

RoPDA: Robust Prompt-Based Data Augmentation for Low-Resource Named Entity Recognition

Sihan Song^{1,2}, Furao Shen^{1,3*}, Jian Zhao^{1,4}

¹National Key Laboratory for Novel Software Technology, Nanjing University, China

²Department of Computer Science and Technology, Nanjing University, China

³School of Artificial Intelligence, Nanjing University, China

⁴School of Electronic Science and Engineering, Nanjing University, China

whalesihan@smail.nju.edu.cn, frshen@nju.edu.cn, jianzhao@nju.edu.cn

Abstract

Data augmentation has been widely used in low-resource NER tasks to tackle the problem of data sparsity. However, previous data augmentation methods have the disadvantages of disrupted syntactic structures, token-label mismatch, and requirement for external knowledge or manual effort. To address these issues, we propose **Robust Prompt-based Data Augmentation (RoPDA)** for low-resource NER. Based on pre-trained language models (PLMs) with continuous prompt, RoPDA performs entity augmentation and context augmentation through five fundamental augmentation operations to generate label-flipping and label-preserving examples. To optimize the utilization of the augmented samples, we present two techniques: self-consistency filtering and mixup. The former effectively eliminates low-quality samples with a bidirectional mask, while the latter prevents performance degradation arising from the direct utilization of label-flipping samples. Extensive experiments on three popular benchmarks from different domains demonstrate that RoPDA significantly improves upon strong baselines, and also outperforms state-of-the-art semi-supervised learning methods when unlabeled data is included.

Introduction

Named Entity Recognition (NER) is a fundamental NLP task which is dedicated to identifying predefined named entities (*e.g.*, persons, organizations and locations) from texts. With the rapid development of deep learning in recent years, fine-tuning pre-trained language models (PLMs) such as BERT (Devlin et al. 2019) for NER tasks has yielded promising results (Zhong and Chen 2021; Chen, Neves, and Solorio 2022). However, fine-tuning PLMs still necessitates a substantial amount of human annotations. NER models, in particular, require each token to be labeled in a sequence, which is a laborious and time-consuming task in the real world. Consequently, NER frequently encounters the challenge of data sparsity, making low-resource NER a pressing priority.

In order to mitigate the issue of data sparsity, various data augmentation methods for low-resource NER have been proposed, such as traditional word-level manipulation (Dai and

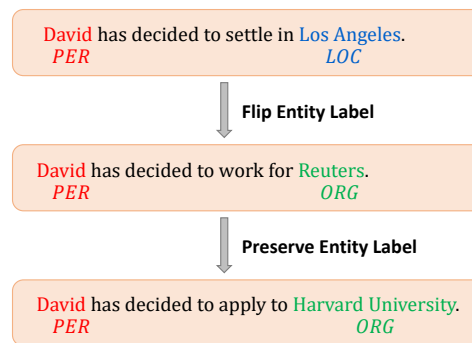


Figure 1: Label-flipping and label-preserving operations.

Adel 2020; Zeng et al. 2020) and more recent PLM-based methods, and the latter have received a lot of attention and have yielded promising results (Zhou et al. 2022b; Wang et al. 2022b). Zhou et al. (2022b) leverage a masked language model to randomly mask entities in the sentence and regenerate them conditioned on labels, thereby enhancing entity diversity. Wang et al. (2022b) employ Soft Prompt for seq2seq PLMs and propose a dual-view augmentation approach to generate sentences conditioned on labels and keywords. However, their approaches have the limitation of only boosting entity diversity but not context diversity, or necessitating external knowledge (*i.e.*, external corpus).

To address these issues, we propose **Robust Prompt-based Data Augmentation (RoPDA)** for low-resource NER. With well-trained continuous prompt (Li and Liang 2021), our model is capable of automatically generating training samples for low-resource NER tasks, eliminating the need for external knowledge or human efforts, which is different from Wang et al. (2022b). To enhance model generalization, we propose five fundamental data augmentation operations. Among these, one operation focuses on context augmentation, aiming to increase context diversity and enrich our training data. The remaining four operations concentrate on entity augmentation, designed to generate adversarial examples and promote entity diversity, including label-flipping and label-preserving operations.

Inspired by Zhou et al. (2022a), a label-flipping operation means regenerating an entity in the sentence into a different

*Corresponding author

type of entity. As shown in Figure 1, after a label-flipping operation, the entity type sequences of the augmented sentence and the original sentence differ only in one entity. Such an augmented sentence can serve as an adversarial example (Goodfellow, Shlens, and Szegedy 2014) for this modified entity type, enhancing the NER model’s capability in distinguishing these two entity types before and after regeneration. Conversely, A label-preserving operation involves regenerating an entity as a new entity of the same type, thereby prompting entity diversity. We argue that NER models can benefit from both label-preserving and label-flipping operations. With these five augmentation operations, we can generate smooth and diverse sentences without relying on external knowledge. To further optimize the utilization of augmented samples, we propose self-consistency (Wang et al. 2022a) filtering and mixup. The former acquires the capability to effectively filter out low-quality samples by fine-tuning the PLM with a bidirectional mask. The latter employs linear interpolation for adversarial examples, resulting in a smoother label distribution that prevents the performance degradation caused by their direct utilization. Through these five augmentation operations followed by self-consistency filtering and mixup, we can generate high-quality augmented samples.

