entity block while “award” is a long entity, posing greater challenges for the NER task. We also extend the input length of a NER method to extract short entities in academic papers, as shown in Figure 2a, suggesting that long input length brings clear benefits to extract entities more precisely due to the perception of longer contexts.

Traditional approaches treat the NER task as a sequence labeling task, assigning a single label to each token, exemplified by the BIOES format. However, these methods are inadequate for recognizing nested entities. To address this issue, later efforts typically employ span-based methods (Su et al. 2022; Yan et al. 2023b), which consider all possible token-pair spans and classify each span. These methods achieve satisfactory accuracy on sentence-based NER tasks but struggle to identify entity blocks across sentences or extract entities from long texts due to substantial redundant computations and GPU memory usage resulting from $\mathcal{O}(L^2)$ computation complexity based on the token-pair span tensor, where L is the input length.

Recently, large language models (LLMs) demonstrate remarkable performance on a spectrum of natural language understanding and generation tasks (Zhao et al. 2023). However, LLMs still fall short and do not align well with information extraction tasks (Qi et al. 2024). On the Scholar-XL dataset (Zhang et al. 2024), extracting author attributes via prompting GPT-4o (Achiam et al. 2023) by providing



(a) Impact of input length for our model SeNER regarding extraction performance on the SciREX dataset (%).



(b) Model parameters of span-based and generation-based methods vs. F1 score on the Scholar-XL dataset.

Figure 2: Performance of entity recognition with respect to input length and the number of model parameters of NER methods.

5 similar demonstrations only achieves a 21.87% F1 score, as shown in Figure 2b. Fine-tuning LLMs to extract entities from long texts is also feasible (Sainz et al. 2023; Qi et al. 2024), but it incurs significant time costs for training and inference compared to span-based methods, without guaranteeing high accuracy.

Therefore, we aim to recognize named entities from long texts in a GPU-memory-friendly way without compromising accuracy. Given these limitations, we propose SeNER, a lightweight span-based method to extract entities from long texts. The main idea of SeNER is to reduce the redundant computations during the encoding and extraction processes of long texts. SeNER presents two core components that have an edge over existing span-based NER methods.

- To encode long texts effectively and efficiently, we employ a bidirectional arrow attention mechanism that encodes both local and global contexts simultaneously. To overcome the entropy instability issue of input texts of varied length, we apply LogN-Scaling (Su 2021) on the [CLS] token to keep the entropy of attention scores stable.
- To reduce superfluous span-based computation and model interactions between token-pair spans, we propose a novel bidirectional sliding-window plus attention (BiSPA) mechanism to efficiently compute horizontal and vertical attention on focused spans.

To enhance the robustness and generalization of our model, we employ the whole word masking strategy (Cui et al. 2021) and the LoRA (Hu et al. 2021) technique during training. Extensive experimental results on three datasets highlight the superiority of our proposed method. SeNER achieves state-of-the-art accuracy while maintaining relatively small model parameters, as depicted in Figure 2b. Additionally, under the same hardware and configuration, our model is capable of handling texts 6 times longer than previous advanced span-based NER methods.

Related Work

NER methods are generally categorized into span-based methods, generation-based methods, and other methods.

Span-based Methods

Span-based methods (Li et al. 2022; Yuan et al. 2022; Su et al. 2022; Zhu and Li 2022) reframe the NER task as a token-pair span classification task. They identify spans based on start and end positions, enumerate all possible candidate spans in a sentence, and perform classification. Most existing methods focus on obtaining high-quality span representations and modeling interactions between spans. CNN-NER (Yan et al. 2023a) utilizes Convolutional Neural Networks (CNNs) to model spatial relations in the token-pair span tensor. UTC-IE (Yan et al. 2023b) further incorporates axis-aware interaction with plus-shaped self-attention for the token-pair span tensor on top of CNN-NER. These methods offer parallel extraction, simple decoding, and advantages in handling nested entity recognition, leading to widespread use and excellent performance. However, calculating all span representations and aggregating interactions between token-pair spans requires substantial computational resources, which limits their effectiveness for long texts.

Generation-based Methods

Generation-based methods extract entities from text in an end-to-end manner, where the generated sequence can be text (Lu et al. 2022; Jiang et al. 2024), entity pointers (Yan et al. 2021), or code (Sainz et al. 2023). With the rise of large language models (LLMs), such methods (Wang et al. 2023a; Xie et al. 2023; Ashok and Lipton 2023) achieve good performance with only a few examples due to their generalization abilities. Some methods (Wang et al. 2023b; Dagdelen et al. 2024) enhance general extraction capabilities by using powerful LLMs, high-quality data, diverse extraction tasks, and comprehensive prior knowledge. GoLLIE (Sainz et al. 2023) ensures adherence to annotation guidelines through strategies such as class order shuffling, class dropout, guideline paraphrasing, representative candidate sampling, and class name masking. ADELIE (Qi et al. 2024) performs instruction tuning on a high-quality alignment corpus and further optimizes it with a Direct Preference Optimization (DPO) objective. However, compared to span-based methods, these methods often require significant computational resources and may perform poorly in generating accurate

longish entities from long texts. The construction of instructions and use of examples can compress input text length, leading to low text utilization. Additionally, autoregressive generation can result in long decoding times.

