

Inverse is Better! Fast and Accurate Prompt for Few-shot Slot Tagging

Yutai Hou*, Cheng Chen*, Xianzhen Luo, Bohan Li, Wanxiang Che†

Research Center for Social Computing and Information Retrieval,
Harbin Institute of Technology

{ythou, cchen, xzluo, bhli, car}@ir.hit.edu.cn

Abstract

Prompting methods recently achieve impressive success in few-shot learning. These methods modify input samples with prompt sentence pieces, and decode label tokens to map samples to corresponding labels. However, such a paradigm is very inefficient for the task of slot tagging. Since slot tagging samples are multiple consecutive words in a sentence, the prompting methods have to enumerate all n-grams token spans to find all the possible slots, which greatly slows down the prediction. To tackle this, we introduce an inverse paradigm for prompting. Different from the classic prompts mapping tokens to labels, we reversely predict slot values given slot types. Such inverse prompting only requires a one-turn prediction for each slot type and greatly speeds up the prediction. Besides, we propose a novel Iterative Prediction Strategy, from which the model learns to refine predictions by considering the relations between different slot types. We find, somewhat surprisingly, the proposed method not only predicts faster but also significantly improves the effect (improve over 6.1 F1-scores on 10-shot setting) and achieves new state-of-the-art performance.

1 Introduction

Few-shot learning (FSL) aims at learning a model from only a few examples and is regarded as one of the key steps toward more human-like artificial intelligence (Wang et al., 2020). Recently, prompt-based methods achieve impressive results and show promising prospects for few-shot learning of Natural Language Processing (NLP) (Liu et al., 2021a; Zhao et al., 2021).

Prompt-based methods reformulate a target task into the language modeling problem, which takes advantages of the powerful pretrained Language Models (LM) (Devlin et al., 2019; Liu et al., 2019;

*Equal contributions.

†Corresponding author.

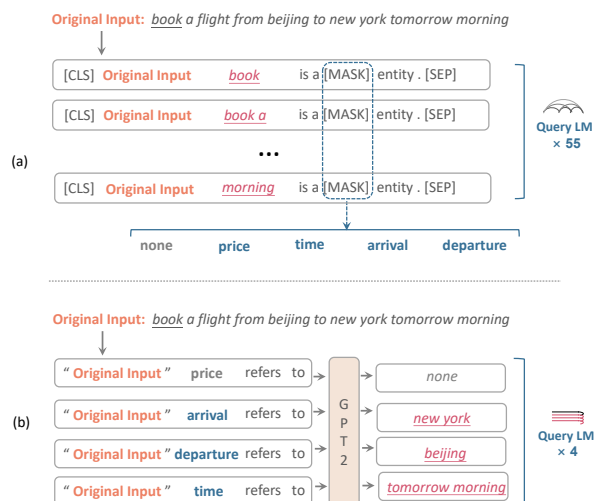


Figure 1: An example of normal (a) and inverse (b) prompting methods for slot tagging. For normal prompts, identifying all slots in the query sentence requires enumeration of all spans, while inverse prompt only needs 1-time prediction for each label.

Lewis et al., 2020; Brown et al., 2020). For example, when classifying the sentiment of the movie review "no reason to watch", prompting methods insert a piece of text "It was", i.e. prompts, to the input example, getting "No reason to watch. It was ___". It is natural to expect a higher probability from the LM to fill the template with "terrible" than "great", and the original task is then converted to a language modeling task. Such conversion reduces the gap between pretraining and target tasks, which allows less dependency on target task data and helps to achieve better performance in low data scenarios (Gao et al., 2021).

However, while achieving great success in sentence level tasks, prompting-based methods show incompatibility for sequence labeling tasks, such as slot tagging. Firstly, the aforementioned prompting paradigm is quite inefficient for slot tagging tasks. Different from the sentence-level tasks that classify samples of whole sentences, slot tagging sam-

ples are multiple consecutive words in a sentence. Therefore, as shown in Fig. 1, to find all the possible slots, prompt-based methods have to enumerate all n-gram word spans, and then query LM for each of them, which greatly slows down the prediction (Cui et al., 2021). Further, as a structure prediction problem, slot tagging benefits from taking the dependencies between labels into account (Ma and Hovy, 2016; Hou et al., 2020). For example in Fig. 1, where the `arrival` entity often appears after a `departure` entity. Such label dependency is hard to be captured by current prompting methods since they predict labels one-by-one independently.