To summarize, our contributions are as follows:

- We propose a robust prompt-based data augmentation method RoPDA for low-resource NER, which contains five augmentation operations. By simultaneously augmenting entities and contexts, RoPDA can generate augmented examples with high diversity.
- We propose self-consistency filtering to improve the quality of augmented samples through bidirectional masking.
- We utilize the mixup technique to interpolate adversarial examples with corresponding original examples, effectively maximizing the utility of adversarial examples.
- Experiments on three popular benchmarks show the significant performance gain of RoPDA over current state-of-the-art baselines.

Related Work

Word-level manipulation is a prevalent data augmentation method that manipulates words in the original text to generate synthetic text utilizing predefined rules (Dai and Adel 2020; Wei and Zou 2019; Kobayashi 2018). Dai and Adel (2020) generate new examples through token substitutions, including synonym replacement and mention replacement. But these methods run the risk of destroying the sentence structure and making labels inconsistent with modified tokens. Recently, leveraging the powerful generative capability and rich knowledge of PLMs to generate augmented data has been explored (Kumar, Choudhary, and Cho 2020; Ding et al. 2020). Zhou et al. (2022b) perform entity replacement on corrupted sentences based on the masked language model. Wang et al. (2022b) and Anaby-Tavor et al. (2020) leverage seq2seq PLMs to generate synthetic data conditioned on entity types.

Adversarial augmentation (Volpi et al. 2018) is commonly used to improve model robustness, but recent research has discovered that it can also improve model generalization (Xie et al. 2020). Reich et al. (2022) adopt expert-guided heuristics for adversarial augmentation, focusing on modifying entities and contexts in accordance with predefined rules. However, their approach lacks diversity in entities and relies on expert knowledge. Lin et al. (2021) create entity attacks by replacing the target entity with entities of the same type from an external knowledge base. On the other hand, our approach, which is based on PLMs, can produce high-quality adversarial examples without the need for external knowledge or human effort.

Method

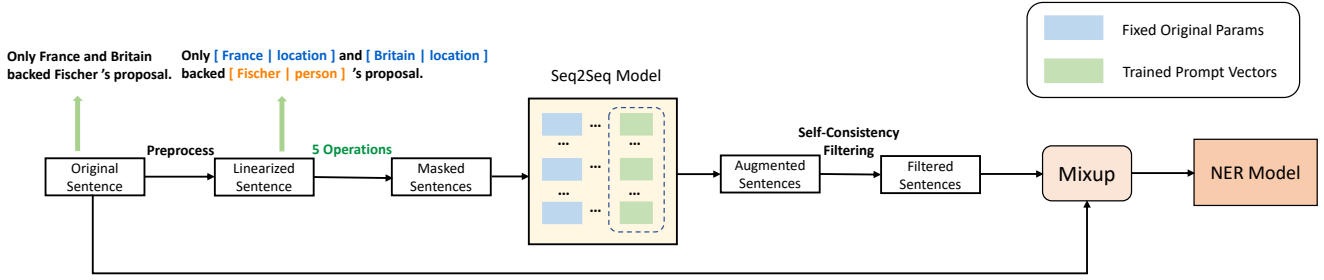
The overview of RoPDA is shown in Figure 2. Firstly, We begin by preprocessing the original data into linearized sentences with entities constrained by their types. Secondly, we add prompt vectors to the seq2seq PLM and fine-tune it with the linearized few-shot data. Thirdly, linearized sentences undergo strategic masking using five fundamental augmentation operations. Fourthly, we regenerate the masked sentences using the fine-tuned PLM to produce augmented sentences. Subsequently, self-consistency filtering is applied to filter out noisy low-quality samples. Lastly, to better utilize generated adversarial examples, a mixup is employed to interpolate the adversarial examples and the original data during the NER model training.

Data Preprocessing

Similar to Paolini et al. (2021) and Zhou et al. (2022b), we adopt a linearization strategy to convert a sentence and its corresponding tags into a linearized sequence. As shown in Figure 2(a), given a sentence $X = [x_1, x_2, x_3, \dots, x_L]$, for each entity $e_{ij} = x_i \dots x_j$ with type l_k (e.g., $l_k = \text{“PER”}$), we convert it into “[$x_i \dots x_j \mid O(l_k)$]”, where $O(l_k)$ is the natural language form of label l_k (e.g., $O(l_k) = \text{“person”}$). Before being sent to PLMs, labeled sentences must be processed using this template. This data preprocessing enables the PLM to explicitly take label information into account when generating tokens, thereby constraining entity and entity label consistency. Furthermore, such a template is reversible, enabling us to recover the sentence and its labels from the linearized sequence.

