

# One Model for All Domains: Collaborative Domain-Prefix Tuning for Cross-Domain NER

Xiang Chen<sup>1</sup>, Lei Li<sup>1</sup>, Shuofei Qiao<sup>1</sup>, Ningyu Zhang<sup>1\*</sup>, Chuanqi Tan<sup>3</sup>,  
Yong Jiang<sup>3</sup>, Fei Huang<sup>3</sup> and Huajun Chen<sup>1,2\*</sup>

<sup>1</sup>Zhejiang University & AZFT Joint Lab for Knowledge Engine

<sup>2</sup>Donghai Laboratory

<sup>3</sup>DAMO Academy, Alibaba Group

{xiang\_chen, leili21, shuofei, zhangningyu, huajunsir}@zju.edu.cn  
{chuanqi.tcq, jiangyong.ml, f.huang}@alibaba-inc.com

## Abstract

Cross-domain NER is a challenging task to address the low-resource problem in practical scenarios. Previous typical solutions mainly obtain a NER model by pre-trained language models (PLMs) with data from a rich-resource domain and adapt it to the target domain. Owing to the mismatch issue among entity types in different domains, previous approaches normally tune all parameters of PLMs, ending up with an entirely new NER model for each domain. Moreover, current models only focus on leveraging knowledge in one general source domain while failing to successfully transfer knowledge from multiple sources to the target. To address these issues, we introduce Collaborative Domain-Prefix Tuning for cross-domain NER (CP-NER) based on text-to-text generative PLMs. Specifically, we present text-to-text generation grounding domain-related instructors to transfer knowledge to new domain NER tasks without structural modifications. We utilize frozen PLMs and conduct collaborative domain-prefix tuning to stimulate the potential of PLMs to handle NER tasks across various domains. Experimental results on the Cross-NER benchmark show that the proposed approach has flexible transfer ability and performs better on both one-source and multiple-source cross-domain NER tasks.

## 1 Introduction

Named entity recognition (NER) is an important task for Knowledge Graph (KG) construction [Zhang *et al.*, 2022] and natural language processing (NLP). Owing to the data scarcity issue in practical scenarios, obtaining adequate domain-related data is usually labour-intensive. The naive idea of training models with rich-resource domain data (source) and transferring knowledge to a new specific domain (target) may struggle handling the semantic gap and limited data problem. Hence, cross-domain NER, capable of learning information from the source domain to specific target domains with limited data, has been proposed to alleviate this issue.

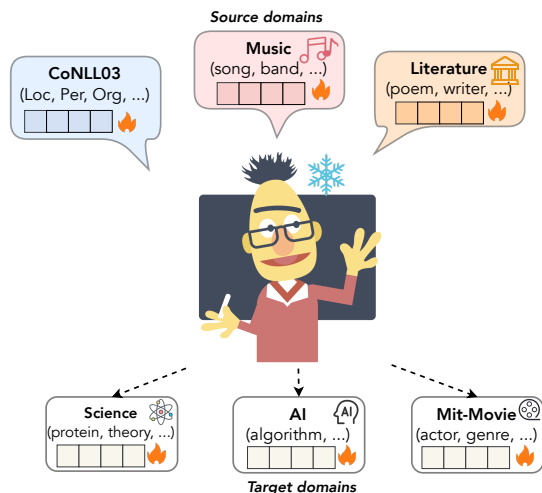


Figure 1: One model for all domains with collaborative prefixes.

To address this task, some previous studies utilize adding auxiliary objects [Liu *et al.*, 2020a; Wang *et al.*, 2020] or designing new model architectures [Jia *et al.*, 2019a; Liu *et al.*, 2020b; Jia and Zhang, 2020] to train with both source and target domain data. Liu *et al.* [2021a] leverages continue pre-training on the target domain for a better understanding of the domain-specific data. Another line of researches [Zheng *et al.*, 2022; Hu *et al.*, 2022] focus on modeling the label relationship across domains to improve label information transfer. Typically, LANER [Hu *et al.*, 2022], the current state-of-the-art (SOTA) method on CrossNER benchmark [Liu *et al.*, 2021b], utilizes an architecture to better leverage the semantic relationships among domains for boosting the cross-domain performance.

Despite the empirical success of previous works, several issues remain which have not been appropriately solved. **Firstly**, previous methods often rely on task-specific architectures for various domains with different entity categories. For instance, LANER [Hu *et al.*, 2022] designs different model architectures to adapt a PLM to other target domains, which restricts the model’s usability in more applications. **Secondly**, most current methods are computationally inefficient

and require tuning all parameters of the PLMs, which ends up with an entirely new NER model for each domain. For example, storing a full copy of pretrained BERT\_Large [Devlin *et al.*, 2019] (340M parameters) for each domain is non-trivial, not to mention the billion-scale large LMs. We introduce a summary of the above two points **from the perspective of practical applications**: *when a new domain data comes, it is not practical for users to design model architectures and tune all parameters in terms of computational resources and time.* **Thirdly**, semantic knowledge transfer is essential for cross-domain NER, but previous works typically can only transfer from single source domain to the target domain and lack the ability of utilizing knowledge from multiple domains to help each other.

All the considerations above lay down our goal to investigate simple yet effective ways to keep one frozen set of PLM parameters for all domains and synthesize knowledge from multiple domains to enhance the target domain performance, as shown in Figure 1. In this work, we enable collaborative domain-prefix tuning for cross-domain NER tasks. As an alternative to fine-tuning the whole PLM, our method, namely CP-NER<sup>1</sup>, tunes no existing PLM parameters. We establish this by casting the NER into a “text-to-text generation grounding domain-related instructor”. Compared with LIGHTNER [Chen *et al.*, 2022a] that generates indexes of entity spans, our formulation enables eliciting general knowledge about named entities from PLMs and avoids the external modification of model architecture. By introducing collaborative domain-prefix tuning consisting of a few trainable parameters, we can flexibly adapt the knowledge from multiple source domains to target domains. The specific process consists of three steps: (i) *Domain-specific Warmup*. This step leverages domain corpus to warm up the corresponding prefix; (ii) *Dual-query Domain Selector*. Identify which domain could benefit each other from both the perspectives of label similarity and prefix similarity; (iii) *Intrinsic Decomposition for Collaborative Domain-Prefix*. This module helps synthesize the final powerful prefix from multiple source domains. In a nutshell, the contributions of our CP-NER are as follows:

- We reformulate NER as the “text-to-text generation grounding domain-related instructor”. This new formulation helps trigger PLMs to produce general knowledge about NER tasks and can handle different entity categories without any domain-specific modification of parameters for different domains, laying the groundwork of **one model for all domains**.
- Our proposal of collaborative domain-prefix tuning for PLMs aims to adapt the knowledge to cross-domain NER tasks. To the best of our knowledge, our work takes the first step to leverage knowledge transfer from multiple source domains specifically for cross-domain NER. The comprehensive nature of the results is crucial for a wide range of cross-domain information extraction applications.
- We obtain SOTA performance on the CrossNER bench-

<sup>1</sup>Codes will be available in <https://github.com/zjunlp/DeepKE/tree/main/example/ner/cross>.

mark involving transferring knowledge from a single general domain to the target domain. Extensive experiments illustrate that our CP-NER can also efficiently leverage the knowledge transfer via learning lightweight parameters and unified architecture for all domains.

## 2 Related Work

### 2.1 Cross-Domain NER

Although PLMs [Devlin *et al.*, 2019; Liu *et al.*, 2019; Lewis *et al.*, 2020; Xue *et al.*, 2021] have made waves in NLP, achieving overwhelming performance in widespread NER datasets, yet they require many domain-related labeled data for training when adapting to the target domains. To this end, cross-domain NER algorithms [Kim *et al.*, 2015; Yang *et al.*, 2018; Lee *et al.*, 2018; Lin and Lu, 2018; Liu *et al.*, 2020a] that alleviate the data scarcity issue and boost the models’ generalization ability to target domains have drawn substantial attention recently. One approach to this problem is multitask learning, which includes adding auxiliary objects [Liu *et al.*, 2020a; Wang *et al.*, 2020] or designing new model architectures [Jia *et al.*, 2019a; Liu *et al.*, 2020b; Jia and Zhang, 2020] for improving the NER performance of the target domain by training on data from both source and target domain. However, these methods typically require training with vast amounts of labeled source domain data to obtain satisfactory target domain performance. Yet our approach leverages collaborative domain-prefix tuning, which allows for synthesizing parametric knowledge across domains and holds one frozen PLM for all domains.

Another area of research in cross-domain NER focuses on transferring label information across domains. Zheng *et al.* [2022] models the relation of labels as a probability distribution to better transfer cross-domain knowledge in NER. Hu *et al.* [2022] proposes an approach to better utilize the semantic relations among domains that jointly predict NER labels by the previous labels (from source) and corresponding token. In contrast, our approach uses descriptive instructions with entity categories’ semantics to enable the PLMs to understand NER tasks across different domains without modifying the model architecture. This allows for a single model to be used for all domains, laying the foundation for an efficient and versatile NER system.

### 2.2 Parameter-Efficient Tuning for PLMs

Pre-trained language models (PLMs) [Devlin *et al.*, 2019; Liu *et al.*, 2019; Lewis *et al.*, 2020; Xue *et al.*, 2021] have been widely used to achieve excellent performance in various NLP tasks [Wang *et al.*, 2018; Wang *et al.*, 2019]. However, fine-tuning these models requires lots of computational resources and storage. To overcome this challenge, researchers have developed parameter-efficient methods to transfer PLMs to new tasks with small trainable parameters [Ding *et al.*, 2022]. These methods include prompt-tuning [Shin *et al.*, 2020; Chen *et al.*, 2022b], which fine-tunes the model using task-specific prompts; and adaptors [Houlsby *et al.*, 2019; Wang *et al.*, 2021; Newman *et al.*, 2022] which add task-specific parameters to the PLMs. However, applying these

methods to new domains with unseen entity types can be challenging. One type of solution is to utilize prefix-tuning [Li and Liang, 2021a], which prepends a sequence of task-specific vectors to the input. More works explore task adaptation with soft prompt [Vu *et al.*, 2022; Su *et al.*, 2022].

Different from those approaches, our work, referred to as CP-NER, stands out from those approaches in several ways. Firstly, we use collaborative knowledge updating between source and target prefixes to capture domain-related knowledge rather than just tuning extra parameters. Secondly, our method allows for easy redeployment when switching to new domains with unseen entity types by using entity categories as grounding for text-to-text PLMs, unlike previous parameter-efficient methods. Finally, our approach aims to utilize multiple source domains to enhance the target domain, specifically focusing on NER, as opposed to transfer across tasks.

### 3 Methodology

#### 3.1 Problem Definition

Given sentence  $\mathbf{x} = \{w_1, w_2, \dots, w_n\}$ , the NER task is aimed to extracting all entities  $\mathbf{y}$ , and  $n$  refers to the length of  $\mathbf{x}$ . Formally, we have the  $i$ -th entity  $y_i \in \mathbf{y}$  as  $y_i = (\mathbf{x}_{l:r}, t)$ , where  $t \in \mathcal{T}$  refers to the type of entity.  $r$  and  $l$  refers to the right and left entity boundary indexes, respectively. Given the source domain data  $\mathcal{D}_{src}$  as well as target domain data  $\mathcal{D}_{tgt}$ , the cross-domain NER task aims to obtain transferable knowledge with  $\mathcal{D}_{src}$  to boost the performance of  $\mathcal{D}_{tgt}$ . Previous studies on cross-domain NER mainly focus on transferring from one source domain to one target domain. While we not only consider “**One Source for Target**”, but also include the “**Multiple Source for Target**” setting, which is practical for on-the-fly cross-domain NER in real scenarios.

#### 3.2 Text-to-Text Generation Grounding Domain-related Instructor

We formulate the cross-domain NER task as sequence-to-sequence generation and leverage a **frozen** T5 [Raffel *et al.*, 2020], as shown in Figure 2. Specifically, for the NER task, we construct the inputs containing the following components:

- **Sentence** - the input sentence  $\mathbf{x}$ .
- **Instruction** - the instruction ( $s$ ) tells the model how to solve NER task. We ask the model to generate a sequence according to the introduction.
- **Domain-related Options** - the option ( $o$ ) includes all entity types  $\mathcal{T}$  of the current domain and acts as both a constraint and a hint.