To tackle the above issues, we introduce an inverse paradigm for prompting. Different from the classic prompts mapping tokens to labels, we reversely predict slot values given slot types. For the example in Fig. 1, we use an inverse prompt to modify the input as “*book a flight from Beijing to New York tomorrow morning. arrival refers to ___*”, and then LM is able to decode multi-word span “New York” at a time. Compared to the classic prompts that require predictions for every n-gram word span (55-times in Fig. 1), we only need to perform decoding for V -times, where V is the number of label types (4-times in Fig. 1), which therefore greatly speeds up the prediction. Surprisingly, experiments show the proposed method not only predicts faster but also significantly improves the performance, indicating that prompting LM reversely is a better fit for the slot tagging task. Besides, to further improve the prediction accuracy, we propose a novel Iterative Prediction Strategy, from which the model learns to refine predictions by considering the relations between different slot types.

To summarize the contribution of this work:

- (1) We introduce the idea of inverse prediction to prompting methods for slot tagging tasks, which greatly speeds up the prediction process.
- (2) We propose an Iterative Prediction Strategy for learning and prediction with slot tagging prompts, which allows the prompting model to consider dependency between different slot types and refine prediction.
- (3) We extensively evaluate the proposed method in various few-shot settings, where the proposed method brings significant improvements not only in speed but also in accuracy.

The code and data are available at <https://github.com/AtmaHou/PromptSlotTagging>.

2 Background

In this section, we begin with a formal definition of the few-shot slot tagging task (§2.1), and then introduce the conventional sequence labeling approaches (§2.2) and recent prompts-based methods (§2.3) for this task.

2.1 Few Shot Slot Tagging

Slot tagging aims at finding key slots within a sentence, such as time or location entities. Given an input sentence $\mathbf{x} = (x_1, x_2, \dots, x_n)$ as a sequence of words, a slot tagging model extracts all M slot label-values pairs $\mathbf{y} = \{(l_i, s_i)\}_{i=1}^M$ in the sentence, where l_i is the i th label in the label set L and $s_j^k = \{x_j, \dots, x_k\}$ is a word span starting from x_j and ending with x_k .

In few-shot settings, model are often evaluated on multiple low-resource domains $\{D_L^{(1)}, D_L^{(2)}, \dots\}$, which is called target domain (Wang et al., 2020). Each target domain $D_L^{(j)}$ only contains a few labeled instances called support set $S = \{(\mathbf{x}^{(i)}, \mathbf{y}^{(i)})\}_{i=1}^{N_S}$, which usually includes K examples (K-shot) for each of N labels (N-way). On each target domain, given support set examples as references, few-shot slot tagging models are required to make predictions for query set samples. Optionally, some few-shot settings also include a set of data-rich domains $\{D_H^{(1)}, D_H^{(2)}, \dots\}$ called source domains, which are used for pretraining of few-shot models.

2.2 Conventional Sequence Labeling Approaches

Conventional approaches often formulate slot tagging as a sequence labeling problem, where each word in input is associated with a sequence label. Given sentence $\mathbf{x} = (x_1, x_2, \dots, x_n)$ as input, these method predicts the best-match sequence labels $\mathbf{y} = (y_1, y_2, \dots, y_n)$. To predict slots with multiple words, sequence labeling approaches adopt a “BIO” labeling strategy, which uses “B” to mark the begin word of a slot, “I” to mark the inner words of a slot and “O” to mark non-slot words. For the example in the Fig. 2, `B-time` is tagged to the first word in a time slot, `I-time` is tagged to a non-begin word within a time slot, and `O` label refers to non-slot words. As shown in Fig. 2(a), few-shot



Figure 2: Illustration of conventional sequence labeling method (a) and classic prompting methods (b).

sequence labeling model is usually formulated as:

$$\begin{aligned}
 \mathbf{h}_{1:n} &= \text{Encoder}(\mathbf{x}_{1:n}), \\
 p(\mathbf{y}_i | \mathbf{x}, S) &= \text{Softmax}(\text{Decoder}(\mathbf{h}_i)), \\
 &\quad (i \in [1, 2, \dots, n]), \\
 \mathbf{y}^* &= (y_1, y_2, \dots, y_n) = \arg \max_{\mathbf{y}} p(\mathbf{y} | \mathbf{x}, S),
 \end{aligned}$$

where S is a K-shot support set, Encoder is usually a pretrained language model such as BERT (Devlin et al., 2019), $\mathbf{h}_{1:n}$ is the hidden state of the encoder with a dimension d_h , and Decoder can either be a linear layer, a CRF layer or any other parametric or non-parametric classifier.

