

GapDNER: A Gap-Aware Grid Tagging Model for Discontinuous Named Entity Recognition

1st Yawen Yang

School of Software
Tsinghua University

Beijing, China

yyw19@mails.tsinghua.edu.cn

2nd Fukun Ma

School of Software
Tsinghua University

Beijing, China

mfk22@mails.tsinghua.edu.cn

3rd Shiao Meng

School of Software
Tsinghua University

Beijing, China

msa21@mails.tsinghua.edu.cn

4th Aiwei Liu

School of Software
Tsinghua University

Beijing, China

liuaw20@mails.tsinghua.edu.cn

5th Lijie Wen*

School of Software
Tsinghua University

Beijing, China

wenlj@tsinghua.edu.cn

Abstract—In biomedical fields, one named entity may consist of a series of non-adjacent tokens and overlap with other entities. Previous methods recognize discontinuous entities by connecting entity fragments or internal tokens, which face challenges of error propagation and decoding ambiguity due to the wide variety of span or word combinations. To address these issues, we deeply explore discontinuous entity structures and propose an effective Gap-aware grid tagging model for Discontinuous Named Entity Recognition, named GapDNER. Our GapDNER innovatively applies representation learning on the context gaps between entity fragments to resolve decoding ambiguity and enhance discontinuous NER performance. Specifically, we treat the context gap as an additional type of span and convert span classification into a token-pair grid tagging task. Subsequently, we design two interactive components to comprehensively model token-pair grid features from both intra- and inter-span perspectives. The intra-span regularity extraction module employs the biaffine mechanism along with linear attention to capture the internal regularity of each span, while the inter-span relation enhancement module utilizes criss-cross attention to obtain semantic relations among different spans. At the inference stage of entity decoding, we assign a directed edge to each entity fragment and context gap, then use the BFS algorithm to search for all valid paths from the head to tail of grids with entity tags. Experimental results on three datasets demonstrate that our GapDNER achieves new state-of-the-art performance on discontinuous NER and exhibits remarkable advantages in recognizing complex entity structures.

Index Terms—Discontinuous NER, Gap-Aware, Intra-Span Regularity, Inter-Span Relation

I. INTRODUCTION

Named entity recognition (NER) aims to identify mentions with specific meanings in text and classify them into pre-defined entity categories, including person, location, gene, disease, etc. NER has long been a fundamental task in natural language processing because of its wide applications in downstream tasks like entity linking [1], relation extraction [2] and knowledge graph [3].

Traditional entity definition contains two underlying assumptions that an entity should be composed of consecutive adjacent tokens and should not overlap with other ones [4]. While in practical applications, there are often complex entity structures that do not conform to the above assumptions, including nested, overlapped and discontinuous entities. Among them, the discontinuous entity denotes an entity consisting of a sequence of non-adjacent tokens. As illustrated in Fig. 1, “severe shoulder pain” represents a discontinuous entity separated by two context gaps. Since typical BIO tagging methods fail to handle discontinuous entities with diverse overlapped structures, discontinuous NER still remains a challenge for most NER systems.




Fig. 1. Example of three discontinuous entities (“severe joint pain”, “severe shoulder pain” and “severe upper body pain”) within the same sentence. The blue and red dashed boxes represent the fragments of discontinuous entities and context gaps respectively.

Some meaningful efforts have been devoted to explorations of discontinuous NER. Existing methods could be primarily divided into sequence-labeling, span-based, seq2seq-based and grid-tagging solutions. **1) Sequence labeling methods** extend the common BIO tagging mode to more complex tagging schemes such as BIOHD [5] and the 10-tag scheme [6]. Such task-specific extensions restrict the adaptability of NER models to new entity types and unseen data sets. **2) Span-based models** [7] recognize discontinuous entities by identifying entity fragments and combining them together, which are limited by maximum span length and struggle in decoding

*Corresponding author.

ambiguity. **3) Seq2seq-based methods** [8], [9], [10] transform discontinuous entity recognition into entity span sequence generation task and adopt the pre-trained seq2seq model to solve it. These methods have no need of span connections, thus avoiding the problem of decoding ambiguity. However, they suffer from error accumulation and low inference efficiency caused by the autoregressive mechanism. **4) Grid tagging models** [11], [12], [13], which formulate entity extraction as token-pair grid classification, have achieved excellent performance in discontinuous NER. But they overlook the semantic relations between different spans and still suffer from decoding ambiguity problem during token connection. Most recently, **large language models (LLMs)** have been explored to offer general solutions to entity recognition through prompt engineering and instruction tuning. But their performances [14], [15] on discontinuous NER are still far from current state-of-the-art models. LLMs, which are optimized for next token prediction, are not satisfactory entity extractors currently.

To reduce the above drawbacks, we explore deep into boundaries and structures of discontinuous entities. From Fig. 1, we discover that each range of discontinuous entities follows the rule that entity fragments and context gaps are arranged alternatively. Consequently, identifying context gaps correctly will certainly help locate entity fragments and alleviate decoding ambiguity. Additionally, entity fragments provide important clues for entity type recognition, and context gaps always contain non-entity tokens. It is beneficial to capture the internal regularity of each span before performing span classification. Furthermore, discontinuous entities are often accompanied by overlapped situations. Learning the semantic relations between different spans may further improve the classification performance.