Formally, CP-NER takes the given instructor ( $s$ ), domain-related option ( $o$ ) and the text sequence ( $\mathbf{x}$ ) as input and generates the output ( $\mathbf{y}$ ), which contains the extracted information from input sentence based on domain-related instructor:

$$\mathbf{y} = \text{LM}_\theta(s; o; \mathbf{x}), \tag{1}$$

where  $\theta$  denotes parameters of the  $\text{LM}_\theta$  (we adopt T5 model in this paper). We convert the output sequence into a natural language consistent with the input instructions. To be specific, we leverage simple templates to reformulate the entity  $(\mathbf{x}_{l:r}, t)$

to natural language by concatenating entity types, mentions and several special tokens (separators). For example, given the sentence “Johnson, who played eight seasons in Baltimore, was named Orioles manager in the off-season replacing Phil Regan.”, the model will generate “((person: Johnson) (location: Baltimore) (organization: Orioles))”.

We argue that the text-to-text generation grounding domain-related instructor can: 1) effectively guide the T5 model to generate named entity sequence with **domain-specific entity categories** so that it can be transferred to new domain NER tasks without any structural modification of the T5 model; 2) stimulate the potential of PLM to handle NER tasks with various domains, which laid the foundation for guiding frozen PLM to generate entity sequences through prefix-tuning.

#### 3.3 Insights of Tuned Prefix as Domain Controller

Prefix-tuning [Li and Liang, 2021b] prepends trainable continuous tokens, also known as *soft prompts*, to each Transformer layer’s hidden states. We regard the prefix as the domain controller for cross-domain NER.

As described in Li and Liang [2021b], directly updating the prefix parameters leads to unstable optimization and a slight drop in performance. So the prefix of the  $j$ -th layer,  $\delta^{(j)}$ , is obtained based on the trainable matrix,  $\mathbf{P}^{(j)}$ , while all the other parameters of PLM will be fixed. In particular, we derive the variant formula of self-attention with the prefix at the  $j$ -th layer as<sup>2</sup>:

$$\begin{aligned} \text{head}^{(j)} &= \text{softmax}(\mathbf{x}^{(j)} \mathbf{W}_q^{(j)} [\delta_k^{(j)}; \mathbf{x}^{(j)} \mathbf{W}_k^{(j)}]^\top) \begin{bmatrix} \delta_v^{(j)} \\ \mathbf{x}^{(j)} \mathbf{W}_v^{(j)} \end{bmatrix} \\ &= (1 - \lambda(\mathbf{x}^{(j)})) \underbrace{\text{Attn}(\mathbf{x}^{(j)} \mathbf{W}_q^{(j)}, \mathbf{x}^{(j)} \mathbf{W}_k^{(j)}, \mathbf{x}^{(j)} \mathbf{W}_v^{(j)})}_{\text{standard attention}} \\ &\quad + \lambda(\mathbf{x}^{(j)}) \underbrace{\text{Attn}(\mathbf{x}^{(j)} \mathbf{W}_q^{(j)}, \delta_k^{(j)}, \delta_v^{(j)})}_{\text{attention of domain-specific prefix}}, \end{aligned} \tag{2}$$

where  $\lambda(\mathbf{x}^{(j)})$  refers to the scalar regarding the sum of normalized attention weights on the key and value embedding from prefixes.

For simplicity, we denote the prefix indexes as  $P_{\text{idX}}$ , then prefix-tuning [Li and Liang, 2021b] formulates the forward propagation at the position  $n$  as:

$$h_n^{(j+1)} = \text{LM}_\theta^{(j)} \left( h_n^{(j)} \mid h_{<n}^{(j)} \right), \tag{3}$$

where  $h_i$  is an activation of the trainable parameter  $\mathbf{P}$ ,  $z = [\text{PREFIX}, \mathbf{x}]$ ,  $h_i^{(j)} = \delta^{(j)}[i, \cdot]$  for all  $j = 0$  to  $L - 1$ ,  $i \in P_{\text{idX}}$  and  $h_i^{(0)} = z_i$  for  $i \notin P_{\text{idX}}$ . With the large-scale LM parameters fixed and continuous optimization on training samples, prefix-tuning is expected to steer the LMs to generate correct entity prediction for test data.

**Remark 1.** As shown in Eq. 2, the domain-specific prefix tuning essentially modifies the original head attention through

<sup>2</sup>Without loss of generalization, we ignore the scaling factor  $\sqrt{d}$  of the Softmax for the ease of representation.

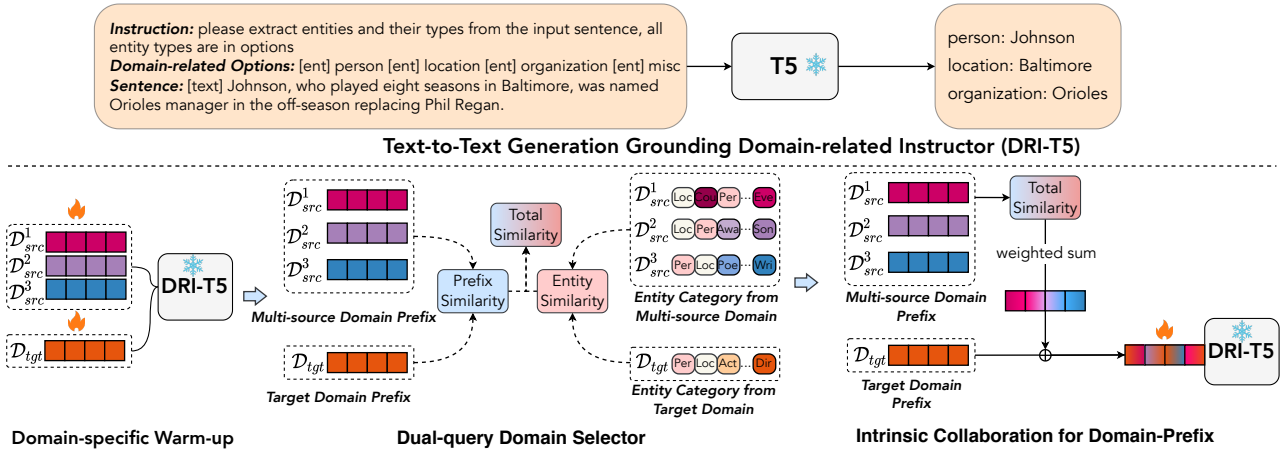


Figure 2: The overall architecture of the proposed CP-NER for cross-domain NER. DRI denote the Domain-Related Instructor.

linear interpolation by a scalar factor (i.e.,  $1 - \lambda$ ), acting as a domain controller to drive the model to output domain-specific predictions. However, prefix-tuning still suffers from knowledge forgotten or dilution in the process of direct domain transfer, which inspires us to further explore more effective ways of joining prefix hints to leverage knowledge from multiple source domains.