2.3 Sequence Labeling with Prompts

Prompt-based methods have been proven effective in many NLU tasks, especially in few-shot settings, but things become complicated when it comes to slot tagging tasks. To identify the slot label for a word span $s_i^j = \{x_i, \dots, x_j\}$ in sentence \mathbf{x} , previous works construct templates, e.g., “[\mathbf{x}] [s_i^j] is a [z] entity”, and prompt a pretrained language model with such templates to predict label-related words [z] (Cui et al., 2021). For example in the Fig. 2(b), predicting the time slot can be achieved as “book a flight from Beijing to New York tomorrow morning. tomorrow morning is a time entity.” However, to find all possible slots, these methods need to traverse all the n-gram spans $s_i^j, i, j \in [1, n]$ in a sentence, which is quite expensive in time and computation.

3 Method

To remedy the high cost of prompt prediction mentioned in the previous section, we introduce a novel inverse paradigm for prompting of slot tagging task, which significantly improves the speed of prediction by transforming the past fill-in-the-blank problem into a generative task. Specifically, we first introduce the construction of our inverse prompts templates (§3.1), and then describe how to use inverse prompts during training and inference (§3.2).

Further, we propose an Iterative Prediction Strategy to refine prediction by considering the relation between different slot types (§3.3). The overview of proposed method is shown in Fig. 3.

3.1 Prompt Creation

In this section, we introduce the creation of the proposed inverse prompts, which includes three main components: the label mapping, the inverse template and the control tokens.

Label Mapping Before prompt construction, we first need to convert each label into a word form that can be easily understood by the pre-trained language model. We employ a label mapping process to achieve this, which use a one-to-one mapping function to convert the label set $L = \{l_1, \dots, l_{|L|}\}$ to a natural language word set $\hat{L} = \{\hat{l}_1, \dots, \hat{l}_{|L|}\}$. For example, in Fig. 3, we convert the label set $L = \{\text{from.Loc, to.Loc, Time, Price}\}$ to a natural language label set $\hat{L} = \{\text{departure, arrival, time, price}\}$.

Inverse Template Prompt template is a piece of sentence with blanks, which is used to modify the original inputs and get prompting inputs for a pretrained language model. To achieve inverse prompting, our template fills in an original sentence and a label as prefixes and subsequently leaves blanks for the LM to generate the corresponding slot values. Specifically, given an input sentence s and a set of mapped labels \hat{L} , for each mapped label $\hat{l}_i \in \hat{L}$, the inverse template is defined as:

$$\text{“} \mathbf{x} \text{” } \hat{l}_i \text{ refers to } _$$

For instance, in Fig. 3, we fill the input “book a flight from beijing to new york tomorrow morning” and each label in \hat{L} into the template to get four prompted inputs p :

“book a flight from beijing to new york tomorrow morning” **departure** refers to _

“book a flight from beijing to new york tomorrow morning” **arrival** refers to _

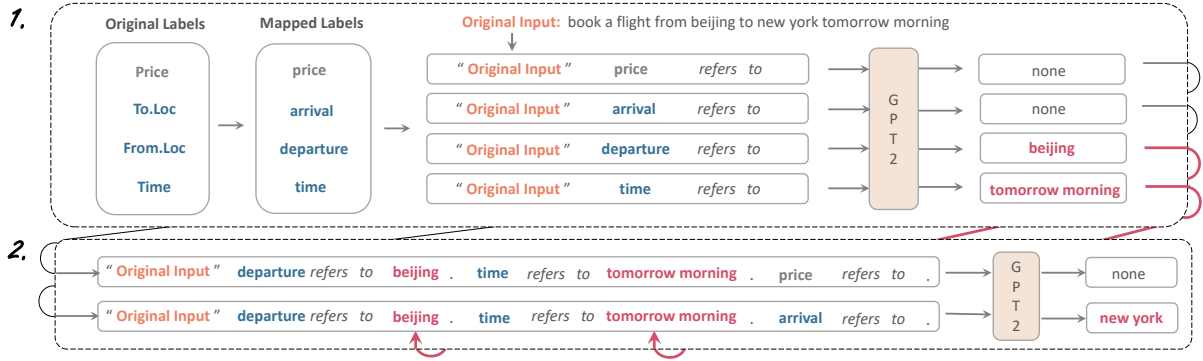


Figure 3: Overview of the proposed method with Inverse Prediction and Iterative Prediction Strategy. We first embed the input sentence with inverse prompts and directly decode slot values given slot types. Then we iteratively refine predictions by reinforcing the prompts with predicted slot-value pairs.