Armed with these observations, we propose an effective **Gap-aware grid tagging model for Discontinuous NER**, called **GapDNER**, to fully utilize gap information. We regard context gaps as special spans and convert span classification into the token-pair grid tagging problem. Next, we design two sequential modules to model semantic features from both intra- and inter-span aspects. The intra-span regularity extraction module empirically employs the Biaffine mechanism on token pairs and skillfully applies linear attention to all tokens within each token pair, aiming to capture the internal regularity of each span (namely entity fragment or context gap). After that, we concatenate the two representations to construct the token-pair feature matrix. We notice that in the overlapped structures of discontinuous entities, several spans share the same head or tail token and thus appear in the same row or column of the token-pair grid. Therefore, the inter-span attention enhancement module calculates the criss-cross attention over the whole feature matrix to measure inter-span relations and facilitate span classification. At the inference stage, we allocate a directed edge to each span, then employ the BFS (Breadth-First Search) algorithm to search for all valid paths from the head to the tail of token pairs with entity type labels. Each path corresponds to an actual entity. The main contributions of this work can be listed as follows:

- We propose an effective gap-aware grid tagging model for discontinuous NER, which innovatively leverages context gap information to enhance entity recognition. To the best of our knowledge, we are the first to apply the representation learning on context gaps in explorations of discontinuous NER.
- We adopt Biaffine mechanism along with linear attention to obtain the intra-span regularity and employ the criss-cross attention to capture the inter-span relations. The integration of intra- and inter-span features significantly promotes span classification.
- We conduct comprehensive experiments on three standard benchmarks for discontinuous NER. Experimental results show that our GapDNER outperforms strong baselines consistently on each dataset and achieves new state-of-the-art results in diverse settings.

II. RELATED WORK

A. Discontinuous NER

Discontinuous NER (DNER) involves recognizing entities composed of a sequence of non-adjacent tokens, which is of great significance for medical record recognition and drug safety monitoring. Existing methods for DNER can be mainly divided into the following categories.

Sequence labeling methods assign a label to each token that indicates the location and type of entities. Tang et al. [5] extended the BIO scheme to BIOHD by introducing two additional tags to distinguish overlapping entity fragments. Corro [6] proposed a new tagging scheme consisting of 10 labels to mark tokens of diverse complex structures. Although effective to some extent, these methods have limitations of tagging inflexibility and decoding ambiguity.

Span-based methods always transform discontinuous NER into a two-stage process that includes span detection and span connection. Li et al. [7] developed a span-based model to jointly recognize overlapped and discontinuous entities by enumerating all possible spans and identifying the neighboring relation between them. Due to the enumeration nature, these methods are limited by the maximum span length and have higher complexity when recognizing long entities.

Seq2seq-based methods uniformly formulate entity recognition as entity span sequence generation. BartNER [8], which incorporated the pre-trained seq2seq model BART and a pointer network, was the first to employ the generative framework to address various NER subtasks. After that, Debias-DDA [9] and Debias-CSL [10] reduced optimization objective bias and fixed-order bias of seq2seq models, respectively. Unfortunately, they face challenges of error propagation and low inference efficiency caused by the autoregressive mechanism.

Grid tagging methods have recently aroused increasing attention by converting entity identification into token-pair grid classification. Wang et al. [11] reformulated DNER as maximal clique discovery and first proposed grid tagging solutions. Another typical approach is W2NER [12], which proposed to model multi-type NER as the word-word relation classification. Inspired by that, Liu et al. [13] put forward another

grid tagging model for DNER by introducing two additional relation tags. Although these models have achieved excellent performance, they still suffer from decoding ambiguity when handling complex overlapped structures.

Other approaches for DNER include the LLM-related [14], [15], hypergraph-based [16] and transition-based [17] models.

B. Optimizations of Grid Attention Computation

In computer vision, it is essential to perform attention calculations on the 2D grid of feature maps. However, capturing global attention often results in excessive computational complexity. To address this difficulty, several optimization methods have been proposed for approximate computation.

Axial attention [18] calculates attention scores along a specific axis (horizontal or vertical) of the feature map, mixing feature information from that axis while keeping information along other axes independent. Most related works combine two axial attention layers of different axes to make the neural model acquire global receptive fields. **Criss-cross attention** [19] fuses feature information in a criss-cross manner, which is similar to the combination of row attention and column attention. It enables each pixel not only to pay attention to the local area but also to transmit information across regions, breaking the limitations of traditional local receptive fields. **Dilated attention** [20], sourced from the dilated convolution, introduces a dilation factor to expand the receptive field when calculating attentions on a feature map. It makes grid elements attend to information at a greater distance without increasing the computational complexity.

In this work, we employ the criss-cross attention to model inter-span relations due to the complex overlapped structures among discontinuous entities.

III. GAP-AWARE GRID TAGGING SCHEME

We regard context gaps as special mentions and define four types of span used for entity recognition. **1) ConE** denotes a continuous entity. **2) DisE** denotes the range of a discontinuous entity (e.g., from “severe” to “pain” in Fig. 1). **3) Frag** represents a continuous fragment of one discontinuous entity (e.g., “severe joint” in Fig. 1). **4) Gap** represents a context interval between two neighboring fragments of the same discontinuous entity (e.g., “and upper body” in Fig. 1).

Equipped with these definitions, we predefine the token pair grid label set C for both continuous and discontinuous entities. Considering one span can be both a continuous entity and a fragment of another discontinuous entity simultaneously, we annotate the (head, tail) grid with the same fragment label (“Frag”) to uniformly represent them. Meanwhile, we distribute entity type labels (e.g., “ADE”, “Disorder”) to the (tail, head) grid of each ConE or DisE span, which serve as the begin flag for entity decoding. Finally, we annotate the (head, tail) grid of a context gap with the “Gap” label. Such operations face a conflict in situations where one token is both an entity and a fragment. To solve this, we choose to annotate the one-word entity with the corresponding entity type regardless of whether it is a fragment, because each

one-word entity will be definitely considered as a potential fragment in the inference stage.