Prompt-based PLM

Fine-tuning is the prevalent way to adapt large PLMs to downstream tasks. In low-resource scenarios, however, the available training data is inadequate relative to the model size, and updating all parameters of the model can readily cause overfitting, leading to a decline in generalization. By prepending instructions to the task input and directly generating the task output from PLMs, prompt can effectively harness the potential of PLMs, especially in low-resource settings (Brown et al. 2020; Liu et al. 2021, 2022). Following Wang et al. (2022b), we add a sequence of continuous trainable vectors, which is also called Soft Prompt, as shown in Figure 2(a), to each layer of the PLM. During training,



(a) Overview

Op1- Preserving:	Only [France location] and [Britain location] <MASK> [<MASK> person] <MASK> proposal.
Op2- Changing:	Only [France location] and [Britain location] <MASK> [<MASK> organization] <MASK>.
Op3- Adding:	Only [France location] and [Britain location] backed [Fischer person] <MASK> [<MASK> person] <MASK> 's proposal.
Op4- Erasing:	Only [France location] and [Britain location] <MASK> proposal.
Op5- Context:	Only [France location] and [Britain location] <MASK> [Fischer person] 's proposal.

(b) Masked Sentences

Op1- Preserving:	Only [France location] and [Britain location] agree to [Merkel person] 's proposal.
Op2- Changing:	Only [France location] and [Britain location] think [European Union organization] 's decision is unreasonable.
Op3- Adding:	Only [France location] and [Britain location] backed [Fischer person] and [John Smith person] 's farm proposal.
Op4- Erasing:	Only [France location] and [Britain location] agreed to consider their proposal.
Op5- Context:	Only [France location] and [Britain location] are not optimistic about [Fischer person] 's proposal.

(c) Augmented Sentences

Figure 2: Overview of our proposed RoPDA. Figure 2(b) represents *Masked Sentences* generated after performing 5 operations in Figure 2(a). Figure 2(c) represents *Augmented Sentences* in Figure 2(a). < MASK > means some words in the linearized sentence being masked.

we only update the parameters of the prompt vectors and fix all PLM parameters. Note that in this paper, we choose the T5 (Raffel et al. 2020) model as our backbone sequence-to-sequence PLM. Different from Wang et al. (2022b), we do not pre-train prompt vectors with additional corpus texts but only utilize a small number of labeled data for training. Although our method does not leverage external knowledge, we still achieve significant improvement compared to their method.

Data Augmentation

Tokens in a sentence can be divided into entity segments and context segments depending on whether or not they are in an entity. Formally, a sentence X is divided into $C_1 E_1 \dots C_n E_n C_{n+1}$, with C_i representing the context segment and E_i representing the entity segment. To enhance data diversity, we propose five fundamental data augmentation operations, one of which is dedicated to context augmentation, while the other four focus on entity augmentation, involving label-flipping and label-preserving operations. The sentence is preprocessed using the linearization strategy before each operation. Then we apply the five operations to the linearized sentence by masking the according part of it. PLMs are then employed to generate the masked parts in order to obtain the synthetic sentences. Figure 2(b) shows examples of each operation.

Op1: Augmenting the Entity-Related Span Choose an entity segment at random from the sentence, mask the entity segment and a portion of its surrounding context segments.

Op2: Changing the Entity Type Choose an entity segment

Strategy	Operations
SA	Op1*M + Op5*N
ELC	Op2*K + Op1*(M-K) + Op5*N
EA	Op3*K + Op1*(M-K) + Op5*N
ER	(Op3 + Op4)*K + Op1*(M-K) + Op5*N

Table 1: Data augmentation strategies. * means that the operation is repeated, followed by the number of repetitions. + means that two operations are performed sequentially. K is the number of label-flipping operations performed on each sentence. M and N represent the times of entity and context augmentation for each sentence.

at random from the sentence and replace its type with a new entity type l_{new} . Then mask the entity segment and a portion of its surrounding context segments.

Op3: Adding an Entity Choose an entity segment E_i from the sentence randomly, and add a new entity segment of type l_{new} and corresponding context segments after E_i . As a result, the processed sentence would be: $\dots C_i E_i < MASK > [< MASK > | O(l_{new})] < MASK > C_{i+1} \dots$.

Op4: Erasing an Entity Choose an entity segment E_i at random, then mask it together with a portion of its surrounding context C_i and C_{i+1} , which is equivalent to removing the entity and associated context from the sentence. As a result, the processed sentence would be: $\dots C_{i-1} E_{i-1} < MASK > E_{i+1} \dots$.

Op5: Augmenting Contextual Spans Choose a context

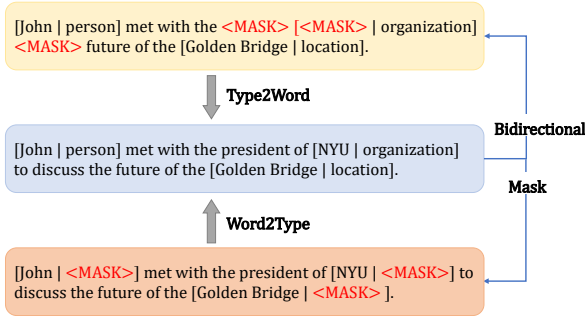


Figure 3: Bidirectional Masking

segment from the sentence at random, and then mask a portion of the context.