“book a flight from beijing to new york tomorrow morning” **time** refers to __

“book a flight from beijing to new york tomorrow morning” **price** refers to __

Control tokens Additionally, we introduce control tokens C to complete the prompts function for the slot tagging task. In order to recognize the case that there’s no corresponding entity of the queried slot type, we introduce $\langle \text{NONE} \rangle$ token to pad the output, and in practice, we use “none” as $\langle \text{NONE} \rangle$ token to make the model output more natural. In order to tag more than one entity of the same slot type, we introduce “;” as $\langle \text{SEP} \rangle$ to divide more than one entity of the same slot type. And we also use “.” as $\langle \text{END} \rangle$ token to indicate the end of a single generation.

3.2 Training and Inference with Inverse Prompts

Till now, we have presented the construction of the inverse prompt. This section will show how to perform training and inference with the prompts.

Training At the training time, we pre-construct the prompt with answers such as “book a flight from beijing to new york tomorrow morning” **departure** refers to **new york** . Then we finetune a pre-trained language model with the answered prompts, and we only calculate loss on the answer tokens (i.e. **new york**) instead of the loss on the whole sentence.

$$L = \sum_{i>|p|} \text{CE}(\hat{y}_i, y_i)$$

where $|p|$ is the length of the prompted input, \hat{y}_i denotes the model predictions, and y_i is the pre-constructed answer.

Inference At the inference time, we feed the prompted inputs into the fine-tuned pre-trained language model and let LM generate the appeared slot values. During generation, we restrict LM to generate only words that appear in the original input sentence or predefined control words. For each prompted input p , the next token $t_k \in \mathcal{x} \cup C$ is determined by language model probability:

$$t_k = \arg \max_{t_k \in \mathcal{x} \cup C} p_{\text{LM}}(t_k | p; t_{1:k-1})$$

Note that restricting the scope of output tokens is crucial to the performance.

3.3 Iterative Prediction Strategy

In the previous section, different slot types are predicted separately. To consider the relations between different slot types, we introduce the Iterative Prediction Strategy, which also provides the model a second chance to revise those unrecognized entities. We assume that different labels are interactive, so the predicted slots could be used as a hint to help predict the missed slots. For example in Fig. 3, it is often easier to generate the “arrival” slot given the results of “departure” and “time”. Motivated by this, as shown in the Fig. 3, we construct another template that concatenates those filled prompts as additional generation condition and use them to revise the slot values that are “none” in the first round of prediction. Below we describe the details of this strategy during training and testing.

Training At the training time, we simulate the cases where the slots are not recognized to enable the model to revise the **none** slot values. We do this by randomly constructing **none** slot value examples. For example, at training time, suppose there

are four training prompts filled with true answers:
“book a flight from beijing to new york tomorrow morning” **departure** refers to **beijing** .

“book a flight from beijing to new york tomorrow morning” **arrival** refers to **new york** .

“book a flight from beijing to new york tomorrow morning” **time** refers to **tomorrow morning** .

“book a flight from beijing to new york tomorrow morning” **price** refers to **none** .

We randomly select some occurred labels (e.g., “arrival”) pretending it was not predicted, and construct a second round prompt:

“book a flight from beijing to new york tomorrow morning” **departure** refers to **beijing** . **time** refers to **tomorrow morning** . **price** refers to **none** . **arrival** refers to ___ .

By using these second round prompts for model training, we encourage the language model to find those unrecognized slots in the first round prediction and allow the model to consider relationships between labels.

Inference During the inference time, we construct the second-round prompts and revise the slots that are not recognized in the first round. For example in the Fig. 3, the model predict **none** value for “price” and “arrival” slot in the first round. We then construct another iteration of the prompted inputs that query the unrecognized slots, given all the labels and slot values that have been predicted:

“book a flight from beijing to new york tomorrow morning” **departure** refers to **beijing** . **time** refers to **tomorrow morning** . **arrival** refers to ___ .

“book a flight from beijing to new york tomorrow morning” **departure** refers to **beijing** . **time** refers to **tomorrow morning** . **price** refers to ___ .

The model is expected to predict the first-round missed slots during the second iteration, considering relations between labels.

4 Experiment

We evaluate the performance of the proposed method on two classic few-shot scenarios: (1) Setting with Only In-domain data, where all training data are only a few labeled support data. (2) Setting with Meta Source Tasks, where some additional data-rich source domains are available for pretraining.