Other Methods

In addition to the two main paradigms, there are a few other types of methods. Some methods (Ma and Hovy 2016; Yan et al. 2019; Straková, Straka, and Hajič 2019) model the NER task as a sequence labeling task. However, these methods struggle with nested entities. Some methods (Li et al. 2019; Tan et al. 2021; Shen et al. 2022) use two independent multi-layer perceptrons (MLPs) to predict the start and end positions of entities separately, which can lead to errors due to treating the entity as separate modules. Some approaches (Lou, Yang, and Tu 2022; Yang and Tu 2022) employ hypergraphs to represent spans, but their decoding processes is complex.

Problem Definition

In this section, we introduce the problem formulation of named entity recognition from long texts.

Problem 1 Named Entity Recognition from Long Texts (Long NER). Given a long input text, the goal is to extract different types of named entities or entity blocks that mark their start and end positions in the text. Note that the input length can exceed 1,000 tokens and the entity length can exceed 100 tokens in our problem.

Taking scholar profiling (Schiaffino and Amandi 2009; Gu et al. 2018) as an example, “birth place” is a kind of entity, while “work experience” often appears as an entity block that involves multiple segments.

Method

As previously discussed, conventional NER methods fall into two main categories: span-based and generation-based. For NER in long texts, span-based methods need to model interactions between token-pair spans, which incurs substantial GPU memory and computation. In contrast, generation-based methods, commonly based on LLMs, are arduous to generate longish entity spans accurately.

In response to these limitations, we propose a lightweight span-based NER model, SeNER, that efficiently encodes long input texts and models token-pair spans interactions. First, we employ a pre-trained language model (PLM) with an arrow attention mechanism to encode long inputs efficiently. To alleviate entropy instability resulting from varied input lengths, we apply LogN-Scaling (Su 2021) to the [CLS] token. Next, we leverage a Biaffine model (Dozat and Manning 2017) to obtain the hidden representation of each token-pair span. Then, we present the token-pair span interaction module, where we propose a novel BiSPA mechanism to significantly reduce redundant candidate token pairs and model interactions between token pairs simultaneously. Finally, we introduce the training strategy and prediction method. An overview of our model is shown in Figure 3.

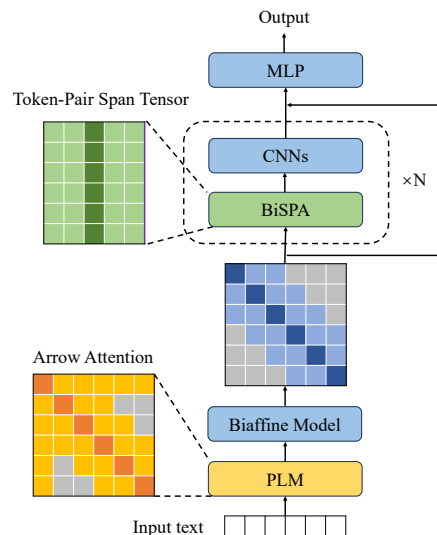


Figure 3: An overview of the SeNER model.




Figure 4: Illustration of arrow attention, full attention, and sliding window attention.

Long Input Encoding

Given a piece of text, we pass it into a PLM to obtain its contextual vector representation.

$$H = [h_1, h_2, \dots, h_L] = \text{PLM}([x_1, x_2, \dots, x_L]) \quad (1)$$

where $H \in \mathbb{R}^{L \times d}$, L is the input length, and d is the output dimension of the PLM.

Traditional NER methods utilize PLMs with full bidirectional attention, incurring a large amount of GPU memory footprint and computation for long texts. Moreover, full attention for long texts is often unnecessary since distant tokens are usually semantically unrelated. In light of this, a straightforward idea is to use sliding window attention (SWA) (Beltagy, Peters, and Cohan 2020; Zaheer et al. 2020), which adopts a fixed window, say w , so that each token attends to w tokens to its left and w tokens to its right. However, SWA ignores the global context, impairing the ability of the Transformer layers to acquire a comprehensive understanding of the entire input text.

To this end, we propose an **Arrow Attention** mechanism, where the [CLS] token uses global attention while other tokens use local sliding window attention, as illustrated in Figure 4. Arrow Attention strikes a balance between global and local attention. Compared to the computational complexity of $\mathcal{O}(L^2)$ for the full attention, arrow attention only requires $\mathcal{O}(wL)$. Furthermore, the global information captured by

the [CLS] token supplements the knowledge of SWA, enhancing the representation of each token and mitigating the information loss caused by the fixed receptive field. Thus, the [CLS] token acts as an attention sink (Xiao et al. 2023) that balances the weights of global and local contexts.

However, varying text lengths can cause entropy instability for the [CLS] token, where the scale of attention scores can change significantly. In this regard, we employ a LogN-Scaling technique on the [CLS] token to stabilize the entropy of attention scores. Specifically, LogN-Scaling is defined as follows:

$$H_{[\text{CLS}]}^t = \text{Attn}_s (H_{[\text{CLS}]}^{t-1} W^Q, H^{t-1} W^K, H^{t-1} W^V) \quad (2)$$

$$\text{Attn}_s (Q, K, V) = \text{softmax} \left(\frac{\log_{512} L}{\sqrt{d}} Q K^\top \right) V \quad (3)$$

where Attn_s is the scaled attention, H^t is the hidden representation of the t -th Transformer layer, and $W^Q, W^K, W^V \in \mathbb{R}^{d \times d}$ are projection matrices.