Op2 to Op4 are the label-flipping operations, since they essentially alter the entity type sequence of the original sentence, thus generating an adversarial example for the altered entity type. There are two ways to choose the new entity type l_{new} for label-flipping operations: one is to choose at random from the label set, and the other is to choose based on the entity type similarity. More details can be seen in the Supplementary Material. Op1, on the other hand, is a label-preserving operation that increases entity diversity by regenerating entities into new entities of the same type without altering the entity sequence. Op5 serves to enhance the diversity of the context, a factor that we consider also crucial for increasing the overall training set diversity.

Both label-flipping and label-preserving operations can improve the performance of NER models and can be effectively utilized in combination. Therefore, we propose four data augmentation strategies based on these fundamental operations, as shown in Table 1. The first, Standard Augmentation (SA), is the label-preserving strategy, and the latter three, Entity Label Change (ELC), Entity Adding (EA) and Entity Replacing (ER), are the label-flipping strategies due to the use of label-flipping operations. As previously mentioned, the data generated by label-flipping strategies can be viewed as adversarial examples.

Self-Consistency Data Filtering

The samples generated by the proposed augmentation strategies may exhibit the entity and type inconsistency issue, especially for label-flipping strategies. To further improve the quality of augmented samples, we propose a novel filtering strategy based on self-consistency that employs a bidirectional mask to fine-tune the T5 model. As shown in Figure 3, the bidirectional mask includes Type2Word and Word2Type, where Type2Word refers to masking the words(entities and contexts) in the linearized sentence and inferring them based on entity types, and Word2Type refers to masking entity types and inferring them based on words. After that, the fine-tuned T5 model will acquire bidirectional inference capabilities between entities/contexts and types and the ability to recognize and verify consistent samples. Specifically, we first generate diverse augmented samples using the type-to-

entity/context inference capability. We then use Word2Type to regenerate each entity type in the augmented samples with the assistance of the entity/context-to-type inference capability. We only keep samples whose regenerated entity types are consistent with the original. We believe that such samples are self-consistent for the model and have a higher degree of confidence.

Mixup

Using the four proposed strategies, we can generate a large number of adversarial samples. However, directly training the NER model with adversarial examples may cause overfitting on adversarial features (Lee, Lee, and Yoon 2020). Mixup (Zhang et al. 2017), as a regularization technique, can not only improve the model’s generalization performance but also its robustness against adversarial attacks (Si et al. 2021; Lee, Lee, and Yoon 2020). To this end, we leverage mixup to prevent the model from overfitting the adversarial samples and help improve generalization ability by linearly interpolating the adversarial samples and the original data.

Given a pair of data points (x, y) and (x', y') , with x denoting a data point and y being its label in a one-hot representation, mixup creates a new data point by linearly interpolating the data points and their labels:

$$\hat{x} = \lambda x + (1 - \lambda)x'$$

$$\hat{y} = \lambda y + (1 - \lambda)y'$$

where the mixing parameter λ is sampled from a Beta distribution, $\lambda \sim \text{Beta}(\alpha, \beta)$.

In this work, we linearly interpolate the label-flipping sample x_f with its corresponding original example x_o in the hidden space. At the m -th layer, the hidden representations for each token in x_f , denoted as h_f^m , are interpolated with h_o^m , the hidden representations for each token in the original example x_o , by a ratio λ :

$$\hat{h}^m = \lambda h_f^m + (1 - \lambda)h_o^m$$

After that, \hat{h}^m is fed to the $(m + 1)$ -th layer. In the meantime, the corresponding labels of the two samples are linearly mixed with the same ratio.

Experiments

Experimental Setup

Datasets We conduct experiments on three datasets. (a) CoNLL03 (Tjong Kim Sang and De Meulder 2003) is a collection of news wire articles from the Reuters Corpus, containing 4 entity types. (b) MIT Restaurant (Liu et al. 2013) is a collection of user utterances in the restaurant domain with 8 entity types. (c) MIT Movie (Liu et al. 2013) consists of user utterances in the movie domain with 12 entity types. These three datasets are from different domains and have varying numbers of entity types, allowing for a comprehensive evaluation of our method. We create four low-resource settings of shot-5/10/20/50 for each dataset. In the shot-K setting, we sample K samples from the Train set for each entity type as the training set and add the remaining to the unlabeled set.