Evaluation To use same evaluation criteria as conventional sequence labeling methods, we need to label tokens reversely and get output in same

format. After generation, we first separate outputs into slot values. For each slot value, we label tokens in the source sentence with three principles: (1) Slot value is complete: only if the whole slot value matches a span in the source sentence, we label it with the corresponding label. (2) Choose the first overlap predicted slot span: if any token in the source sentence has been labeled, we do not relabel this token even when it matches another slot value. (3) Use BIO labels: add “B-” to the beginning token of the slot span, add “I-” to the non-begin token of the slot span, and label non-slot tokens with “O”. After labeling tokens reversely, we evaluate F1 scores within each few-shot episode.¹

4.1 Setting with Only In-domain data

Datasets For few-shot setting without source domain transfer, we conduct experiments on three few-shot datasets with only in-domain data: MIT-Restaurant Review (Liu et al., 2013), MIT-Movie Review (Liu et al., 2013) and MIT-Movie-Hard Review.² We conduct experiments with $K \in \{10, 20, 50, 100, 200, 500\}$ shots settings to fully evaluate the performance of our method in all three datasets. To overcome the randomness associated with support set selection, we sample 10 different support set for each K -shot setting and report averaged results. All models are trained and tested with the same data.

Implements Our model employs the smallest GPT2 (Radford et al., 2019) pre-trained model as the base model for fine-tuning, and no new parameters are introduced. Besides, we set the learning rate as $6.25e - 5$ and batch size as 2 for few-shot training. For all our experiments, we finetune the model only on few-shot support set for 2 epochs (4 on 10/20 shots settings) with the AdamW optimizer and linear decaying scheduler. Since there is no development set, all hyperparameters are roughly set based on experience without tuning. Data and code used are public available.

Baselines In our experiments, we compare with competitive baselines including both conventional sequence labeling methods and recent prompt-based methods.

¹For each episode, we calculate the F1 score on query samples with conl1eval script: <https://www.clips.uantwerpen.be/conl12000/chunking/conl1eval.txt>

²MIT-Movie Review has two datasets: a simple one and a complex one. We denote the simple one as MIT-Movie and combine both as MIT-Movie-Hard.

Model	MIT-Restaurant					
	10	20	50	100	200	500
Wiseman and Stratos (2019) + PT	4.1	3.6	4.0	4.6	5.5	8.1
Ziyadi et al. (2020) + PT	27.6	29.5	31.2	33.7	34.5	34.6
Huang et al. (2020) + PT	46.1	48.2	49.6	50.0	50.1	-
Sequence Labeling BART + PT	8.8	11.1	42.7	45.3	47.8	58.2
Sequence Labeling BERT + PT	27.2	40.9	56.3	57.4	58.6	75.3
Template-based BART + PT	53.1	60.3	64.1	67.3	72.2	75.7
Sequence Labeling BERT	21.8	39.4	52.7	53.5	57.4	61.3
Template-based BART	46.0	57.1	58.7	60.1	62.8	65.0
Ours	49.35	60.48	65.34	70.41	73.69	76.13
Ours + Iterative	52.10	61.49	66.83	70.98	73.97	76.37

Model	MIT-Movie-Hard					
	10	20	50	100	200	500
Wiseman and Stratos (2019) + PT	3.1	4.5	4.1	5.3	5.4	8.6
Ziyadi et al. (2020) + PT	40.1	39.5	40.2	40.0	40.0	39.5
Huang et al. (2020) + PT	36.4	36.8	38.0	38.2	35.4	38.3
Sequence Labeling BART + PT	13.6	30.4	47.8	49.1	55.8	66.9
Sequence Labeling BERT + PT	28.3	45.2	50.0	52.4	60.7	76.8
Template-based BART + PT	42.4	54.2	59.6	65.3	69.6	80.3
Sequence Labeling BERT	25.2	42.2	49.64	50.7	59.3	74.4
Template-based BART	37.3	48.5	52.2	56.3	62.0	74.9
Ours	52.07	59.11	65.63	69.35	72.36	75.03
Ours + Iterative	53.31	60.19	66.13	69.63	72.45	74.83

Model	MIT-Movie					
	10	20	50	100	200	500
Sequence Labeling BERT	50.60	59.34	71.33	-	-	-
NNShot	50.47	58.94	71.17	-	-	-
StructShot	53.19	61.42	72.07	-	-	-
Template-based BART	49.30	59.09	65.13	-	-	-
EntLM	57.31	62.36	71.93	-	-	-
Ours	57.04	67.86	76.81	80.28	82.43	84.55
Ours + Iterative	59.74	70.09	77.60	80.63	82.64	84.51

Table 1: F1 scores of few-shot slot tagging task on three different datasets. 10 indicates 10 instances for each entity type. +PT denotes the model is pre-trained on additional datasets. +Iterative denotes enhance model with Iterative Prediction Strategy.