Note that LogN-Scaling is commonly used for length extrapolation in LLMs and imposed on all input tokens. Here we utilize LogN-Scaling solely on the [CLS] token to improve the stability and robustness of our model.

Biaffine Model

Subsequently, the hidden representation H is fed into a Biaffine model to extract features for each candidate span.

$$H^s, H^e = \text{MLP}_{\text{start}}(H), \text{MLP}_{\text{end}}(H) \quad (4)$$

$$S_{i,j} = (H_i^s)^\top W_1 H_j^e + W_2 (H_i^s \oplus H_j^e) + b \quad (5)$$

where $\text{MLP}_{\text{start}}$ and MLP_{end} are multi-layer perceptrons, $H^s/H^e \in \mathbb{R}^{L \times d}$ are hidden start/end embeddings, $W_1 \in \mathbb{R}^{d \times c \times d}$, $W_2 \in \mathbb{R}^{c \times 2d}$, $b \in \mathbb{R}^c$, and c is the output dimension of the Biaffine model. The symbol \oplus represents the concatenation operation. $S \in \mathbb{R}^{L \times L \times c}$, called token-pair span tensor, denotes the hidden representation of each candidate span. For example, $S_{i,j}$ represents the features of $[x_i, \dots, x_j]$.

Token-Pair Span Interaction Module

Note that the token-pair span tensor S considers each possible candidate span. However, for long input texts, it is unnecessary to consider every candidate span, especially for extremely long spans. Additionally, the GPU memory occupied by tensor S increases quadratically with the input length L . In light of this, we propose preserving only the hidden features of spans whose lengths do not exceed w' as shown in Figure 5. Thus, S is compressed to $S_h \in \mathbb{R}^{L \times w' \times c}$.

Previous studies (Yan et al. 2023a,b) show that modeling the interactions between token pairs, such as plus-shaped and local interaction, should be helpful. Plus-shaped attention applies the self-attention mechanism horizontally and vertically. However, plus-shaped attention cannot be performed directly on the compressed hidden feature tensor S_h since either the original horizontal or vertical dimension is disrupted. Therefore, we propose a novel bidirectional

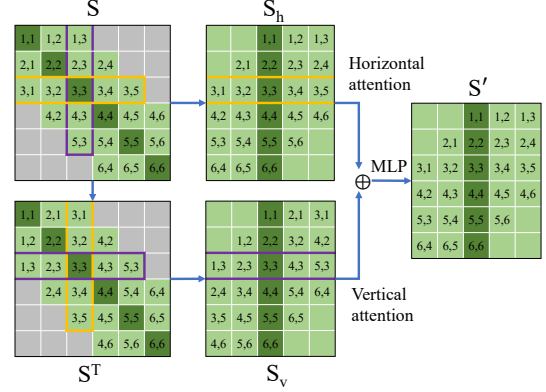


Figure 5: Diagram of the transformation for the token-pair span tensors in BiSPA mechanism.

sliding-window plus attention (**BiSPA**) mechanism to perform plus-shaped attention on the compressed S_h .

Specifically, we first compute the horizontal self-attention on S_h , as shown in the top middle of Figure 5. Next, we propose a transformation method, that transforms the top left matrix S to the bottom middle matrix S_v , and then compute the vertical self-attention based on S_v . Finally, we concatenate the horizontal and vertical attention matrices and feed them into an MLP to aggregate plus-shaped perceptual information. Notably, the computational complexity of the BiSPA mechanism is reduced from $\mathcal{O}(L^3)$ to $\mathcal{O}(L \times (w')^2)$, optimizing the training efficiency significantly.

$$Z_{i,:}^{h/v} = \text{Attn} \left(S_{i,:}^{h/v} W_{h/v}^Q, S_{i,:}^{h/v} W_{h/v}^K, S_{i,:}^{h/v} W_{h/v}^V \right) \quad (6)$$

$$\text{Attn} (Q, K, V) = \text{softmax} \left(\frac{QK^\top}{\sqrt{c}} \right) V \quad (7)$$

$$S' = \text{MLP} (Z^h \oplus Z^v) \quad (8)$$

where $W_h^Q, W_h^K, W_h^V, W_v^Q, W_v^K, W_v^V \in \mathbb{R}^{c \times c}$, Z^h/Z^v is intermediate representation after horizontal/vertical self-attention, and $S' \in \mathbb{R}^{L \times w' \times c}$ is the token-pair span feature after BiSPA mechanism.

The BiSPA mechanism endows the model with the capacity to perceive horizontal and vertical directions. We further use two types of position embeddings to enhance the sense of distances between token pairs and the area the token pair locates (Yan et al. 2023b). (1) Rotary Position Embedding (RoPE) (Su et al. 2024) encodes the relative distance between token pairs, which is used for both horizontal and vertical self-attention. (2) Matrix Position Embedding indicates whether each entry in S' is the original upper or lower triangles, which adds to S_h and S_v .