Under this tagging scheme, the upper triangular region of the grid can only include “<Frag>” and “<Gap>” labels. The lower triangle region can only contain entity type labels like “<ADE>”. Specially, the diagonal positions can contain all the above types since they stand for spans with only one token. For better understanding, we provide an example of grid tagging scheme in Fig. 2, corresponding to the input sentence.

	Severe	joint	,	shoulder	and	upper	body	pain
Severe	Frag	Frag						
joint			Gap		Gap			
,							Gap	
shoulder				Frag				
and							Gap	
upper								Frag
body								
pain	ADE							Frag

Fig. 2. An example of our gap-aware grid tagging scheme. Blue and red fonts represent fragment and gap labels respectively, while green fonts denote entity type labels. The black dotted line covers the diagonal area of the grid.

IV. PROPOSED MODEL

The overall model architecture is illustrated in Fig. 3, which mainly includes three components: Sequence Encoder Layer, Intra-Span Regularity Extraction and Inter-Span Relation Enhancement. During model inference, we decode entities based on classification results of the token-pair feature grid.

A. Sequence Encoder Layer

We adopt BERT [21] and BiLSTM [22] to construct sequence encoder layer. The pre-trained model BERT has been demonstrated to be one of the outstanding models for representation learning in information extraction. BiLSTM excels at modeling sequence order and local dependencies. Given an input sentence of n tokens $s = [t_1, t_2, \dots, t_n]$, we feed it into the BERT network to obtain context-dependent embeddings. After the BERT calculation, one target token may be divided into several word pieces according to BERT dictionary. Following pervious works [12], we employ the max pooling operation to generate token representation from word piece embeddings. Then the sequence of token representations is put into BiLSTM module to produce the final token representation $\mathbf{H} = [h_1, h_2, \dots, h_n] \in \mathbb{R}^{n \times d}$, which can be formulated as:

$$x_i = \text{BERT}(t_i), \quad (1)$$

$$\vec{h}_i = \text{LSTM}_f(\vec{h}_{i-1}, x_i; \theta_f), \quad (2)$$

$$\overleftarrow{h}_i = \text{LSTM}_b(\overleftarrow{h}_{i-1}, x_i; \theta_b), \quad (3)$$

$$h_i = [\vec{h}_i; \overleftarrow{h}_i], \quad (4)$$

where n is the sequence length, d is the hidden dimension of token representations after BiLSTM, $[\cdot; \cdot]$ concatenates two vectors in the last dimension, θ_f and θ_b denote the parameters of the forward and backward LSTM, respectively.




Fig. 3. The overall architecture of proposed GapDNER model. $\otimes$ denotes vector concatenation of the last dimension. $\oplus$ represents the element-wise addition.

B. Intra-Span Regularity Extraction

In this module, we aim to learn span representations that incorporate internal regularities for better classification. Typical span-based methods for NER represent span features via concatenating head and tail embeddings directly. Subsequent revisions include span length feature integration [7], and Biaffine decoder [23] which strengthens the semantic interaction between head and tail tokens. Nevertheless, these methods only capture insufficient and coarse-grained span features since they concentrate mainly on span boundaries but ignore the internal composition regularity.

To effectively learn the semantic regularity of entity fragments and context gaps within discontinuous entities, we employ Biaffine mechanism and linear attention to explore boundary features and internal regularity, respectively. For the former, we first exploit two separate multilayer perceptrons (MLPs) to map the head and tail tokens into different semantic spaces. Then we apply a biaffine decoder over the two representation sequences to generate span boundary feature matrix $\mathbf{H}^{(\text{span})} \in \mathbb{R}^{n \times n \times d}$. Each element is computed as follows:

$$\bar{h}_i = \text{MLP}_{\text{head}}(h_i), \quad (5)$$

$$\bar{h}_j = \text{MLP}_{\text{tail}}(h_j), \quad (6)$$

$$h_{ij}^{(\text{span})} = \bar{h}_i^\top U^{(1)} \bar{h}_j + [\bar{h}_i; \bar{h}_j] U^{(2)} + b_1, \quad (7)$$

where $\bar{h}_i, \bar{h}_j \in \mathbb{R}^{d \times 1}$ are the head and tail representations of token pair (i, j) after semantic mapping. $U^{(1)}$ denotes a $d \times d \times d$ tensor, $U^{(2)}$ represents a $2d \times d$ matrix, and $b_1 \in \mathbb{R}^{1 \times d}$ is the bias. $U^{(1)}, U^{(2)}, b_1$ are all learnable parameters.

It is worth noting that regularity information stems from each token in the span. The RICON method [24] has compared

the performance of several regularity extracting ways (e.g., Linear attention, Max-pooling, Multi-Head attention) on Chinese NER and discovered linear attention outperformed others. Inspired by that, we select the simple but effective linear attention to extract the intra-span regularity $\mathbf{H}^{(\text{reg})} \in \mathbb{R}^{n \times n \times d}$. For each token pair grid (i, j) , assuming $i < j$, the attention score and regularity representation are calculated as:

$$a_t = W_{\text{reg}}^\top h_t + b_{\text{reg}}, \quad (8)$$

$$\alpha_t = \frac{\exp(a_t)}{\sum_{k=i}^j \exp(a_k)}, \quad (9)$$

$$h_{ij}^{(\text{reg})} = \sum_{t=i}^j \alpha_t \cdot h_t, \quad (10)$$

where $t \in \{i, i+1, \dots, j\}$ stands for the token index of the span, $W_{\text{reg}} \in \mathbb{R}^{d \times 1}$ and $b_{\text{reg}} \in \mathbb{R}^1$ are learnable weights and bias. Especially, for the span which has only one token (namely $i = j$), we naturally adopt the hidden representation h_t from BiLSTM to indicate its regularity. Since entity type labels are located in the lower triangular region of the grid where $i > j$, we introduce another weights vector and bias to compute the linear attention and corresponding regularity representation. Then we concatenate span boundary and internal regularity representations as the token-pair grid feature $\mathbf{G} \in \mathbb{R}^{n \times n \times 2d}$.