- **Sequence Labeling BERT** (Devlin et al., 2019) can be seen as a BERT-based sequence labeling baseline which fine-tunes the BERT model with a token-level linear classifier head.

- **Template-based BART** (Cui et al., 2021) is a prompt-based method that query BART-based LM (Lewis et al., 2020) every possible span in sentence if it belong to a certain category and therefore also need to enumerate all label for inference.

- **NNShot and StructShot** (Yang and Katiyar, 2020) are two metric-based few-shot learning approaches for slot tagging and NER. NNShot is an instance-level nearest neighbor classifier for few-shot prediction, and StructShot promotes NNShot with a Viterbi algorithm during decoding.

- **EntLM** (Ma et al., 2021b) is a prompt-based method that leverage substitution between words of the same type to achieve one pass prediction.

Results Table 1 shows the results of the proposed method only finetuned on few-shot in-domain data. Among these results, we can observe that:

- (1) Our proposed method performs consistently better than all the baseline methods on all three datasets. It outperforms the strongest baseline Template-based BART which uses BART-large by average F1 scores on three datasets of 11.96 in 10-shot setting even with a much smaller pre-trained language model (the smallest GPT2).

- (2) Our proposed method is even comparable or outperforms those baselines with data-rich domain pre-training.

- (3) Our proposed method performs much better than baselines in fewer labeled samples settings, especially in 10 and 20 shot settings, which indicates our method can leverage information from limited labeled data more efficiently.

Model	5-shot Slot Tagging							
	We	Mu	Pl	Bo	Se	Re	Cr	Ave.
Bi-LSTM	25.44	39.69	45.36	73.58	55.03	40.30	40.49	45.70
SimBERT	53.46	54.13	42.81	75.54	57.10	55.30	32.38	52.96
TransferBERT	56.01	43.85	50.65	14.19	23.89	36.99	14.29	34.27
MN	38.80	37.98	51.97	70.61	37.24	34.29	72.34	49.03
WPZ+BERT	69.06	57.97	44.44	71.97	74.62	51.01	69.22	62.61
TapNet+CDT	67.83	68.72	73.74	86.94	72.12	69.19	66.54	72.15
L-WPZ+CDT	78.23	62.36	59.74	76.19	83.66	69.69	71.51	71.62
L-TapNet+CDT	69.58	64.09	74.93	85.37	83.76	69.89	73.80	74.49
ConVEx*	71.5	77.6	79.0	84.5	84.0	73.8	67.4	76.8
Ours	70.44	71.63	78.67	87.37	81.38	71.77	74.42	76.53
Ours + Iterative	70.63	71.97	78.73	87.34	81.95	72.07	74.44	76.73

Table 2: F1 score results on 5-shot Snips. * denotes using additional Reddit data for pretraining. Our methods achieve the best performance among those using same training data.

(4) Our method significantly outperformed Sequence Labeling BERT whose performance is quite poor on 10 and 20 shot settings, which indicates that the number of labeled data is too scarce for conventional sequence labeling tasks, and proves that the prompt-based method is effective in few-shot slot tagging tasks.

(5) The proposed Iterative Prediction Strategy consistently improves the slot tagging performance. The improvements become greater with fewer learning shots and the averaged improvement in 10 and 20 shot setting on three datasets are 2.23 and 1.44. This shows that when there is less data, the iterative revising mechanism is more important.

4.2 Setting with Meta Source Tasks

Datasets We also evaluate the model ability of transferring from data-rich domains to unseen few-shot domains and conduct experiments on SNIPS (Coucke et al., 2018) dataset. Following the data split provided by Hou et al. (2020), we construct 5-shot SNIPS datasets from the original SNIPS datasets. The few-shot SNIPS dataset consists of 7 domains with different label sets: GetWeather (We), Music (Mu), PlayList (Pl), Rate-Book (Bo), SearchScreenEvent (Se), BookRestaurant (Re), and SearchCreativeWork (Cr). Each domain contains 100 few-shot episodes, and each episode consists of a support set and a query.

Implements Following Henderson and Vulic (2021), we conduct our cross-domain experiments with 5-shot few-shot settings to evaluate the ability of our model to transfer from rich-data domains to unseen few-shot domains. For our proposed method, same as in-domain settings, we use the

smallest GPT2 as the base model, and no new parameters are introduced. We pretrain the model in source domains and fine-tune it on the target few-shot domain. We set learning rate as $6.25e - 5$ and batch size as 16 for pretraining and batch size as 2 for 5-shot finetuning. During finetuning, we use the same AdamW optimizer and linear decaying scheduler. The hyper-parameters are decided according to performance on the dev set. Data and code used are public available.