### 3.4 Collaborative Domain-prefix Tuning

Directly transferring knowledge through prefix-tuning can be challenging because the knowledge learned in the source prefix can easily be forgotten or diluted when fine-tuned on data from the target domain due to differences in text styles and entity categories. To address this issue and help the T5 model’s domain controller more effectively capture knowledge from source domains, we propose a collaborative domain-prefix tuning method.

The overall process is illustrated in Figure 2: (1) domain-related warm-up to capture knowledge from the corpus of the current domain; (2) dual-query domain selector to determine the proportion of prefix knowledge from multiple source domains; (3) intrinsic decomposition for collaborative domain-prefix, which can flexibly synthesize prefix knowledge from source and target domains.

**Domain-specific warm-up.** To flexibly organize prefixes with knowledge from different domains, we first leverage domain corpus to warm up the corresponding prefixes. Formally, we warm up the newly initialized trainable parameter matrix  $\{\mathbf{P}^{(0)}, \dots, \mathbf{P}^{(L-1)}\}$  for all  $L$  layers of frozen T5 where the final prefix of  $l$ -th layer is drawn from the parameter matrix  $\mathbf{P}^{(l)}$ . We optimize the parameter matrix of prefix at each layer  $\{\mathbf{P}^{(0)}, \dots, \mathbf{P}^{(L-1)}\}$  using the training set  $\mathcal{D}_{tr}$ . We follow the log-likelihood objective in T5 [Raffel *et al.*, 2020] and abbreviate the optimization objective as follows:

$$\min_{\{\mathbf{P}^{(0)}, \dots, \mathbf{P}^{(L-1)}\}} \mathbb{E}_{(x,y) \sim \mathcal{D}_{tr}} \left[ \sigma \left( h_n^{(L)}, y \right) \right] \quad (4)$$

where  $\sigma$  as the softmax scoring function that output  $h_o^{(L)}$  to a probability vector over the vocabulary.  $h_i^{(j)}$  refers to the

intermediate activation of the  $j$ -th layer at step  $i$ . We utilize the above optimization objective over the domain corpus to warm up the domain-specific parameter matrix and then generate the domain controller  $\{\delta^{(0)}, \dots, \delta^{(L-1)}\}$ .

**Dual-query domain selector.** When the source and target domains share part of the same entity categories, it is intuitive to leverage those shared labels having similar semantic information for flexible adaptation to the target domain. Besides, domain-specific prefixes maintain different grammar styles and themes. Based on the above consideration, we design a dual-query domain selector from both the perspective of label similarity and prefix similarity to identify which domains could benefit each other.

- **Entity Similarity:** we measure entity similarity through the embedding of the T5 model. Denote the word embedding of T5 as  $LM_w$ ; we feed entity categories of source and target domain into the  $LM_w$  to obtain the semantic category representations of the  $i$ -th source domain as  $E_{src}^i \in \mathbb{R}^{|\mathcal{T}_{src}^i| \times d}$  and target domain as  $E_{tgt} \in \mathbb{R}^{|\mathcal{T}_{tgt}| \times d}$ . Then we conduct pooling over the representation and calculate their similarity matrix as follows:

$$\mathcal{S}_e^i = \cos(\text{Pool}(E_{src}^i), \text{Pool}(E_{tgt})), \quad (5)$$

where  $\text{Pool}$  is the pooling operation function and  $\mathcal{S}_e^i \in \mathbb{R}$ ;

- **Prefix Similarity:** given the prefix  $\delta_{src} = \{\delta_{src}^{(0)}, \dots, \delta_{src}^{(L-1)}\}$  and  $\delta_{tgt} = \{\delta_{tgt}^{(0)}, \dots, \delta_{tgt}^{(L-1)}\}$  of source and target domain, we also apply cosine similarity function over prefixes to calculate prefix similarity as:

$$\mathcal{S}_p^i = \text{Avg}(\cos(\delta_{src}, \delta_{tgt})), \quad (6)$$

where  $\mathcal{S}_p^i \in \mathbb{R}$  and  $\text{Avg}$  represents the average function.

For the target domain  $\mathcal{D}_{tgt}$ , we denote  $\alpha$  as the hyper-parameter  $\alpha \in [0 : 1]$  to determine the ratio between the dual-query similarities as follows:

$$\mathcal{S}_{total}^i = \alpha^i \mathcal{S}_e^i + (1 - \alpha^i) \mathcal{S}_p^i. \quad (7)$$

As for leveraging  $I$  source domains for collaborative prefix tuning, we apply the softmax function over the total similarity of multiple source domains as follows:

$$\mathbf{S}_{total} = \sigma([\mathbf{S}_{total}^1; \dots; \mathbf{S}_{total}^I]), \quad (8)$$

To pursue compositional generalization and differentiable optimization, we encourage the prefix to be continuously activated and combined. Every single prefix from source domains can be viewed as some fundamental skill and domain-specific knowledge, and NER tasks from another domain can be solved by combining such modular prefixes. Though similar ideas have been proposed in other names and contexts [Sun *et al.*, 2020; Ponti *et al.*, 2022], this is the first work that implements the cross-domain NER with multiple source domains to drive PLMs to recognize named entities.