$$\mathbf{G} = [\mathbf{H}^{(\text{span})}; \mathbf{H}^{(\text{reg})}] \quad (11)$$

C. Inter-Span Relation Enhancement

As mentioned above, different spans have rich spatial and semantic relations such as neighboring, nested, overlapped and head-tail. It should be beneficial to span identification if we

are capable of leveraging these special correlations. To achieve this goal, we calculate the inter-span attention over the token-pair grid and obtain attention-enhanced grid representations for subsequent predictions.

We observe that discontinuous entities often appear together with overlapped structures in biomedical documents. Such co-occurrence situations possibly make entity fragments or context gaps of different entities share the same head or tail token. Take Fig. 2 for example, “severe” and “severe joint” are both entity fragments sharing the same head boundary, thus they are located in the same row of the grid. Similarly, “shoulder and upper body” and “and upper body” denote two context gaps containing the same tail token, and are distributed in the same column of the token-pair grid.

Considering the obvious axial distribution characteristics of different spans, we introduce criss-cross attention [19] to measure the inter-span relations and reduce computational complexity. As depicted in Fig. 3, the criss-cross attention module collects interactive information in horizontal and vertical directions to enhance span-level representative capability. To be specific, we first utilize one MLP to map the concatenated span regularity features into the same semantic space, and reduce vector dimension simultaneously.

$$\mathbf{M} = \text{MLP}(\mathbf{G}) \quad (12)$$

where $\mathbf{M} \in \mathbb{R}^{n \times n \times d}$. Next, the attention module adopts two convolutional layers with 1×1 filters on $\mathbf{M}$ to generate the query and key matrix $\mathbf{Q}$ and $\mathbf{K}$ respectively, where $\mathbf{Q}, \mathbf{K} \in \mathbb{R}^{n \times n \times d'}$. Another convolutional layer with 1×1 filters is applied on $\mathbf{M}$ to produce the value matrix $\mathbf{V} \in \mathbb{R}^{n \times n \times d}$. For each token pair grid q_{ij} of $\mathbf{Q}$, we obtain the corresponding criss-cross set $\mathbf{K}_{cc} \in \mathbb{R}^{(n+n-1) \times d'}$, $\mathbf{V}_{cc} \in \mathbb{R}^{(n+n-1) \times d}$ by extracting feature vectors which are in the same row or column with grid q_{ij} from $\mathbf{K}$ and $\mathbf{V}$. Particularly, for the grid position (i, j) , the corresponding component m'_{ij} of the criss-cross attention output $\mathbf{M}' \in \mathbb{R}^{n \times n \times d}$ is defined as follows:

$$\begin{aligned} m'_{ij} &= \text{CC_Attention}(q_{ij}, \mathbf{K}_{cc}, \mathbf{V}_{cc}) \\ &= \text{Softmax}\left(\frac{q_{ij} \mathbf{K}_{cc}^\top}{\sqrt{d'}}\right) \mathbf{V}_{cc}, \quad 1 \leq i, j \leq n, \end{aligned} \quad (13)$$

where n is the length of input sentence and also denotes the height and width of the token pair grid. $\mathbf{K}_{cc}$ and $\mathbf{V}_{cc}$ represent keys and values selected from the feature vectors $\mathbf{K}$ and $\mathbf{V}$. Given the query q_{ij} located in (i, j) , keys and values positioned at the union set of same row and column elements (formally $\{(i', j') | i' = i \text{ or } j' = j\}$) will be selected to perform attention computation. In this way, the proposed inter-span relation enhancement module can model the overlapped and neighboring relations between multiple spans, and can capture long-range dependency effectively.

After that, we conduct the element-wise addition to integrate the attention output with initial feature matrix.

$$\mathbf{M}'' = \mathbf{M}' + \mathbf{M}, \quad (14)$$

where $\mathbf{M}'' \in \mathbb{R}^{n \times n \times d}$ denotes the relation enhanced output.

D. Model Training

Based on the grid representation enhanced by inter-span relations, we exploit a linear layer to obtain the classification logits and adopt the softmax function to calculate the label probability $\hat{y}_{ij}$ for the token pair (t_i, t_j) as:

$$\hat{y}_{ij} = \text{Softmax}(\text{Linear}(\mathbf{M}''_{ij})), \quad (15)$$

where $y_{ij} \in \mathbb{R}^{|C|}$ denotes the probability vector of the grid tags predefined in C . For each sentence $s = [t_1, t_2, \dots, t_n]$, our training target is to minimize the negative log-likelihood (NLL) loss between the predicted probability and the corresponding gold labels, formulated as:

$$\mathcal{L} = -\frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \sum_{c=1}^{|C|} y_{ij}^c \log \hat{y}_{ij}^c, \quad (16)$$

where n is the number of sentence tokens, y_{ij} is the one-hot binary vector with regards to gold class labels, $\hat{y}_{ij}$ represent the predicted probability vector. The superscript c indicates the c -th class of the predefined grid tag set C .