Dataset Shot	CoNLL03				MIT Restaurant				MIT Movie			
	5	10	20	50	5	10	20	50	5	10	20	50
Baseline	65.9	73.6	78.5	82.8	50.3	59.2	66.1	70.5	68.0	70.8	76.0	81.2
SDANER	68.7	74.3	79.4	83.4	51.2	59.8	66.2	70.7	72.8	75.6	78.1	81.8
MELM	67.1	74.6	78.1	82.9	50.7	60.1	66.2	70.4	69.2	71.3	76.5	81.5
PromDA	71.1	78.2	81.0	84.2	51.8	60.3	66.7	70.7	75.2	76.4	78.5	82.4
RoPDA	74.1	79.8	81.9	85.1	55.3	62.2	67.6	71.4	75.5	77.2	79.8	82.6
MetaST*	70.5	76.4	79.8	83.6	55.2	62.4	68.6	72.5	71.7	77.7	79.0	81.9
RoPDA+self-training*	72.1	80.2	84.3	86.2	56.2	62.7	69.1	72.8	72.3	78.0	80.9	83.3
RoPDA*	75.0	81.9	85.2	86.4	56.6	64.1	69.5	73.0	75.9	78.3	81.6	83.7

Table 2: F1-score on benchmark datasets. * denotes the use of unlabeled data.

Baselines We use the model trained only on the original data as the baseline. We also adopt several other state-of-the-art methods for comparison: (1) **SDANER** (Dai and Adel 2020) explores different token replacement techniques for NER tasks. (2) **MELM** (Zhou et al. 2022b) performs entity replacement on corrupted sentences based on the masked language model. (3) **PromDA** (Wang et al. 2022b) employs soft prompt for seq2seq PLMs and proposes a dual-view data augmentation method that generates data conditioned on labels or keywords. (4) **MetaST** (Wang et al. 2021) is a state-of-the-art semi-supervised method, which makes use of unlabeled data through self-training and mitigates errors from noisy pseudo-labels through adaptive re-weighting.

Leveraging Unlabeled Data Unlabeled data contains a wealth of information that, if harnessed effectively, can enhance model performance significantly. A common solution for utilizing unlabeled data is self-training, which involves pseudo-annotating unlabeled data in the i -th iteration with the model from the $(i-1)$ -th iteration. To maximize the utilization of knowledge in unlabeled data, we propose a method called RoPDA* for generating augmented samples for unlabeled data. After pseudo-annotating unlabeled data, we only retain those with high-confidence pseudo-labels and use PLMs to synthesize augmented data for them. Finally, we train the NER model using original data, adversarial examples, pseudo-labeled data, and augmented data for unlabeled data in an iterative fashion.

Main Results

Without unlabeled data, RoPDA achieves significant performance gains and outperforms the state-of-the-art baselines with a large margin at every shot. As shown in Table 2, RoPDA consistently outperforms SDANER, MELM, and PromDA, and also outperforms MetaST in most cases, which uses additional unlabeled data. RoPDA yields an improvement of 2.3-8.3% on CoNLL03, while achieving an improvement of 0.6-4.8% on MIT Restaurant and 1.4-7.5% on MIT Movie.

With unlabeled data, our approach achieves further improvements as shown in RoPDA* in Table 2. RoPDA* far outperforms MetaST, the state-of-the-art semi-supervised method, across all benchmarks. In comparison to RoPDA, RoPDA* achieves another improvement by 1.9%, 1.7%, and 1.1% on three benchmarks, respectively. To better understand the role of unlabeled data, we also conduct ex-

Methods	CoNLL03	Resta	Movie
RoPDA	81.9	67.6	79.8
w/o. prompt	80.5	67.1	79.0
w/o. filter	81.5	67.1	79.5
w/o. mixup	81.3	67.2	79.4

Table 3: Ablation study for the combination of four types under shot-20.

Model	Dataset	SA	ELC	ER	EA	All
w/o filter	CoNLL03	0.2	0.8	0.4	0.6	0.4
	Resta	0.3	1.0	0.3	1.5	0.5
	Movie	0.2	0.7	0.6	0.4	0.3
	Avg	0.2	0.8	0.4	0.8	0.4
w/o mixup	CoNLL03	0.2	1.3	0.5	0.9	0.6
	Resta	0.1	1.3	0.7	0.8	0.4
	Movie	0.5	0.9	1.2	1.0	0.4
	Avg	0.3	1.2	0.8	0.9	0.5

Table 4: Ablation study on every single type. Values in the table represent the *drop of F1*. Avg represents the average of three benchmarks. All represents the combination of the four types.

periments on RoPDA+self-training. As shown in Table 2, RoPDA* consistently outperforms RoPDA+self-training, which shows that data augmentation on unlabeled data can further improve model performance, reflecting the superiority of our method.

Ablation Study

As shown in Tables 3 and 4, we conduct ablation studies to quantify the contribution of various components, i.e., soft prompt/self-consistency filtering/mixup. Additionally, we analyze the impact of filtering and mixup on each individual strategy.

Soft Prompt To examine the impact of soft prompt, we compare the soft prompt model with the standard T5 model, which undergoes fine-tuning using a learning rate of $1e-4$. As depicted in Table 3, eliminating the soft prompt leads to significant performance decreases across all three benchmarks, highlighting the efficacy of soft prompt in low-

Dataset	SA	ELC	ER	EA
CoNLL03	78.2	63.1	57.5	57.9
Resta	72.8	62.3	55.5	48.8
Movie	72.1	61.7	57.7	54.2
Avg	74.4	62.4	56.9	53.6

Table 5: The proportion(%) retained for each type after self-consistency filtering.