**Intrinsic collaboration for domain-prefix.** Based on the selection of source domains to benefit target domains, we propose intrinsic collaboration for domain-prefix, which can synthesize the final powerful prefix from multiple source domains. Specifically, for the warm-up parameter matrix  $\mathbf{P}^{(j)} \in \mathbb{R}^d$  from the source domain, we argue it is in the same intrinsic space as the prefix from target domain. We perform *intrinsic refactor* such as addition operation to the warm-up parameter matrices of the target domain, as shown in Figure 2. The specific parameter matrix  $\mathbf{P}_{tgt}^{(j)}$  of target prefix at the  $j$ -th layer is refactored as:

$$\hat{\mathbf{P}}_{src}^{(j)} = \mathbf{S}_{total} \cdot [\mathbf{P}_{src_1}^{(j)}; \dots; \mathbf{P}_{src_I}^{(j)}], \quad (9)$$

$$\hat{\mathbf{P}}_{tgt}^{(j)} \leftarrow (\mathbf{P}_{tgt}^{(j)} + \hat{\mathbf{P}}_{src}^{(j)}) / 2, \quad (10)$$

where  $\hat{\mathbf{P}}_{src}^{(j)}$  denote the aggregated matrix from multiple domains<sup>3</sup>. We further follows the objective function in Eq.4 to update  $\hat{\mathbf{P}}_{tgt}^{(j)}$  while keeping  $\text{LM}_\theta$  frozen.

**Remark 2.** *From the optimal control (OC) perspective, prefix-tuning can be formalized as seeking the OC of the pre-trained LM for specific domains. And our collaborative domain-prefix tuning can be interpreted as seeking the close-loop control for leveraging knowledge from multiple-source to enhance the target domain NER performance.*

## 4 Experiments

### 4.1 Experimental Settings

**Dataset.** We conduct experiments using four publicly available datasets, including CrossNER [Liu *et al.*, 2021b], CoNLL 2003 [Sang and Meulder, 2003], MIT Restaurant [Liu *et al.*, 2013a], and MIT Movie [Liu *et al.*, 2013b]. CoNLL 2003 is a well-known NER dataset that serves as the source domain, which is labeled with four categories: PER, LOC, ORG, and MISC. Additionally, CrossNER contains five separate domain datasets: politics, natural science, music, AI, and literature. We adhere to the official splits for training, validation, and test sets, and the statistics for these datasets can be found in the supplementary materials.

<sup>3</sup>For scenarios with only one source domain, we let  $\hat{\mathbf{P}}_{src}^{(j)} = \mathbf{P}_{src}^{(j)}$

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

- DAPT [Liu *et al.*, 2021b]: DAPT employs a large unlabeled corpus related to a specific domain, which is based on BERT and fine-tunes it on the CrossNER benchmark.
- COACH [Liu *et al.*, 2020b]: This method utilizes patterns of slot entities and combines the features for each slot entity in order to improve the accuracy of entity type predictions.
- CROSS-DOMAIN LM [Jia *et al.*, 2019b]: This method utilizes a parameter generation network to merge cross-domain language modeling with NER, resulting in improved model performance.
- FLAIR [Akbik *et al.*, 2018]: This method utilizes the internal states of a character-level language model to generate contextual string embeddings, which are integrated into the NER model.
- BARTNER [Yan *et al.*, 2021]: This approach uses the pre-trained BART model to generate entity spans, treating the NER task as a sequence generation problem.
- LST-NER [Zheng *et al.*, 2022]: This approach models the relationship between labels as a probability distribution and builds label graphs in both the source and target label spaces for cross-domain NER tasks.
- LANER [Hu *et al.*, 2022]: This method introduces a new approach for cross-domain named entity recognition by utilizing an autoregressive framework to strengthen the connection between labels and tokens.
- LIGHTNER [Chen *et al.*, 2022a]: This utilizes a plug-gable prompting method to improve NER performance in low-resource settings.

Due to all baselines being reported in **base PLMs**, we utilize **T5-base** to conduct a fair comparison for all experiments.

### 4.2 Transfer from Single Source Domain

As we can see from the **w/o DAPT** column of Table 1, CP-NER can consistently obtain better performance than SoTA method LANER across all target domains, with a 3.56% F1-score improvement in the *Source & Target* setting. Moreover, we find that CP-NER surpasses LANER on the science and music domains, with a 6% to 8% F1-score performance gains on each domain. Experimental results demonstrate that our CP-NER technique achieves superior performance on the CrossNER benchmark. Our approach beats explicitly naive baselines and is on par with the LANER+DAPT model, which adopts an external domain-related corpus for pre-training.

On the other hand, we also apply DAPT on our CP-NER to conduct a fair comparison with other baselines. As **Introducing DAPT** column of Table 1 illustrates, our method (CP-NER) can better recognize and categorize entities on these domains with DAPT. This is because it leverages many domain-related data, which can boost performance. When introducing DAPT, we observe that CP-NER also outperforms other baselines in all domains, revealing that our method is compatible with DAPT.

MODELS	CONLL2003					AVERAGE
	POLITICS	SCIENCE	MUSIC	LITERATURE	AI	
<b>w/o DAPT</b>						
COACH [Liu <i>et al.</i> , 2020b]	61.50	52.09	51.66	48.35	45.15	51.75
CROSS-DOMAIN LM [Jia <i>et al.</i> , 2019b]	68.44	64.31	63.56	59.59	53.70	61.92
FLAIR [Akbik <i>et al.</i> , 2018]	69.54	64.71	65.60	61.35	52.48	62.73
BARTNER-BASE [Yan <i>et al.</i> , 2021]	69.90	65.14	65.35	58.93	53.00	62.46
LST-NER [Zheng <i>et al.</i> , 2022]	70.44	66.83	72.08	67.12	60.32	67.36
LANER [Hu <i>et al.</i> , 2022]	71.65	69.29	73.07	67.98	61.72	68.74
LIGHTNER [Chen <i>et al.</i> , 2022a]	72.78	66.74	72.28	65.17	35.82	62.56
<b>CP-NER</b>	<b>73.41</b>	<b>74.65</b>	<b>78.08</b>	<b>70.84</b>	<b>64.53</b>	<b>72.30</b>
<b>Introducing DAPT</b>						
DAPT [Liu <i>et al.</i> , 2021b]	72.05	68.78	75.71	69.04	62.56	69.63
LST-NER+DAPT	73.25	70.07	76.83	70.76	63.28	70.84
LANER+DAPT	74.06	71.83	78.78	71.11	65.79	72.31
<b>CP-NER+DAPT</b>	<b>74.25</b>	<b>75.82</b>	<b>79.10</b>	<b>72.17</b>	<b>67.95</b>	<b>73.86</b>

Table 1: Comparisons of existing studies and our CP-NER in terms of F1 scores are provided. The AVERAGE indicates the average F1 score across five domains in the CrossNER benchmark using only a single source (CONLL2003). The baseline results are cited directly from the LANER and LST-NER papers.