After the BiSPA mechanism, we employ CNN with kernel size 3×3 on S' to model the local interactions between token-pair spans.

$$S'' = \text{Recover}(\text{Conv}(\sigma(\text{Conv}(S')))) \quad (9)$$

where $S'' \in \mathbb{R}^{L \times L \times c}$ is recovered to the square size, and σ is the activation function.

We name the module encompassing the BiSPA mechanism and the convolutional module as the BiSPA Transformer block. The BiSPA Transformer blocks will be repeatedly used to ensure full interaction between token pairs.

Training and Prediction

We utilize MLP layers to transform the output of the final BiSPA Transformer block into output scores. We use binary cross-entropy as the loss function.

$$\hat{Y} = \text{MLP}(S'' + S) \quad (10)$$

$$\mathcal{L} = - \sum_{i,j=1}^L \sum_{r=1}^R \left(Y_{i,j}^r \log \left(\hat{Y}_{i,j}^r \right) \right) \quad (11)$$

$$+ \left(1 - Y_{i,j}^r \right) \log \left(1 - \hat{Y}_{i,j}^r \right) \quad (12)$$

where $\hat{Y} \in \mathbb{R}^{L \times L \times R}$ represents the scores of candidate entities, and R is the number of entity types.

To improve the robustness and generalization of our model, we employ the whole word masking strategy (Cui et al. 2021) during training and utilize LoRA (Hu et al. 2021) technique to train the PLM parameters.

During prediction, our model uses the average of the upper triangular and lower triangular values as the final prediction score, as follows:

$$P_{i,j}^r = \frac{\left(\hat{Y}_{i,j}^r + \hat{Y}_{j,i}^r \right)}{2}, i \leq j \quad (13)$$

All text spans that satisfy $P_{i,j}^r > 0$ are outputted. If the boundaries of multiple candidate spans conflict, the span with the highest prediction score is selected.

Experiment

Datasets

We conduct experiments on three NER datasets: Scholar-XL (Zhang et al. 2024), SciREX (Jain et al. 2020), and Profiling-07 (Tang, Zhang, and Yao 2007; Tang et al. 2008). The statistics of all datasets are detailed in Table 1. As shown in Table 1, the input lengths and entity lengths of the three datasets are longer than those of traditional named entity recognition datasets, presenting greater challenges.

Baselines

We compare our model with several recent NER methods:

Span-based Methods: **CNN-NER** (Yan et al. 2023a): is a span-based method that utilizes Convolutional Neural Networks (CNN) to model local spatial correlations between spans. **UTC-IE** (Yan et al. 2023b): models axis-aware interaction with plus-shaped self-attention and local interaction with CNN on top of the token-pair span tensor.

Others Methods: **DiffusionNER** (Shen et al. 2023): formulates the NER task as a boundary-denoising diffusion process and thus generates named entities from noisy spans.

	Scholar-XL	SciREX	Profiling-07
Input avg. len.	433.42	5678.29	785.09
Input max. len.	692	13731	17382
Input num.	2099	438	1446
Entity num.	20994	156931	17416
Entity type	12	4	13
Entity avg. len.	12.45	2.28	8.88
Entity max. len.	480	18	307

Table 1: Statistics of the datasets (in words).

Generation-based Methods: **UIE** (Lu et al. 2022): uniformly encodes different extraction structures via a structured extraction language, adaptively generates target extractions, and captures the common IE abilities via a large-scale pre-trained text-to-structure model. **InstructUIE** (Wang et al. 2023b): leverages natural language instructions and instruction tuning to guide large language models for IE tasks. **GOLLIE** (Sainz et al. 2023): is based on Code-Llama (Roziere et al. 2023) and fine-tunes the foundation model to adhere to specific annotation guidelines. **ADELIE** (Qi et al. 2024): builds a high-quality instruction tuning dataset and utilizes supervised fine-tuning (SFT) followed by direct preference optimization (DPO). **ToNER** (Jiang et al. 2024): firstly employs an entity type matching model to discover the entity types that are most likely to appear in the sentence, and then adds multiple binary classification tasks to fine-tune the encoder in the generative model. **GPT-4o** (Achiam et al. 2023): employs the gpt-4o-2024-08-06 API, utilizing a 5-shot in-context learning approach to enhance performance. **Claude-3.5** (Anthropic 2024): uses the claude-3-5-sonnet-20241022 API, also adopting 5-shot in-context learning.

Experimental Setup

All experiments are conducted on an 8-card 80G Nvidia A100 server. The entire text is used for the Scholar-XL dataset, while the other two datasets are truncated to 5120 using a sliding window approach, as a trade-off due to limited GPU memory. For prediction, the prediction of the text segment is mapped to the starting/ending position of the original text. Hyper-parameters are selected based on the F1 score on the validation set. For each experiment, we run 3 times with different random seeds and report the average results. We choose DeBERTa-V3-large (He, Gao, and Chen 2023) as the PLM for span-based methods and DiffusionNER. We use AdamW (Loshchilov, Hutter et al. 2017) optimizer with a weight decay of $1e^{-2}$. The unilateral window sizes of the arrow attention and BiSPA mechanism are both set to 128. We only use low-rank adaptation on the Q and V matrix of the self-attention mechanism with a rank of 8.