SA	ELC	EA	ER	CoNLL03	Resta	Movie	Avg
✓	✓	✓		81.0	66.8	79.5	75.8
✓	✓		✓	81.3	67.0	79.4	75.9
✓		✓	✓	81.6	67.9	79.6	76.4
	✓	✓	✓	81.7	67.4	79.6	76.2
✓	✓	✓	✓	81.9	67.6	79.8	76.4

Table 6: Experiment results for data combinations under shot-20.

resource scenarios.

Self-Consistency Filtering As shown in Table 3, after removing self-consistency filtering, there is a moderate decline in model performance across all benchmarks. Subsequently, we proceed to examine the impact of self-consistency filtering on each individual strategy. As shown in Table 4, removing self-consistency filtering brings greater performance degradation to label-flipping strategies with the biggest decline in ELC having a value of 0.8, but minor degradation to the label-preserving strategy. We believe that this is because label-flipping strategies introduce more semantic and structural alterations to the original sentence, thereby introducing more noise and necessitating a higher degree of filtering. This hypothesis can be supported by Table 5, where it is evident that self-consistency filtering retains the largest proportion of data for SA, while effectively filtering out a much greater amount of low-quality data for label-flipping strategies, specifically ELC.

Mixup As shown in Table 3, removing mixup leads to performance degradation across all benchmarks with an average decline of 0.5. Furthermore, an analysis of mixup’s impact on each individual strategy, as shown in Table 4, reveals that mixup also significantly affects label-flipping strategies to a greater extent than the label-preserving strategy. This phenomenon can be attributed to the fact that interpolating adversarial examples with the original data prevents overfitting adversarial features and improves generalization. Conversely, mixup on label-preserving data does not yield the same benefits. Notably, self-consistency filtering and mixup have analogous impacts on these strategies, as they both essentially enhance label-flipping samples.

Analysis

Different Combinations of Strategies As shown in Table 6, the performance of the model decreases when one strategy is removed from the combination of all, except ELC. This implies that the contribution of ELC is comparatively

Dataset	F-Mixup	P-Mixup	Both-Mixup	Joint-Mixup
CoNLL03	81.9	81.2	80.8	81.5
Resta	67.6	67.1	67.2	66.7

Table 7: Different mixup methods under shot-20.

Dataset	Random	F+S	P+S	F+D	P+D
CoNLL03	81.9	81.4	80.8	80.6	81.2
MIT	67.6	66.9	66.3	66.0	66.4

Table 8: Ablation study for label flipping schemes. F is Fixed, P is Probability, S means that the more similar between two entity types, the higher the flip probability and D means the opposite.

limited when incorporating the other three strategies. We believe this is because directly changing entity types introduce too much noise, which is inferior to the two indirect ways of changing entity types, EA and ER. In conclusion, all four strategies can improve model performance, but ELC contributes less.

Different Mixup Choices Our original mixup method (**F-Mixup**) is to mix up the label-flipping examples, also called adversarial examples, with the corresponding original data. We also carry out experiments on the following mixup methods to verify the superiority of F-Mixup: (1) **P-Mixup** mixes up the label-preserving data with the corresponding original data. (2) **Both-Mixup** mixes both label-flipping data and label-preserving data with the corresponding original data. (3) **Joint-Mixup** mixes the label-flipping data with both the corresponding label-preserving data and original data. As shown in Table 7, **F-Mixup** is undoubtedly superior to other mixups, demonstrating the importance of interpolating adversarial examples with the original data.

Label Flipping Schemes There are two ways to choose l_{new} for label-flipping operations. The first one **Random** is to choose randomly from the label set, and the other is to choose based on entity type similarity. For similarity-based methods, we propose two options: **Fixed** means fixed flipping of each entity type to the most/least similar type. **Probability** indicates that the flip probability is calculated based on the entity type similarity using the softmax function. Table 8 shows that **Random** outperforms all similarity-based schemes. We contend that this is due to the fact that random selection has the potential to maximize sentence diversity, whereas selection based on similarity will result in a much higher probability of flipping to a specific type than other types for each type, thus reducing sentence diversity.