SOURCE	METHODS	SCI.	AI	MIT Mov.			MIT RES.		
		All	All	10	20	50	10	20	50
NONE	LANER [Hu <i>et al.</i> , 2022]	67.79	58.65	49.25	54.48	70.96	45.15	49.14	65.10
	LIGHTNER [Chen <i>et al.</i> , 2022a]	49.94	15.02	12.90	39.18	67.03	5.70	23.03	51.15
	<b>CP-NER</b>	<b>70.24</b>	<b>60.24</b>	<b>60.01</b>	<b>67.21</b>	<b>81.57</b>	<b>49.42</b>	<b>63.24</b>	<b>75.92</b>
CONLL.+ POL. + LIT. + MUS.	LANER <Chain Transfer>	67.82	55.60	50.65	60.63	72.60	38.32	45.56	58.38
	LANER <Ensemble>	69.97	59.47	53.55	63.01	72.85	44.79	50.99	62.30
	LIGHTNER <Chain Transfer>	67.49	40.91	51.77	66.19	74.44	48.13	52.97	65.13
	LIGHTNER <Ensemble>	58.69	26.25	32.01	52.63	67.29	26.50	57.09	67.07
	<b>CP-NER &lt;Multiple Source&gt;</b>	<b>75.33</b>	<b>64.90</b>	<b>70.74</b>	<b>73.22</b>	<b>81.90</b>	<b>65.14</b>	<b>70.46</b>	<b>77.34</b>

Table 2: Model performance, measured by the F1 score, in the setting of transferring from multiple source domains to the target domains.

### 4.3 Transfer from Multiple Source Domains

Previous work for cross-domain NER only focuses on transferring from a single source domain to a target domain, failing to leverage multiple source domains, even though numerous source domains are more in line with practical application scenarios. As shown in Table 2, we experiment with transferring from multiple source domains. We use the CoNLL 2003, Politics, Literature, and Music datasets as source domains and test on the Mit-Movie, AI, Science, and Mit-Restaurant target domains. To compare our CP-NER with other baselines, we propose the following setting for multiple source domains.

**Setting.** (i) <Ensemble>. Train on a single source domain and transfer to the target domain, then ensemble the models transferred from multiple source domains. (ii) <Chain Transfer>. Continuously train on multiple source domains in sequence and then transfer to the target domain. (iii) CP-NER<Multiple Source>. Leverage the warm-up domain-specific prefix to conduct collaborative domain-prefix tuning.

**Multiple source makes better transfer learning.** According to the results of Table 1 and Table 2, it can be observed that using multiple sources provides consistent improvement for baseline models. However, only LIGHTNER<Ensemble> and LANER<Chain Transfer> achieve little improvements compared with the results in the single source domain setting, revealing that <Ensemble> may be more suitable for lightweight tuning. In contrast, <Chain Transfer> is ideal for fully-tuned models. Furthermore, our CP-NER using collab-

orative prefix tuning empowers the effectiveness of multiple sources and achieves excellent improvement compared to using a single source. The results in Table 2 show that our model outperforms all baseline methods in both the None source domain and multiple source domain settings, demonstrating that the proposed approach can be beneficial in transferring knowledge learned from the source domain.

**Effect of data size.** To investigate the influence of the size of the target domain data on performance, we conduct experiments using varying number of target training data (ranging from 10 to 50 samples) based on the settings that yielded the best results. From Table 2, we notice that most of the models achieve performance improvement as the amount of data increases, indicating the essential role of labeled data. Furthermore, it can be seen that our proposed method (CP-NER) consistently outperforms the baselines (LANER and LIGHTNER) across all data sizes, further demonstrating the effectiveness of the approach in incorporating information from multiple source domains through collaborative domain-prefix tuning for cross-domain NER.

### 4.4 Analysis

**Analysis of similarity-based selector.** We visualize the entity and prefix similarity of the dual-query domain selector to analyze the correlation between source and target domains. From Figure 3, we find that the similarities in the left column are highly correlated with the entity semantic category. For example, the Music domain is most related to the target



Target	Sentence	Prediction	Reference Source
Mit-Movie	was vivien leigh in a <b>musical</b> that was rated r and received nine stars and above	None: <genre> Null Multi-Source: <genre> musical	<b>Literature:</b> The following year they collaborated on a <b>musical</b> film version of The Little Prince , based on the classic children 's tale by Antoine de Saint-Exupry.
Science	These include the TRUE Blue Crew, Antipodean Resistance, the Australian Defence League, National Action ( <b>Australia</b> ), the Q Society, Reclaim Australia and the Lads Society (formerly United Patriots Front).	None: <organization> Null Multi-Source: <organization> Australiza	<b>Music:</b> In addition to relentless touring in the U.S. and Canada, PUSA made multiple tours of Europe, <b>Australia</b> , New Zealand and Japan."
AI	Other films between 2016 to 2020 that captured with IMAX camera's were Zack Snyder's Batman v Superman: Patty Jenkins' Wonder Woman 1984, Cary Joji Fukunaga's No Time to Die and <b>Joseph</b> Kosinski's Top Gun: Maverick.	None: <researcher> Joseph Multi-Source: <person> Joseph	<b>Literature:</b> Her stage credits include Norman Mailer's The Deer Park, Israel Horowitz's The Indian Wants the Bronx, Neil Simon's The Good Doctor and <b>Joseph</b> Papp's 1974 Richard III at the Lincoln Center.

Table 3: Cases analysis between None source domain transfer and transfer with multiple source domains.




Figure 3: Similarity visualization of the source and target domains.

Movie domain, and they have the same entity category “song” and similar entity categories “genre” and “musicgenre”. Furthermore, the entity similarity and the prefix similarity have similar trends, where the source domains with the highest similarity values are roughly the same. These observations reveal that the dual-query selector effectively identifies the most valuable source domains for the target domain.