Evaluation Metrics

We report the micro-F1 score for all attributes. An entity is considered correct only if both the entity type and the entity span are predicted correctly. Precision (P) is the portion of correctly predicted spans over predicted spans, while Recall

(R) is the portion of correctly predicted spans over ground-truth entity spans.

Main Results

Table 2 provides a holistic comparison of different NER methods on three datasets. Generally speaking, span-based methods (CNN-NER, UTC-IE, and our model SeNER) outperform other types of NER methods.

Generation-based methods utilize generation loss to fine-tune the foundation model to adapt to the long NER task, achieving unfavorable performance. UIE outperforms InstructUIE, possibly because UIE defines a structured extraction language that suits the long NER problem better than naively performing instruction tuning. GOLLIE and ADELIE achieve similar performance, except for Profiling-07 dataset, which is due to the fact that this dataset is sliced and diced with a high number of empty data and thus makes GOLLIE overfit these empty examples. ToNER obtains unsatisfactory performance, possibly since the two-stage framework leads to error propagation and the usage of small language models for generation limits its potential. GPT-4o and Claude-3.5-sonnet are less effective, suggesting that the proprietary model prompting does not perform the long text NER task well.

The span-based NER methods (CNN-NER, UTC-IE, and SeNER) outperform other types of NER methods, including Diffusion-NER. DiffusionNER is a diffusion-based method that recovers the boundaries of the entities from a fixed amount of Gaussian noise and it is hard to recover longish entities from long texts. CNN-NER models fine-grained span interactions via CNN, achieving decent extraction performance. UTC-IE further improves CNN-NER by introducing plus-shaped attention on the token-pair span tensor, achieving consistent outperformance over CNN-NER.

Our model SeNER exhibits noticeable improvements or is on par with the best baseline, suggesting that with the design of arrow attention coupled with LogN-Scaling on the [CLS] in the PLM encoder, as well as the BiSPA mechanism on the token-pair span tensor, our model is capable of saving computation and memory resources without degrading the extraction accuracy. In addition, longer text with more focused attention can effectively help the model understand the semantic information of the text in more detail and extract the corresponding entities.

Ablation Study

Table 3 presents a justification for the effectiveness of each component in our model. Removing either the arrow attention or BiSPA mechanism results in a decrease in model performance on Scholar-XL and Out-of-Memory (OOM) errors on SciREX and Profiling-07. It indicates that both modules effectively reduce explicit memory usage, enabling the model to handle longer texts and thereby improving overall performance. Specifically, BiSPA significantly reduces compute and memory footprint by reducing negative samples. In contrast, the arrow attention has a limited ability to reduce memory usage for short text on the Scholar-XL dataset. Substituting the arrow attention with sliding window attention (SWA) leads to a significant performance drop, highlighting

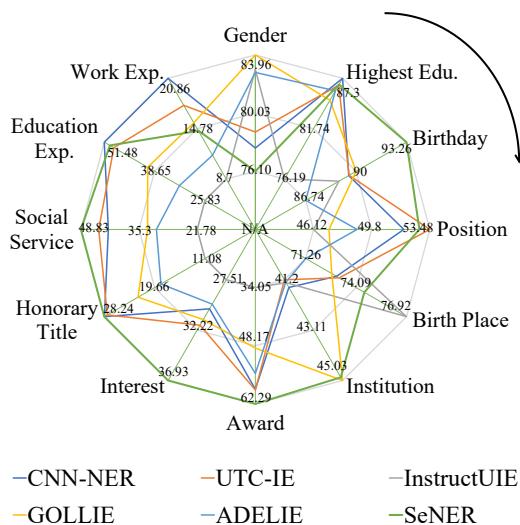


Figure 6: Performance of different entity types on the Scholar-XL dataset (%). The average length of entities increases clockwise from Gender to Work Exp. (Highest Edu.: Highest Education, Education Exp.: Education Experience, Work Exp.: Work Experience.)

the necessity of imposing attention scores on the [CLS] token to absorb global contextual information. Adding LogN-Scaling consistently improves the performance, thereby enhancing model stability and robustness. Although removing LoRA does not cause OOM errors, the F1 score decreases across all datasets to some extent, demonstrating that LoRA can effectively reduce training parameters and prevent overfitting. Whole Word Masking (WWM) increases the diversity of input texts, thus improving the generalization capacity of the model.

Detailed Analysis for Entity Types

In this subsection, we focus on comparing the performance of our method with span-based methods (CNN-NER and UTC-IE) and LLM-based methods (InstructUIE, GOLLIE, and ADELIE) across entities of varying lengths and types.

The results on the Scholar-XL dataset are depicted in Figure 6, with the average length of the entity types increasing clockwise from “Gender” to “Work Experience”. Generative methods, leveraging the powerful capabilities of LLMs, achieve superior performance in extracting “Gender” and “Birth Place” types. However, for other types of entities, span-based methods demonstrate consistent superiority. Our model SeNER outperforms CNN-NER and UTC-IE in most types of entities, with particularly notable improvements for longish entities. Specifically, for “Social Service”, our method achieves an improvement of 6.38% over CNN-NER and 4.14% over UTC-IE, respectively. The performance of SeNER for entity types “Education Experience” and “Work Experience” falls behind the leading ones a little, indicating that the approximation strategy in our model inevitably loses some information, especially on very long entities.