Hyperparameter Setting for Data Augmentation There are several important parameters for data augmentation: K , M and N . K is the number of label-flipping operations performed on each sample. K is initially set to 1, and when K increases, the F1 score on CoNLL03 falls by 1.1%. M and N represent the times of entity and context augmentation for each sentence, respectively. As shown in Figure 4, when context augmentation is removed, that is, when $N=0$,

Strategy	Example
Original	[Bonds person] came out of Wednesday’s game against [New York organization] in the ninth inning after suffering a mild hamstring strain .
SA	[Matt Carpenter person] was carted out of the game for [Baltimore Orioles organization] in the ninth <u>with</u> a mild hamstring strain.
ELC	[Federer person] pulled out of Wednesday’s [European miscellaneous] final after suffering back strain.
ER	[Robinson Cano person] <u>left the field of</u> Wednesday’s game <u>in</u> [Los Angeles location] <u>with two outs</u> in the ninth after a hard groundout.
EA	[Marquez person] <u>was the</u> [Boston organization] <u>starting pitcher and left the</u> game against [New York organization] in the ninth after suffering a mild hamstring strain.
Original	find me a [nice rating] place to eat that is [not too expensive price]
SA	find me a [nice rating] place to have dinner that is [reasonably priced price]
ELC	<u>where is the</u> [best rating] place to get [chicken wings dish]
ER	find me <u>some</u> [good rating] food <u>with</u> [parking amenity] nearby
EA	find me a [nice rating] place to eat that is [not too expensive price] <u>and has</u> [free wifi amenity] please

Table 9: Augmented data from RoPDA. The top half of the table is from CoNLL03 and the bottom half is from MIT Restaurant. Bold text chunks are novel entities. Underlined text chunks are novel contexts.




Figure 4: Experiment results with different M and N on CoNLL03.

the performance declines by 1.2%, indicating the importance of enhancing the contexts of sentences. When M and N take fixed values, the performance suffers when they are too large or too small, with M=2 and N=3 providing the best results. In addition, the best performance is attained when M and N are both randomly chosen from {1, 2, 3}. In our original experiment, when performing entity augmentation, we not only mask and regenerate the entity but also regenerate a portion of the surrounding context, as we believe this will make the generated entity more in tune with its context. We run experiments on CoNLL03 to verify our hypothesis. The F1 score decreases by 0.8% when only regenerating the entity without regenerating its context, indicating the importance of regenerating its context simultaneously.

RoPDA in Data-Rich Settings Despite being designed for low-resource settings, RoPDA can still offer appreciable performance improvements in high-resource settings. As shown in Table 10, in data-rich scenarios, RoPDA improves model performance on all three benchmarks, and as the amount of data increases, the performance gain brought by RoPDA continues to decline. In Addition, we discover that the worse the baseline performance of a certain dataset, the greater the performance gain brought by RoPDA.

Dataset	1000	2000	All Data
CoNLL03	87.5(+0.4)	88.7(+0.3)	90.3(+0.3)
Resta	75.3(+1.6)	78.2(+0.6)	80.0(+0.5)
Movie	83.9(+1.2)	84.7(+0.5)	87.9(+0.4)

Table 10: Gains brought by RoPDA in data-rich settings.

Case Study

Table 9 shows some generated examples from RoPDA. RoPDA can generate high-quality new entities and increase entity diversity by leveraging the powerful inference ability and rich knowledge contained in PLMs. Taking the sentence from CoNLL03 as an example, “New York” is usually represented as a location-type entity, but in our example sentence, “New York” represents an organization type. In EA strategy, the trained language model learns the true meaning of “New York” through contextual inference and its conditioned label and thus, generates a new organization entity “Boston” which has a similar meaning to “New York”. In addition, as can be seen in Table 9, in ELC and ER strategies, the original entity can be regenerated as a new type of entity that is harmonious with the context. Moreover, RoPDA improves not only entity diversity but also context diversity significantly.

Conclusion

In this paper, we propose a robust prompt-based data augmentation method RoPDA for low-resource NER. RoPDA performs entity augmentation and context augmentation through five fundamental augmentation operations to generate label-flipping and label-preserving examples. To optimize the utilization of augmented data, we introduce two techniques: self-consistency filtering and mixup. Self-consistency filtering efficiently eliminates low-quality samples, while mixup mitigates performance degradation arising from the direct utilization of adversarial samples. Extensive experiments on three benchmark datasets showcase the effectiveness of RoPDA.

Acknowledgments

This work was supported by the National Science Foundation of China under Grant 62276127.