E. Entity Decoding

In the inference stage, the predictions of our model are token pairs and their corresponding grid labels. Following the grid tagging scheme, the entity type label indicates the tail-head token pair of one or more entities belonging to the corresponding category. The “<Frag>” label represents the head-tail token pair of one continuous entity or fragment of a discontinuous entity. Moreover, the “<Gap>” tag stands for the head-tail token pair of a context gap between two neighboring fragments of a discontinuous entity.

During entity decoding, we treat the head and tail token of continuous or discontinuous entities as the start and end nodes, assign a directed edge from the head to tail of each fragment and gap span. Then the entity recognition process is converted into the directed path-searching problem. In order to make the end of the previous edge correspond to the beginning of the next edge, we uniformly move the end of each directed edge to the next token. Considering that several discontinuous entities can possibly share the same tail-head token pair, we employ the BFS (Breadth-First Search) algorithm to search for all valid paths, each of which stands for an entity mention.

One valid path should meet two requirements: 1) the path contains an odd number of spans and the first edge cannot be a gap; 2) fragments and gaps should be arranged alternately if the path has more than one edge. After path searching operation, we obtain the final entity by connecting entity fragments within each valid path in turn. Fig. 4 shows the decoding process of continuous and discontinuous entities.

V. EXPERIMENTS

A. Datasets

To evaluate the performance of our proposed model on discontinuous NER, we choose three public NER datasets in the biomedical domain, namely CADEC [25], ShArE13 [26] and ShArE14 [27]. The CADEC dataset, short for CSIRO Adverse




Fig. 4. Unified decoding process of continuous and discontinuous entities, including path searching and fragment connection. The end nodes are uniformly moved to the next token to match the start nodes of the next edge.

TABLE I
DATA STATISTICS OF THREE NER DATASETS.

	CADEC	ShARe13	ShARe14
Total documents	1250	299	431
Total Sentences	7597	18767	34618
Total Entities	6318	11148	19073
Entity Type	ADE	Disorder	Disorder
Disc. Entities	679	1088	1658
- Percentage	10.75%	9.76%	8.69%
- Number of tokens	1-13	1-7	1-7

Drug Event Corpus, is a rich annotated corpus of patient-reported adverse drug events (ADEs). Its data are sourced from medical forum posts and contain annotated mentions of symptoms, diseases and ADEs. Following previous researches [12], [17], we only use the adverse drug events (ADE) entities for assessments. ShARe13 and ShARe14 datasets, derived from the Shared Annotated Resources (ShARe) project and the CLEF Initiative, are used for the ShARe/CLEF eHealth Challenge in 2013 and 2014, respectively. They consist of many narrative clinical reports extracted from the MIMIC database, which are annotated for the entity recognition and normalization of disorders.

We utilize the preprocessing scripts provided by Dai et al. [17] for data cleaning and splitting. Table 1 presents the statistics of each discontinuous NER dataset. It is worth noting that around 10% entities are discontinuous in each corpus.

B. Baselines and Metrics

We compare the proposed model with baseline methods designed for discontinuous or multi-type NER tasks.

- **Sequence-labeling methods:** **BIOHD** [5] extended the BIO tagging scheme to BIOHD mode and combined LSTM and CRF network to identify discontinuous entities by sequence labeling. **Corro** [6] was recently proposed to improve the sequence tagging scheme through a two-layer representation of discontinuous entities.

- **Span-based methods:** **JointNER** [7] proposed a span-based model which could jointly tackle overlapped and discontinuous entities by learning the neighboring relation between spans.

- **Seq2seq-based methods:** **BartNER** [8] developed an effective seq2seq model with the pointer network to solve three types of NER subtasks. **MAPtr** [28] proposed a generative model for discontinuous NER based on pointer networks, where the pointer indicated the role of current token and the location of next constituent token. **Debias-DDA** [9] and **Debias-CSL** [10] explored the incorrect biases in the generation process and further improved the performance through deconfounding data augmentation and calibrating sequence likelihood, respectively.

- **Grid-tagging methods:** **MAC** [11] first introduced the grid tagging scheme to predict entity fragments and their connecting relations within a discontinuous entity. **W2NER** [12] formulated different NER subtasks into the classification problem of word pair relations and decoded discontinuous entities with grid tags defined as Next-Neighboring-Word (NNW) and Tail-Head-Word (THW). **TOE** [13] improved the tagging scheme of W2NER by adding two grid labels, namely Previous-Neighboring-Word (PNW) and Head-Tail-Word (HTW), to obtain more fine-grained relations and alleviate error propagation.

- **LLM-related methods:** **LLaMA2-13B** [14] developed a cascade instruction tuning method to transfer domain-specific LLMs to biomedical NER tasks. **LLaMA-7B** [15] proposed an effective prompt template for LLMs, which incorporates label-injected instructions for few-shot NER applications.

- **Others:** **SegHyper** [16] employed segmental hypergraphs to extract all candidate spans, then merged them into discontinuous entities. **TBM** [17] was a transition-based end-to-end model with generic neural encoding for discontinuous NER.

For the evaluation metrics, we employ the span-level precision, recall and F1 score, based on the exact matching of entity token indexes and categories.

C. Implementation details

For each dataset, we predefine the token pair grid label set C with “None”, “Frag”, “Gap” and each entity type (e.g., “ADE”). Thus the length of grid label set is $|T| + 3$, where T denotes the initial label set of entity categories.

With regard to model components, we employ the pretrained BioBERT [29] for CADEC, Clinical BERT [30] for ShARe13 and ShARe14, which is the same as other baselines [12], [13] for fair comparisons. We use the AdamW optimizer with the learning rate set to 5e-6 for BERT and 1e-3 for other modules in model training. The total epoch number is 15 and the batch size is set to 12 for CADEC, and 8 for ShARe13 and ShARe14. To avoid overfitting, we apply the dropout rate of 0.5 to each dataset. The best-performing model is taken according to the micro-F1 score on the validation data set. For each experiment, we report the average F1 results with 5 runs of training and testing. All experiments are conducted on a single GeForce RTX 3090 GPU.