Method type	Method	Scholar-XL			SciREX			Profiling-07		
		P	R	F1	P	R	F1	P	R	F1
Generation-based Methods	UIE	43.32	36.80	39.80	65.88	56.44	60.80	65.92	57.51	61.43
	ToNER	40.08	29.48	33.97	57.43	31.56	40.73	48.80	41.21	44.68
	InstructUIE	34.63	36.50	35.54	56.31	54.60	55.44	59.09	63.19	61.07
	GOLLIE	43.74	40.88	42.26	71.56	71.50	71.53	64.51	9.46	16.50
	ADELIE	45.60	39.05	42.07	70.10	71.84	70.96	65.75	49.53	56.50
	Claude-3.5	17.15	30.16	21.87	57.78	7.97	14.01	34.40	43.07	38.25
	GPT-4o	18.24	27.31	21.87	40.44	7.69	12.92	36.96	43.73	40.06
Others Methods	DiffusionNER	<u>55.33</u>	29.87	38.80	77.11	62.36	68.96	70.51	44.16	54.31
Span-based Methods	CNN-NER	50.92	44.72	47.59	72.13	74.56	73.32	69.19	62.79	65.56
	UTC-IE	53.17	<u>46.01</u>	<u>49.10</u>	71.90	<u>75.09</u>	<u>73.42</u>	<u>69.79</u>	<u>65.28</u>	67.43
	SeNER (Ours)	57.41	46.80	51.56	<u>72.89</u>	76.17	74.49	67.52	67.17	<u>67.34</u>

Table 2: Main results on three long NER datasets (%). The best results are **boldfaced** and the second best results are underlined.

Method	Scholar-XL		SciREX		Profiling-07	
	F1	Mem	F1	Mem	F1	Mem
SeNER	51.56	16.95	74.49	69.23	67.34	63.36
w/o Arrow	51.34	16.94	-	OOM	-	OOM
w/o LogN	50.78	17.18	74.27	68.71	67.11	63.47
w SWA	49.95	16.46	73.50	68.33	65.60	63.16
w/o LoRA	49.93	17.98	73.98	70.79	66.53	66.58
w/o BiSPA	50.48	41.39	-	OOM	-	OOM
w/o WWM	50.45	17.14	74.24	69.00	64.97	63.37

Table 3: Ablation studies on three long NER datasets. Mem means memory usage (GB), SWA denotes sliding window attention and WWM is whole word masking.

Analysis for Maximum Input Length

We examine the maximum input length supported by training each NER method on a single Nvidia A100 with a batch size of 1, as shown in Figure 7. Generation-based methods often employ various lightweight strategies, such as quantization, FlashAttention (Dao et al. 2022), and Zero Redundancy Optimizer (ZeRO) (Rajbhandari et al. 2020) enabling models like GOLLIE and ADELIE to handle long texts. In contrast, span-based methods need to model token-pair span tensor, resulting in supporting shorter input length. Our method, SeNER, demonstrates substantial improvements over CNN-NER and UTC-IE, supporting input lengths that are 3 times and 6 times longer, respectively.

Efficiency Performance for Inference Time

Figure 7 also displays the inference time of span-based methods and LLM-based methods. It can be observed that LLM-based methods lead to 10 times longer inference time than span-based methods. Our method SeNER achieves a similar inference time compared with CNN-NER, achieving significantly better extraction accuracy simultaneously. SeNER can save 20% inference time compared with UTC-IE and encode longer input texts, still maintaining state-of-the-art extraction accuracy.




Figure 7: Blue bar: Maximum input length comparison of different methods (k). Orange bar: Inference time comparison of different methods (second). Both are conducted on the longest SciREX dataset.

Conclusion

In this paper, we tackle the problem of extracting entities from long texts, a less explored area in Named Entity Recognition (NER). Current span-based and generation-based NER methods face issues such as computational inefficiency and memory overhead in span enumeration, along with inaccuracy and time costs in text generation. To address these challenges, we introduce SeNER, a lightweight span-based approach that featuring a bidirectional arrow attention mechanism and LogN-Scaling for effective long-text embedding. Additionally, we propose a bidirectional sliding-window plus-shaped attention (BiSPA) mechanism that significantly reduces redundant candidate token-pair spans and models their interactions. Extensive experiments show that SeNER achieves state-of-the-art accuracy in extracting entities from long texts across three NER datasets, while maintaining GPU-memory efficiency. Our innovations in arrow attention and the BiSPA mechanism have the potential to advance future research in information extraction tasks.

Acknowledgments

This work is supported by NSFC for Distinguished Young Scholar 62425601, Tsinghua University Initiative Scientific Research Program and the New Cornerstone Science Foundation through the XPLOER PRIZE. This work is also supported by the Natural Science Foundation of China (NSFC) 62406164, the Postdoctoral Fellowship Program of CPSF under Grant Number GZB20240358 and 2024M761680.