References

- Anaby-Tavor, A.; Carmeli, B.; Goldbraich, E.; Kantor, A.; Kour, G.; Shlomov, S.; Tepper, N.; and Zwerdling, N. 2020. Do not have enough data? Deep learning to the rescue! In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 7383–7390.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33: 1877–1901.
- Chen, S.; Neves, L.; and Solorio, T. 2022. Style Transfer as Data Augmentation: A Case Study on Named Entity Recognition. *arXiv preprint arXiv:2210.07916*.
- Dai, X.; and Adel, H. 2020. An analysis of simple data augmentation for named entity recognition. *arXiv preprint arXiv:2010.11683*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 4171–4186. Minneapolis, Minnesota: Association for Computational Linguistics.
- Ding, B.; Liu, L.; Bing, L.; Kruengkrai, C.; Nguyen, T. H.; Joty, S.; Si, L.; and Miao, C. 2020. DAGA: Data augmentation with a generation approach for low-resource tagging tasks. *arXiv preprint arXiv:2011.01549*.
- Goodfellow, I. J.; Shlens, J.; and Szegedy, C. 2014. Explaining and Harnessing Adversarial Examples. *CoRR*, abs/1412.6572.
- Kobayashi, S. 2018. Contextual augmentation: Data augmentation by words with paradigmatic relations. *arXiv preprint arXiv:1805.06201*.
- Kumar, V.; Choudhary, A.; and Cho, E. 2020. Data augmentation using pre-trained transformer models. *arXiv preprint arXiv:2003.02245*.
- Lee, S.; Lee, H.; and Yoon, S. 2020. Adversarial vertex mixup: Toward better adversarially robust generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 272–281.
- Li, X. L.; and Liang, P. 2021. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv preprint arXiv:2101.00190*.
- Lin, B. Y.; Gao, W.; Yan, J.; Moreno, R.; and Ren, X. 2021. RockNER: A Simple Method to Create Adversarial Examples for Evaluating the Robustness of Named Entity Recognition Models. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 3728–3737. Online and Punta Cana, Dominican Republic: Association for Computational Linguistics.
- Liu, J.; Pasupat, P.; Cyphers, S.; and Glass, J. 2013. Asgard: A portable architecture for multilingual dialogue systems. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, 8386–8390. IEEE.
- Liu, X.; Ji, K.; Fu, Y.; Tam, W.; Du, Z.; Yang, Z.; and Tang, J. 2022. P-Tuning: Prompt Tuning Can Be Comparable to Fine-tuning Across Scales and Tasks. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 61–68. Dublin, Ireland: Association for Computational Linguistics.
- Liu, X.; Zheng, Y.; Du, Z.; Ding, M.; Qian, Y.; Yang, Z.; and Tang, J. 2021. GPT understands, too. *arXiv preprint arXiv:2103.10385*.
- Paolini, G.; Athiwaratkun, B.; Krone, J.; Ma, J.; Achille, A.; Anubhai, R.; Santos, C. N. d.; Xiang, B.; and Soatto, S. 2021. Structured prediction as translation between augmented natural languages. *arXiv preprint arXiv:2101.05779*.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; Liu, P. J.; et al. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140): 1–67.
- Reich, A.; Chen, J.; Agrawal, A.; Zhang, Y.; and Yang, D. 2022. Leveraging Expert Guided Adversarial Augmentation For Improving Generalization in Named Entity Recognition. *arXiv preprint arXiv:2203.10693*.
- Si, C.; Zhang, Z.; Qi, F.; Liu, Z.; Wang, Y.; Liu, Q.; and Sun, M. 2021. Better robustness by more coverage: Adversarial and mixup data augmentation for robust finetuning. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, 1569–1576.
- Tjong Kim Sang, E. F.; and De Meulder, F. 2003. Introduction to the CoNLL-2003 Shared Task: Language-Independent Named Entity Recognition. In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*, 142–147.
- Volpi, R.; Namkoong, H.; Sener, O.; Duchi, J. C.; Murino, V.; and Savarese, S. 2018. Generalizing to unseen domains via adversarial data augmentation. *Advances in neural information processing systems*, 31.
- Wang, X.; Wei, J.; Schuurmans, D.; Le, Q.; Chi, E.; Narang, S.; Chowdhery, A.; and Zhou, D. 2022a. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*.
- Wang, Y.; Mukherjee, S.; Chu, H.; Tu, Y.; Wu, M.; Gao, J.; and Awadallah, A. H. 2021. Meta self-training for few-shot neural sequence labeling. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 1737–1747.
- Wang, Y.; Xu, C.; Sun, Q.; Hu, H.; Tao, C.; Geng, X.; and Jiang, D. 2022b. PromDA: Prompt-based Data Augmentation for Low-Resource NLU Tasks. *arXiv preprint arXiv:2202.12499*.
- Wei, J.; and Zou, K. 2019. Eda: Easy data augmentation techniques for boosting performance on text classification tasks. *arXiv preprint arXiv:1901.11196*.

Xie, C.; Tan, M.; Gong, B.; Wang, J.; Yuille, A. L.; and Le, Q. V. 2020. Adversarial examples improve image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 819–828.

Zeng, X.; Li, Y.; Zhai, Y.; and Zhang, Y. 2020. Counterfactual generator: A weakly-supervised method for named entity recognition. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 7270–7280.

Zhang, H.; Cisse, M.; Dauphin, Y. N.; and Lopez-Paz, D. 2017. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*.

Zhong, Z.; and Chen, D. 2021. A Frustratingly Easy Approach for Entity and Relation Extraction. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 50–61. Online: Association for Computational Linguistics.

Zhou, J.; Zheng, Y.; Tang, J.; Jian, L.; and Yang, Z. 2022a. FlipDA: Effective and Robust Data Augmentation for Few-Shot Learning. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 8646–8665. Dublin, Ireland: Association for Computational Linguistics.

Zhou, R.; Li, X.; He, R.; Bing, L.; Cambria, E.; Si, L.; and Miao, C. 2022b. MELM: Data Augmentation with Masked Entity Language Modeling for Low-Resource NER. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2251–2262.