TABLE II
OVERALL PERFORMANCE ON DISCONTINUOUS NER DATASETS. † REPRESENTS OUR REPRODUCTION RESULTS VIA THE SOURCE CODES. **BOLD** DENOTES THE BEST SCORE WHILE UNDERLINE INDICATES THE SECOND BEST.

Model		CADEC			ShARe13			ShARe14		
		P	R	F1	P	R	F1	P	R	F1
Sequence-labeling	BIOHD [5]	67.80	64.99	66.36	-	-	-	-	-	-
	Corro [6]	-	-	<u>72.90</u>	-	-	82.10	-	-	81.80
Span-based	JointNER [7]	-	-	69.90	-	-	<u>82.50</u>	-	-	-
Seq2seq-based	BartNER [8]	70.08	71.21	70.64	82.09	77.42	79.69	77.20	83.75	80.34
	MAPtr [28]	75.50	71.80	72.40	87.90	77.20	80.30	-	-	-
	Debias-DDA [9]	71.35	71.86	71.60	81.09	78.13	79.58	77.88	83.77	80.72
	Debias-CSL [10]	72.33	71.01	71.66	81.86	78.48	80.14	78.68	83.63	81.01
LLM-related	LLaMA2-13B [14]	-	-	-	-	-	76.87	-	-	77.54
	LLaMA-7B [15]	-	-	47.10	-	-	-	-	-	-
Others	SegHyper [16]	72.10	48.40	58.00	83.80	60.40	70.30	79.10	70.70	74.70
	TBM [17]	68.90	69.00	69.00	80.50	75.00	77.70	78.10	81.20	79.60
Grid-tagging	MAC [11]	70.50	72.50	71.50	84.30	78.20	81.20	78.20	84.70	81.30
	W2NER† [12]	74.52	70.91	72.67	85.05	79.92	82.41	79.04	84.30	81.58
	TOE† [13]	76.23	69.60	72.77	86.54	78.38	82.26	79.76	84.07	<u>81.86</u>
GapDNER (ours)		74.48	72.83	73.65	85.52	80.39	82.87	81.36	84.13	82.72

D. Main Results

The overall performance of our proposed GapDNER on discontinuous NER is presented in Table II. We could observe that GapDNER consistently outperforms all recent baselines in each dataset. Compared with previous best results, our GapDNER achieves 0.75%, 0.37% and 0.86% F1 improvements in the CADEC, ShARe13 and ShARe14 datasets, respectively, leading to new state-of-the-art results for discontinuous NER. The maximum improvement achieved on the complex dataset ShARe14 demonstrates the superior adaptability of our model to various real-world scenarios. Meanwhile, GapDNER obtains the best recall in CADEC and ShARe13, and the highest precision in ShARe14. Experimental results demonstrate the outstanding capability of the proposed model in recognizing discontinuous entities.

We attribute the performance improvements to several core factors. Different from other grid tagging and span-based methods, we novelly add the learning of context gaps, which can help locate entity fragments and alleviate decoding ambiguity of discontinuous entities. We perform the linear attention computation on internal tokens of each span and capture the structural regularity within both continuous entities and entity fragments. The application of criss-cross attention models the semantic relations of overlapped spans, further enhancing span classification performance. We finally adopt the BFS algorithm to find all possible paths, which generates more valid discontinuous entities and promotes entity recall.

E. Ablation Study

We conduct necessary ablation experiments on three datasets to demonstrate the effectiveness of each component. As presented in Table III, both the biaffine and linear attention

TABLE III
MODEL ABLATION STUDIES WITH SPAN-LEVEL F1 SCORES. **BOLD** AND UNDERLINE INDICATE THE LARGEST AND SECOND LARGEST DECREASES, RESPECTIVELY. “W/O BIAFFINE” DENOTES REPLACING BIAFFINE WITH CONCATENATING HEAD AND TAIL REPRESENTATIONS.

Model	CADEC	ShARe13	ShARe14
GapDNER	73.65	82.87	82.72
w/o BiLSTM	73.47 (0.18↓)	82.63 (0.24↓)	82.44 (0.28↓)
w/o Biaffine	71.85 (<u>1.80</u> ↓)	80.92 (1.95 ↓)	81.30 (1.42↓)
w/o Linear Attn	71.68 (1.97 ↓)	81.35 (<u>1.52</u> ↓)	81.13 (1.59 ↓)
w/o Criss-cross Attn	72.53 (1.12↓)	82.11 (0.76↓)	81.79 (0.93↓)

modules have a significant impact on the model. Removing either of them leads to an obvious decline in performance, which indicates the intra-span regularity module captures typical classification features. Meanwhile, deleting the criss-cross attention module results in a noticeable drop in F1 score, revealing the learning of inter-span relations can further enhance span classification. In addition, the application of BiLSTM also makes a slight contribution to F1 performance.

F. Further Analysis

Effectiveness on Discontinuous Entities As mentioned above, discontinuous entities account for about 10% of all entities in each data set. The evaluation results on total entities may not sufficiently show the strengths of our model in recognizing discontinuous mentions. In this case, we experiment on the subset of original test data where each sentence contains at least one discontinuous entity. From Fig. 5, we find that grid-tagging models significantly outperform seq2seq-based and other methods in identifying discontinuous mentions. More importantly, our GapDNER achieves the highest F1 score compared with grid-tagging and other baselines, demonstrating its effectiveness on discontinuous entity identification.