Baselines We provided competitive strong baselines, including traditional finetune-based methods and advanced few-shot learning methods.

- **Bi-LSTM** (Schuster and Paliwal, 1997) uses GLoVe (Pennington et al., 2014) embedding for slot tagging and is trained on the support sets.
- **SimBERT** is a metric-based method using cosine similarity of BERT-based embedding to label tokens with the most similar token’s label.
- **Matching Network (MN)** (Vinyals et al., 2016) is a few-shot sequence labeling model based on the matching network and uses BERT embedding.
- **TransferBERT** is a domain transfer-based conventional NER model using BERT, which is first pre-trained on source domains and then fine-tuned on the target domain support set.
- **WPZ** (Fritzler et al., 2019) is a metric-based few-shot slot tagging method similar to MN, but is based on the prototypical network (Snell et al., 2017).
- **TapNet+CDT, L-TapNet+CDT, L-WPZ+CDT** (Hou et al., 2020) are metric-based few-shot learning methods designed for slot tagging, which in-

roduces a CRF-based framework to consider the relation between different slots.

- **ConVEx** (Henderson and Vulic, 2021) is a fine-tuning-based method that models slot tagging as a cloze task and is first pre-trained on Reddit data then fine-tuned on few-shot slot tagging data. Note that the Reddit data is not used by our method and other baselines during the experiments.

Results Table 2 shows the results of the cross-domain few-shot setting, from which we can observe that:

(1) Our proposed method outperforms all the baselines except ConVEx which uses extra Reddit data in the cross-domain 5-shot setting. Despite using less training data, our model still achieves comparable results with Covex, proving its superiority.

(2) We outperform TransferBERT by 42.36 F1 scores which strongly proved that the prompt-based method can transfer more knowledge from the source domain and is more data-efficient than conventional methods.

(3) Our method outperforms metric-based few-shot learning baselines, for example, 2.24 F1 scores higher than L-TapNet+CDT, which proves its competitiveness compared to classical few-shot learning methods.

(4) Our Iterative Prediction Strategy improved Our method by about 0.5 F1 scores, demonstrating that the revising ability is likely to be transferable and is effective under cross-domain scenarios.

4.3 Analysis

Effects of Iterative Prediction Strategy As shown in Table 1, the proposed Iterative Prediction Learning brings consistent improvement, especially in low-resource settings. It works by revising predictions with a second-round query to recognize those missing slots, which can bring an increase in recall score. To confirm that, we make a detailed analysis with precision score (P), recall score (R) and F1 score (F) in Table 3.

When Iterative Revise Strategy is added, we can get a rise in recall score about 4 points in 10-shot, 2~4 points in 20 shot and more than 1 points in other shot settings in exchange for a slight precision drop, resulting in a rise in overall F1 score by about 2 points in 10 and 20 shots.

Model	MIT-Restaurant			MIT-Movie			
	P	R	F	P	R	F	
10	Ours	67.7	42.4	52.1	84.0	46.4	59.7
	w/o Iter	69.4	38.3	49.4	85.9	42.7	57.0
	w/o Joint	68.8	38.9	49.7	85.6	43.0	57.2
20	Ours	70.1	54.7	61.5	83.5	60.4	70.1
	w/o Iter	71.6	52.3	60.5	86.3	55.9	67.9
	w/o Joint	70.92	53.45	61.0	85.6	56.9	68.3
50	Ours	73.6	61.2	66.8	83.6	72.4	77.6
	w/o Iter	75.4	57.6	65.3	85.9	69.5	76.8
	w/o Joint	74.3	59.2	65.7	84.7	70.8	77.1
100	Ours	76.1	66.5	71.0	84.4	77.2	80.6
	w/o Iter	78.0	64.2	70.4	86.3	75.0	80.3
	w/o Joint	76.7	66.0	71.0	85.0	76.5	80.5
200	Ours	77.8	70.5	74.0	85.4	80.0	82.6
	w/o Iter	79.5	68.7	73.7	87.1	78.2	82.4
	w/o Joint	78.0	70.1	73.8	85.1	79.9	82.4
500	Ours	79.4	73.5	76.4	86.3	82.8	84.5
	w/o Iter	81.0	71.8	76.1	87.9	81.4	84.6
	w/o Joint	79.6	73.4	76.4	86.6	82.1	84.3

Table 3: Ablation analysis Iterative Prediction Strategy **w/o Iter** denotes removing iterative strategy and **w/o joint** denotes using two separate models for the two iterative steps.