References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Anthropic. 2024. The Claude 3 Model Family: Opus, Sonnet, Haiku.
- Ashok, D.; and Lipton, Z. C. 2023. Promptner: Prompting for named entity recognition. *arXiv preprint arXiv:2305.15444*.
- Beltagy, I.; Peters, M. E.; and Cohan, A. 2020. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*.
- Cui, Y.; Che, W.; Liu, T.; Qin, B.; and Yang, Z. 2021. Pre-training with whole word masking for chinese bert. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29: 3504–3514.
- Dagdelen, J.; Dunn, A.; Lee, S.; Walker, N.; Rosen, A. S.; Ceder, G.; Persson, K. A.; and Jain, A. 2024. Structured information extraction from scientific text with large language models. *Nature Communications*, 15(1): 1418.
- Dao, T.; Fu, D.; Ermon, S.; Rudra, A.; and Ré, C. 2022. Flashattention: Fast and memory-efficient exact attention with io-awareness. *Advances in Neural Information Processing Systems*, 35: 16344–16359.
- Dozat, T.; and Manning, C. D. 2017. Deep Biaffine Attention for Neural Dependency Parsing. In *International Conference on Learning Representations*.
- Gu, X.; Yang, H.; Tang, J.; Zhang, J.; Zhang, F.; Liu, D.; Hall, W.; and Fu, X. 2018. Profiling Web users using big data. *Social Network Analysis and Mining*, 8: 1–17.
- He, P.; Gao, J.; and Chen, W. 2023. DeBERTaV3: Improving DeBERTa using ELECTRA-Style Pre-Training with Gradient-Disentangled Embedding Sharing.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Jain, S.; van Zuylen, M.; Hajishirzi, H.; and Beltagy, I. 2020. SciREX: A Challenge Dataset for Document-Level Information Extraction. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 7506–7516.
- Jiang, G.; Luo, Z.; Shi, Y.; Wang, D.; Liang, J.; and Yang, D. 2024. ToNER: Type-oriented Named Entity Recognition with Generative Language Model. *arXiv preprint arXiv:2404.09145*.
- Li, J.; Fei, H.; Liu, J.; Wu, S.; Zhang, M.; Teng, C.; Ji, D.; and Li, F. 2022. Unified named entity recognition as word-word relation classification. In *proceedings of the AAAI conference on artificial intelligence*, 10965–10973.
- Li, X.; Feng, J.; Meng, Y.; Han, Q.; Wu, F.; and Li, J. 2019. A unified MRC framework for named entity recognition. *arXiv preprint arXiv:1910.11476*.
- Loshchilov, I.; Hutter, F.; et al. 2017. Fixing weight decay regularization in adam. *arXiv preprint arXiv:1711.05101*.
- Lou, C.; Yang, S.; and Tu, K. 2022. Nested Named Entity Recognition as Latent Lexicalized Constituency Parsing. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*, 6183–6198.
- Lu, Y.; Liu, Q.; Dai, D.; Xiao, X.; Lin, H.; Han, X.; Sun, L.; and Wu, H. 2022. Unified structure generation for universal information extraction. *arXiv preprint arXiv:2203.12277*.
- Ma, X.; and Hovy, E. 2016. End-to-end sequence labeling via bi-directional lstm-cnns-crf. *arXiv preprint arXiv:1603.01354*.
- Miwa, M.; and Bansal, M. 2016. End-to-end relation extraction using lstms on sequences and tree structures. *arXiv preprint arXiv:1601.00770*.
- Mollá, D.; Van Zaanen, M.; and Smith, D. 2006. Named entity recognition for question answering. In *Australasian Language Technology Association Workshop*, 51–58.
- Qi, Y.; Peng, H.; Wang, X.; Xu, B.; Hou, L.; and Li, J. 2024. ADELIE: Aligning Large Language Models on Information Extraction. *arXiv preprint arXiv:2405.05008*.
- Rajbhandari, S.; Rasley, J.; Ruwase, O.; and He, Y. 2020. Zero: Memory optimizations toward training trillion parameter models. In *International Conference for High Performance Computing, Networking, Storage and Analysis*, 1–16.
- Roziere, B.; Gehring, J.; Gloeckle, F.; Sootla, S.; Gat, I.; Tan, X. E.; Adi, Y.; Liu, J.; Remez, T.; Rapin, J.; et al. 2023. Code llama: Open foundation models for code. *arXiv preprint arXiv:2308.12950*.
- Sainz, O.; García-Ferrero, I.; Agerri, R.; de Lacalle, O. L.; Rigau, G.; and Agirre, E. 2023. Gollie: Annotation guidelines improve zero-shot information-extraction. *arXiv preprint arXiv:2310.03668*.
- Schiaffino, S.; and Amandi, A. 2009. Intelligent user profiling. In *Artificial Intelligence An International Perspective: An International Perspective*, 193–216.
- Shen, Y.; Song, K.; Tan, X.; Li, D.; Lu, W.; and Zhuang, Y. 2023. Diffusionner: Boundary diffusion for named entity recognition. *arXiv preprint arXiv:2305.13298*.
- Shen, Y.; Wang, X.; Tan, Z.; Xu, G.; Xie, P.; Huang, F.; Lu, W.; and Zhuang, Y. 2022. Parallel instance query network for named entity recognition. *arXiv preprint arXiv:2203.10545*.
- Straková, J.; Straka, M.; and Hajič, J. 2019. Neural architectures for nested NER through linearization. *arXiv preprint arXiv:1908.06926*.
- Su, J. 2021. Analyzing the Scale Operation of Attention from the Perspective of Entropy Invariance.