**Case analysis.** As shown in Table 3, the models that are transferred from multiple source domains benefit from the source domains and perform more robustly. For example, some entities have already appeared in the source domains and can be easily identified in the target domain. Moreover, the same entities from source domains, such as “Australia” and “Joseph”, provide direct support for the target domain, illustrating the effectiveness of our collaborative domain-prefix tuning.

**Ablation study.** As Figure. 4 illustrates, we conduct an ablation study to evaluate the impact of various model components and source domains on domain transfer performance. Specifically, we systematically eliminate the entity similarity, prefix similarity, warm-up, and domain-related options from the model, respectively. Our results show that the performance of the models decreases significantly, particularly when the entity similarity component is removed. This highlights the crucial role played by our dual-query domain selector in facilitating effective domain transfer. Additionally, we investigate




Figure 4: Ablation Study on the setting with multiple source domains. The dashed line denotes the results of our method in Table 2.

the effect of different source domains on the domain transfer performance. Our findings indicate that each source domain plays a unique role in the domain transfer process and exhibits varying levels of importance for different target domains, verifying the feasibility of leveraging knowledge from multiple domains to improve the target domain performance.

## 5 Conclusion and Future Work

In this paper, we target cross-domain NER and propose a lightweight approach dubbed CP-NER. The proposed approach uses collaborative domain-prefix tuning to better utilize knowledge from multiple domains, thus, leading to better performance for on-the-fly cross-domain NER. Through comprehensive experiments on benchmark datasets, we have shown that CP-NER achieves better results than a series of state-of-the-art cross-domain learning methods on NER. Furthermore, from the perspective of the application, our CP-NER can realize one model for all domains (Keeping the base LM frozen), which is more in line with the needs of industrial applications. In the future, we plan to extend the proposed approach to multilingual NER and other information extraction tasks. Besides, it is promising to build one unified large-scale generation model with pluggable and programmable modules to address the generalization issues.

## Acknowledgments

We would like to express gratitude to the anonymous reviewers for their kind comments. This work was supported by the National Natural Science Foundation of China (No.62206246 and U19B2027), Zhejiang Provincial Natural Science Foundation of China (No. LGG22F030011), Ningbo Natural Science Foundation (2021J190), and Yongjiang Talent Introduction Programme (2021A-156-G), CAAI-Huawei MindSpore Open Fund, and NUS-NCS Joint Laboratory (A-0008542-00-00).

## Contribution Statement

Xiang Chen, Lei Li and Shuofei Qiao make an equal contribution and share co-first authorship. Ningyu Zhang and Huajun Chen are both the corresponding authors.

## References

- [Akbik *et al.*, 2018] Alan Akbik, Duncan Blythe, and Roland Vollgraf. Contextual string embeddings for sequence labeling. In *Proceedings of COLING*, 2018.
- [Chen *et al.*, 2022a] Xiang Chen, Lei Li, Shumin Deng, Chuanqi Tan, Changliang Xu, Fei Huang, Luo Si, Huajun Chen, and Ningyu Zhang. Lightner: A lightweight tuning paradigm for low-resource NER via pluggable prompting. In *Proceedings of COLING 2022*, pages 2374–2387. International Committee on Computational Linguistics, 2022.
- [Chen *et al.*, 2022b] Xiang Chen, Ningyu Zhang, Xin Xie, Shumin Deng, Yunzhi Yao, Chuanqi Tan, Fei Huang, Luo Si, and Huajun Chen. Knowprompt: Knowledge-aware prompt-tuning with synergistic optimization for relation extraction. In *WWW '22: The ACM Web Conference 2022, Virtual Event, Lyon, France, April 25 - 29, 2022*, pages 2778–2788. ACM, 2022.
- [Devlin *et al.*, 2019] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of ACL 2019*, pages 4171–4186, 2019.
- [Ding *et al.*, 2022] Ning Ding, Yujia Qin, Guang Yang, Fuchao Wei, Zonghan Yang, Yusheng Su, Shengding Hu, Yulin Chen, Chi-Min Chan, Weize Chen, Jing Yi, Weilin Zhao, Xiaozhi Wang, Zhiyuan Liu, Hai-Tao Zheng, Jianfei Chen, Yang Liu, Jie Tang, Juanzi Li, and Maosong Sun. Parameter-efficient fine-tuning of large-scale pre-trained language models. *Nature Machine Intelligence*, abs/2203.06904, 2022.
- [Houlsby *et al.*, 2019] Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for NLP. In *ICML*, pages 2790–2799, 2019.
- [Hu *et al.*, 2022] Jinpeng Hu, He Zhao, Dan Guo, Xiang Wan, and Tsung-Hui Chang. A label-aware autoregressive framework for cross-domain NER. In *Findings of NAACL 2022*, 2022.
- [Jia and Zhang, 2020] Chen Jia and Yue Zhang. Multi-cell compositional lstm for ner domain adaptation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5906–5917, 2020.
- [Jia *et al.*, 2019a] Chen Jia, Xiaobo Liang, and Yue Zhang. Cross-domain ner using cross-domain language modeling. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2464–2474, 2019.
- [Jia *et al.*, 2019b] Chen Jia, Liang Xiao, and Yue Zhang. Cross-domain NER using cross-domain language modeling. In *Proceedings of ACL 2019*, 2019.
- [Kim *et al.*, 2015] Young-Bum Kim, Karl Stratos, Ruhi Sarikaya, and Minwoo Jeong. New transfer learning techniques for disparate label sets. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 473–482, 2015.
- [Lee *et al.*, 2018] Ji Young Lee, Franck Dernoncourt, and Peter Szolovits. Transfer learning for named-entity recognition with neural networks. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, 2018.
- [Lewis *et al.*, 2020] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, 2020.
- [Li and Liang, 2021a] Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. In *ACL*, 2021.
- [Li and Liang, 2021b] Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. In *Proceedings of ACL/IJCNLP 2021*, 2021.
- [Lin and Lu, 2018] Bill Yuchen Lin and Wei Lu. Neural adaptation layers for cross-domain named entity recognition. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2012–2022, 2018.
- [Liu *et al.*, 2013a] Jingjing Liu, Panupong Pasupat, Scott Cyphers, and James R. Glass. Asgard: A portable architecture for multilingual dialogue systems. In *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2013, Vancouver, BC, Canada, May 26-31, 2013*, pages 8386–8390. IEEE, 2013.
- [Liu *et al.*, 2013b] Jingjing Liu, Panupong Pasupat, Yining Wang, Scott Cyphers, and James R. Glass. Query understanding enhanced by hierarchical parsing structures. In *2013 IEEE Workshop on Automatic Speech Recognition and Understanding, Olomouc, Czech Republic, December 8-12, 2013*, pages 72–77. IEEE, 2013.
- [Liu *et al.*, 2019] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy,



- Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv:1907.11692*, 2019.
- [Liu *et al.*, 2020a] Zihan Liu, Genta Indra Winata, and Pascale Fung. Zero-resource cross-domain named entity recognition. In *Proceedings of the 5th Workshop on Representation Learning for NLP*, pages 1–6, Online, July 2020. Association for Computational Linguistics.
- [Liu *et al.*, 2020b] Zihan Liu, Genta Indra Winata, Peng Xu, and Pascale Fung. Coach: A coarse-to-fine approach for cross-domain slot filling. In *Proceedings of ACL 2020*, pages 19–25. Association for Computational Linguistics, 2020.
- [Liu *et al.*, 2021a] Zihan Liu, Feijun Jiang, Yuxiang Hu, Chen Shi, and Pascale Fung. NER-BERT: A pre-trained model for low-resource entity tagging. *arXiv:2112.00405*, 2021.
- [Liu *et al.*, 2021b] Zihan Liu, Yan Xu, Tiezheng Yu, Wenliang Dai, Ziwei Ji, Samuel Cahyawijaya, Andrea Madotto, and Pascale Fung. Crossner: Evaluating cross-domain named entity recognition. In *AAAI*, pages 13452–13460. AAAI Press, 2021.
- [Newman *et al.*, 2022] Benjamin Newman, Prafulla Kumar Choubey, and Nazneen Rajani. P-adapters: Robustly extracting factual information from language models with diverse prompts. In *ICLR*, 2022.
- [Ponti *et al.*, 2022] Edoardo Maria Ponti, Alessandro Sordani, and Siva Reddy. Combining modular skills in multi-task learning. *arXiv:2202.13914*, 2022.
- [Raffel *et al.*, 2020] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21:140:1–140:67, 2020.
- [Sang and Meulder, 2003] Erik F. Tjong Kim Sang and Fien De Meulder. Introduction to the conll-2003 shared task: Language-independent named entity recognition. In Walter Daelemans and Miles Osborne, editors, *Proceedings of the Seventh Conference on Natural Language Learning, CoNLL 2003, Held in cooperation with HLT-NAACL 2003, Edmonton, Canada, May 31 - June 1, 2003*, pages 142–147. ACL, 2003.
- [Shin *et al.*, 2020] Taylor Shin, Yasaman Razeghi, Robert L. Logan IV, Eric Wallace, and Sameer Singh. AutoPrompt: Eliciting Knowledge from Language Models with Automatically Generated Prompts. In *EMNLP*, pages 4222–4235, 2020.
- [Su *et al.*, 2022] Yusheng Su, Xiaozhi Wang, Yujia Qin, Chi-Min Chan, Yankai Lin, Huadong Wang, Kaiyue Wen, Zhiyuan Liu, Peng Li, Juanzi Li, Lei Hou, Maosong Sun, and Jie Zhou. On transferability of prompt tuning for natural language processing. In *Proceedings of the ACL 2022*, pages 3949–3969, Seattle, United States, July 2022.
- [Sun *et al.*, 2020] Tianxiang Sun, Yunfan Shao, Xiaonan Li, Pengfei Liu, Hang Yan, Xipeng Qiu, and Xuanjing Huang. Learning sparse sharing architectures for multiple tasks. In *AAAI*, pages 8936–8943. AAAI Press, 2020.
- [Vu *et al.*, 2022] Tu Vu, Brian Lester, Noah Constant, Rami Al-Rfou’, and Daniel Cer. SPoT: Better frozen model adaptation through soft prompt transfer. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5039–5059, Dublin, Ireland, May 2022. Association for Computational Linguistics.
- [Wang *et al.*, 2018] Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding. In *EMNLP*, pages 353–355, 2018.
- [Wang *et al.*, 2019] Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. Superglue: A stickier benchmark for general-purpose language understanding systems. In *NeurIPS*, pages 3261–3275, 2019.
- [Wang *et al.*, 2020] Jing Wang, Mayank Kulkarni, and Daniel Protiuc-Pietro. Multi-domain named entity recognition with genre-aware and agnostic inference. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8476–8488, 2020.
- [Wang *et al.*, 2021] Ruize Wang, Duyu Tang, Nan Duan, Zhongyu Wei, Xuanjing Huang, Jianshu Ji, Guihong Cao, Daxin Jiang, and Ming Zhou. K-Adapter: Infusing Knowledge into Pre-Trained Models with Adapters. In *Findings of ACL-IJCNLP*, pages 1405–1418, 2021.
- [Xue *et al.*, 2021] Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. mt5: A massively multi-lingual pre-trained text-to-text transformer. In *Proceedings of the NAACL-HLT 2021, Online, June 6-11, 2021*, pages 483–498. Association for Computational Linguistics, 2021.
- [Yan *et al.*, 2021] Hang Yan, Tao Gui, Junqi Dai, Qipeng Guo, Zheng Zhang, and Xipeng Qiu. A unified generative framework for various NER subtasks. In *Proceedings of ACL/IJCNLP 2021*, 2021.
- [Yang *et al.*, 2018] Jie Yang, Shuailong Liang, and Yue Zhang. Design challenges and misconceptions in neural sequence labeling. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3879–3889, 2018.
- [Zhang *et al.*, 2022] Ningyu Zhang, Xin Xu, Liankuan Tao, Haiyang Yu, Hongbin Ye, Shuofei Qiao, Xin Xie, Xiang Chen, Zhoubo Li, and Lei Li. Deepke: A deep learning based knowledge extraction toolkit for knowledge base population. In *Proceedings of the The 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022 - System Demonstrations, Abu Dhabi, UAE, December 7-11, 2022*, pages 98–108. Association for Computational Linguistics, 2022.
- [Zheng *et al.*, 2022] Junhao Zheng, Haibin Chen, and Qianli Ma. Cross-domain named entity recognition via graph matching. In *Findings of ACL 2022*, 2022.