Fig. 5. Results comparisons on sentences containing discontinuous entities.

Performance on Overlapped Entities We observe that overlapped structures always appear together with discontinuous entities in three biomedical data sets. To further explore the ability of our model to recognize overlapped structures, we conduct analysis experiments on the test data subset where each sentence has at least one pair of overlapped entities. As shown in Fig. 6, our model still achieves the optimal results on each dataset, and shows the most significant improvement on ShARe14, demonstrating its superior ability to handle complex overlapped situations. Moreover, all types of methods perform better on overlapped NER than on discontinuous NER, indicating the former are less challenging than the latter.




Fig. 6. Results comparisons on sentences having overlapped entities.

Impact of Gap Length Longer lengths of context gaps will likely pose more challenges to discontinuous entity recognition. To verify the advantages of our gap-aware model in diverse settings, we further analyze the experimental results of test data on different gap lengths. From Fig. 7, we discover that as the gap length increases, the F1 curves of three methods show similar changing trends and our model performs best in most cases, especially for long gaps. Meanwhile, each curve performs well at two smaller gap lengths (e.g., 2 and 4 for CADEC) and shows a noticeable decline in performance when the gap length is greater than 5. Interestingly, each method performs poorly at the smallest gap length 1, which may be because discontinuous entities are easily mistaken for continuous ones by the model when they have only one word to separate neighboring fragments.

Case Study for Linear and Criss-cross Attention Distributions To demonstrate the intra- and inter-span modules capture useful features for span classification, we select some cases and visualize the linear and criss-cross attention distributions in Fig. 8 and Fig. 9, respectively. As depicted in Fig. 8, we observe that the linear attention module assigns high attention scores to the core words within a type-specific span, such as “pain” in the entity fragment, and the non-entity words




Fig. 7. F1 performance of different gap lengths within discontinuous entities.

“,” and “and” in the context gap. This result indicates the linear attention successfully extracts the internal composition regularity of each span. From Fig. 9, we discover the fragment “pain”, which is in the same column as the target fragment “upper body pain”, is assigned the highest attention score. Meanwhile, the gap “joint ,” located in the same row as the target gap “joint , shoulder and”, obtains the highest score in the right part. In addition, several adjacent cells of the target spans also receive notable scores. These observations confirm that the criss-cross attention module effectively models the semantic relations among overlapped spans.




Fig. 8. Linear attention distributions of the internal tokens within an entity fragment and a context gap.




Fig. 9. Criss-cross attention distributions of an entity fragment and a context gap over the token-pair feature grid, where the blank cells are not involved in the attention computation.

VI. CONCLUSION

In this paper, we propose an effective gap-aware grid tagging model for discontinuous NER, which novelly introduces context gap information to enhance entity fragment classification and resolve ambiguities in entity decoding. To accurately locate and classify spans, we develop two modules to acquire span features from intra- and inter-span aspects. The intra-span regularity extraction module comprehensively

obtains span boundary interactions and internal composition regularities with the biaffine decoder and linear attention computation, respectively. The inter-span relation enhancement module adopts the criss-cross attention to model semantic relations between different spans and further promote span classifications. Experimental results on three biomedical datasets demonstrate that our model obtains new state-of-the-art performance on discontinuous NER and exhibits strong capabilities in recognizing complex entity structures.

ACKNOWLEDGMENT

This work is supported by the National Key Research and Development Program of China (No.2024YFB3309702), the National Nature Science Foundation of China (No.62021002), Tsinghua BNRist and Beijing Key Laboratory of Industrial Big Data System and Application.