- Su, J.; Ahmed, M.; Lu, Y.; Pan, S.; Bo, W.; and Liu, Y. 2024. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568: 127063.
- Su, J.; Murtadha, A.; Pan, S.; Hou, J.; Sun, J.; Huang, W.; Wen, B.; and Liu, Y. 2022. Global pointer: Novel efficient span-based approach for named entity recognition. *arXiv preprint arXiv:2208.03054*.
- Tan, Z.; Shen, Y.; Zhang, S.; Lu, W.; and Zhuang, Y. 2021. A sequence-to-set network for nested named entity recognition. *arXiv preprint arXiv:2105.08901*.
- Tang, J.; Zhang, D.; and Yao, L. 2007. Social network extraction of academic researchers. In *Seventh IEEE International Conference on Data Mining*, 292–301.
- Tang, J.; Zhang, J.; Yao, L.; Li, J.; Zhang, L.; and Su, Z. 2008. Arnetminer: extraction and mining of academic social networks. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, 990–998.
- Wang, S.; Sun, X.; Li, X.; Ouyang, R.; Wu, F.; Zhang, T.; Li, J.; and Wang, G. 2023a. Gpt-ner: Named entity recognition via large language models. *arXiv preprint arXiv:2304.10428*.
- Wang, X.; Zhou, W.; Zu, C.; Xia, H.; Chen, T.; Zhang, Y.; Zheng, R.; Ye, J.; Zhang, Q.; Gui, T.; et al. 2023b. Instructuie: Multi-task instruction tuning for unified information extraction. *arXiv preprint arXiv:2304.08085*.
- Xiao, G.; Tian, Y.; Chen, B.; Han, S.; and Lewis, M. 2023. Efficient streaming language models with attention sinks. *arXiv preprint arXiv:2309.17453*.
- Xie, T.; Li, Q.; Zhang, J.; Zhang, Y.; Liu, Z.; and Wang, H. 2023. Empirical Study of Zero-Shot NER with ChatGPT. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 7935–7956.
- Xu, B.; Xu, Y.; Liang, J.; Xie, C.; Liang, B.; Cui, W.; and Xiao, Y. 2017. CN-DBpedia: A never-ending Chinese knowledge extraction system. In *International Conference on Industrial, Engineering and Other Applications of Applied Intelligent Systems*, 428–438.
- Yan, H.; Deng, B.; Li, X.; and Qiu, X. 2019. TENER: adapting transformer encoder for named entity recognition. *arXiv preprint arXiv:1911.04474*.
- Yan, H.; Gui, T.; Dai, J.; Guo, Q.; Zhang, Z.; and Qiu, X. 2021. A unified generative framework for various NER sub-tasks. *arXiv preprint arXiv:2106.01223*.
- Yan, H.; Sun, Y.; Li, X.; and Qiu, X. 2023a. An Embarrassingly Easy but Strong Baseline for Nested Named Entity Recognition. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics*, 1442–1452.
- Yan, H.; Sun, Y.; Li, X.; Zhou, Y.; Huang, X.; and Qiu, X. 2023b. UTC-IE: A Unified Token-pair Classification Architecture for Information Extraction. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics*, 4096–4122.
- Yang, S.; and Tu, K. 2022. Bottom-Up Constituency Parsing and Nested Named Entity Recognition with Pointer Networks. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*, 2403–2416.
- Yuan, Z.; Tan, C.; Huang, S.; and Huang, F. 2022. Fusing Heterogeneous Factors with Triaffine Mechanism for Nested Named Entity Recognition. In *Findings of the Association for Computational Linguistics*, 3174–3186.
- Zaheer, M.; Guruganesh, G.; Dubey, K. A.; Ainslie, J.; Al-berti, C.; Ontanon, S.; Pham, P.; Ravula, A.; Wang, Q.; Yang, L.; et al. 2020. Big bird: Transformers for longer sequences. *Advances in neural information processing systems*, 33: 17283–17297.
- Zhang, F.; Shi, S.; Zhu, Y.; Chen, B.; Cen, Y.; Yu, J.; Chen, Y.; Wang, L.; Zhao, Q.; Cheng, Y.; et al. 2024. OAG-Bench: A Human-Curated Benchmark for Academic Graph Mining. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*.
- Zhao, W. X.; Zhou, K.; Li, J.; Tang, T.; Wang, X.; Hou, Y.; Min, Y.; Zhang, B.; Zhang, J.; Dong, Z.; et al. 2023. A survey of large language models. *arXiv preprint arXiv:2303.18223*.
- Zhu, E.; and Li, J. 2022. Boundary smoothing for named entity recognition. *arXiv preprint arXiv:2204.12031*.