REFERENCES

- [1] Phong Le and Ivan Titov, "Improving Entity Linking by Modeling Latent Relations between Mentions," in Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Melbourne, Australia, 2018, pp. 1595–1604.
- [2] William Hogan, Jiacheng Li, and Jingbo Shang, "Fine-grained Contrastive Learning for Relation Extraction," in Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, Abu Dhabi, United Arab Emirates, 2022, pp. 1083–1095.
- [3] Yuhao Yang, Chao Huang, Lianghao Xia, and Chenliang Li, "Knowledge Graph Contrastive Learning for Recommendation," in Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, New York, USA, 2022, pp. 1434–1443.
- [4] Xiang Dai, "Recognizing Complex Entity Mentions: A Review and Future Directions," in Proceedings of ACL 2018, Student Research Workshop, Melbourne, Australia, 2018, pp. 37–44.
- [5] Buzhou Tang, Jianglu Hu, Xiaolong Wang, Qingcai Chen, and Tianyong Hao, "Recognizing Continuous and Discontinuous Adverse Drug Reaction Mentions from Social Media Using LSTM-CRF," *Wirel. Commun. Mob. Comput.*, vol. 2018, Jan. 2018.
- [6] Caio Filippo Corro, "A Fast and Sound Tagging Method for Discontinuous Named-Entity Recognition," in Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, Miami, Florida, USA, 2024, pp. 19506–19518.
- [7] Fei Li, ZhiChao Lin, Meishan Zhang, and Donghong Ji, "A Span-Based Model for Joint Overlapped and Discontinuous Named Entity Recognition," in Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), Online, 2021, pp. 4814–4828.
- [8] Hang Yan, Tao Gui, Junqi Dai, Qipeng Guo, Zheng Zhang, and Xipeng Qiu, "A Unified Generative Framework for Various NER Subtasks," in Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), Online, 2021, pp. 5808–5822, Association for Computational Linguistics.
- [9] Shuai Zhang, Yongliang Shen, Zeqi Tan, Yiquan Wu, and Weiming Lu, "De-Bias for Generative Extraction in Unified NER Task," in Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Dublin, Ireland, 2022, pp. 808–818, Association for Computational Linguistics.
- [10] Yu Xia, Yongwei Zhao, Wenhao Wu, and Sujian Li, "Debiasing Generative Named Entity Recognition by Calibrating Sequence Likelihood," in Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers), Toronto, Canada, 2023, pp. 1137–1148.
- [11] Yucheng Wang, Bowen Yu, Hongsong Zhu, Tingwen Liu, Nan Yu, and Limin Sun, "Discontinuous Named Entity Recognition as Maximal Clique Discovery," in Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), Online, 2021, pp. 764–774.
- [12] Jingye Li, Hao Fei, Jiang Liu, Shengqiong Wu, Meishan Zhang, Chong Teng, Donghong Ji, and Fei Li, "Unified Named Entity Recognition as Word-Word Relation Classification," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, no. 10, pp. 10965–10973, 2022.
- [13] Jiang Liu, Donghong Ji, Jingye Li, Dongdong Xie, Chong Teng, Liang Zhao, and Fei Li, "TOE: A Grid-Tagging Discontinuous NER Model Enhanced by Embedding Tag/Word Relations and More Fine-Grained Tags," *IEEE/ACM Trans. Audio, Speech and Lang. Proc.*, vol. 31, pp. 177–187, Nov. 2022.
- [14] J. Zhao, C. Liu, J. Liang, Z. Li and Y. Xiao, "A Novel Cascade Instruction Tuning Method for Biomedical NER," in ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing, Seoul, Korea, Republic of, 2024, pp. 11701–11705.
- [15] Xingyu Zhu, Feifei Dai, Xiaoyan Gu, Bo Li, Meiou Zhang, and Weiping Wang, "GL-NER: Generation-Aware Large Language Models for Few-Shot Named Entity Recognition," in *Artificial Neural Networks and Machine Learning – ICANN 2024*, Cham, 2024, pp. 433–448.
- [16] Bailin Wang and Wei Lu, "Combining Spans into Entities: A Neural Two-Stage Approach for Recognizing Discontiguous Entities," in Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, Hong Kong, China, 2019, pp. 6216–6224.
- [17] Xiang Dai, Sarvnaz Karimi, Ben Hachey, and Cecile Paris, "An Effective Transition-based Model for Discontinuous NER," in Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, Online, 2020, pp. 5860–5870.
- [18] Jonathan Ho, Nal Kalchbrenner, Dirk Weissenborn, and Tim Salimans, "Axial Attention in Multidimensional Transformers," *arXiv preprint arXiv:1912.12180*, 2019.
- [19] Z. Huang et al., "CCNet: Criss-Cross Attention for Semantic Segmentation," in Proceedings of the IEEE/CVF international conference on computer vision, Seoul, Korea (South), 2019, pp. 603–612.
- [20] J. Jiao et al., "DilateFormer: Multi-Scale Dilated Transformer for Visual Recognition," in *IEEE Transactions on Multimedia*, vol. 25, pp. 8906–8919, 2023.
- [21] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," in Proceedings of NAACL 2019: Human Language Technologies, Volume 1 (Long and Short Papers), Minneapolis, Minnesota, 2019, pp. 4171–4186.
- [22] Zhiheng Huang, Wei Xu, and Kai Yu, "Bidirectional LSTM-CRF Models for Sequence Tagging," *arXiv preprint arXiv:1508.01991*, 2015.
- [23] Juntao Yu, Bernd Bohnet, and Massimo Poesio, "Named Entity Recognition as Dependency Parsing," in Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, Online, 2020, pp. 6470–6476.
- [24] Yingjie Gu, Xiaoye Qu, Zhefeng Wang, Yi Zheng, Baoxing Huai, and Nicholas Jing Yuan, "Delving Deep into Regularity: A Simple but Effective Method for Chinese Named Entity Recognition," in Findings of NAACL 2022, Seattle, United States, 2022, pp. 1863–1873.
- [25] Sarvnaz Karimi, Alejandro Metke-Jimenez, Madonna Kemp, and Chen Wang, "CadeC: A corpus of adverse drug event annotations," *Journal of biomedical informatics*, vol. 55, pp. 73–81, 2015.
- [26] Sameer Pradhan, Noemie Elhadad, Brett R South, David Martinez et al., "Task 1: ShARe/CLEF eHealth Evaluation Lab 2013," *CLEF (working notes)*, vol. 1179, 2013.
- [27] Danielle L. Mowery, Sumithra Velupillai, Brett R. South, Lee Christensen, David Martinez, et al., "Task 2: ShARe/CLEF eHealth Evaluation Lab 2014," in Proceedings of CLEF 2014, Sheffield, United Kingdom, Sep 2014.
- [28] Hao Fei, Donghong Ji, Bobo Li, Yijiang Liu, Yafeng Ren, and Fei Li, "Rethinking Boundaries: End-To-End Recognition of Discontinuous Mentions with Pointer Networks," in Proceedings of the AAAI Conference on Artificial Intelligence, 2021, vol. 35, pp. 12785–12793.
- [29] Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang, "BioBERT: a pre-trained biomedical language representation model for biomedical text mining," *Bioinformatics*, Vol. 36, no. 4, pp. 1234–1240, 2019.
- [30] Emily Alsentzer, John Murphy, William Boag, Wei-Hung Weng, Di Jindi, Tristan Naumann, and Matthew McDermott, "Publicly Available Clinical BERT Embeddings," in Proceedings of the 2nd Clinical Natural Language Processing Workshop, Minneapolis, Minnesota, USA, 2019, pp. 72–78.