We further explore the effect of jointly learning of the first-round prediction and the second-round revising, and learn two abilities separately with two models. As shown in Table 3, w/o Joint model outperforms the no-revising model but lags behind the joint model. This indicates that joint learning the revising ability may act as data augmentation and brings more improvements than simple revising.

Efficiency Study Unlike Template-based BART that queries every n-gram span in the source sentence for each label (with $O(n^2 * m)$ where n is the length of the source sentence and m is the size of the label set) time complexity, our proposed method queries labels in the label set and directly generate slot values (with $O(n * m)$ time complexity). In theory, our method is much faster than Template-based BART, especially dealing with long sentences with sparse slots. To prove this, we conduct efficiency experiments by calculating the decoding time of each method on a Titan XP GPU with batch size as 8, and we set our max generation length at 40. As shown in Table 4, our method is about 8 times as fast as the Template-based BART method, and more than 3 times as fast as theirs with Iterative Prediction Strategy. During experiments, we find that as the number of labels increases, the model does become linearly slower, which may become limitations. However, the number of label types is usually smaller than the sentence length

Model	MIT-Movie	MIT-Restaurant
Baseline (Normal Prompt)	408.0	236.0
Ours	51.2	33.2
Ours + Iterative	119.4	71.4

Table 4: Comparison of the decoding time (second).

and much smaller than the number of spans, so that this growth does not affect the value of our method in practice. Besides, we find no significant correlation between the number of labels and our performance.

5 Related Work

Prompt-based learning Prompt-based learning approaches have been a broadly discussed topic since large language models like GPT models (Brown et al., 2020) are hard to fine-tune in low-resource scenarios. Early attempts (Schick and Schütze, 2021a,b) introduce manual prompts to text classification tasks. For natural language understanding (NLU) tasks, automatically searching discrete prompts methods are proposed such as Jiang et al. (2020); Shin et al. (2020); Gao et al. (2021). Meanwhile, due to the continuity of parameters in neural networks, continuous prompts for both text classification and generation tasks (Li and Liang, 2021; Liu et al., 2021b; Han et al., 2021) have been proposed. Unlike sentence-level tasks, prompting methods are very complicated for slot tagging and NER tasks. Cui et al. (2021) proposes a template-based method querying every slot span with each label which is expensive for decoding. Different from them, we introduce an inverse paradigm for prompting slot tagging tasks. Note that inverse prompting (Zou et al., 2021) has a similar name to our work but is entirely different in method and task. They aim to generate prompt templates inversely. Amendable generation (Tian et al., 2021) share a similar idea of using Iterative Prediction Strategy to generate and revise dialog state. By contrast, we focus on a different task for sequence labeling and first introduce an Iterative Prediction Strategy to prompting models. There are also generation-based methods for sequence labeling (Yan et al., 2021), which is not a prompting method, since it re-initializes decoding layers and learns a generative model from scratch.

Few-shot slot tagging Previous few-shot slot tagging methods focus on metric learning based methods, which classify tokens by word-label similarity

(Snell et al., 2017; Vinyals et al., 2016). Hou et al. (2020) leverage label name semantics to get better label representation and model label dependency in few-shot settings. Yang and Katiyar (2020) make a prediction based on the nearest neighbor sample instead of the nearest label representation. Besides, some works also explore training a model with additional data from non-slot-tagging task (Huang et al., 2020; Henderson and Vulic, 2021). Hou et al. (2021) improves few-shot slot tagging performance by jointly learning it with intent detection. Different from directly learning the few-shot slot tagging model, some researches explore to reformulate the slot tagging into other NLP tasks. Ma et al. (2021a) reforms slot tagging into a reading comprehension task. Yu et al. (2021) treats slot tagging as a retrieval task, Coope et al. (2020) uses span extracting task to extract slot and predict corresponding label and Cui et al. (2021) leverages prompts for few-shot NER. Different from those methods above, we are the first to reformulate the slot tagging task into a prompt-based generation task.

6 Conclusion

In this paper, to liberate the prompting methods from the burdensome prediction of slot-tagging tasks, we introduce a novel inverse prediction manner to prompting methods of slot-tagging, which significantly improves both the efficiency and accuracy. To further improve performance, we propose an Iterative Prediction Strategy for learning, which enables the prompting model to consider dependency between labels and refine prediction. Extensive experiments verify the effectiveness of the proposed method in various few-shot settings, indicating inverse prediction is a better fit for prompting of slot tagging task.

Acknowledgments

We are grateful for the helpful comments and suggestions from the anonymous reviewers. This work was supported by the National Key R&D Program of China via grant 2020AAA0106501 and the National Natural Science Foundation of China (NSFC) via grant 61976072 and 62176078.

Ethics Section

We analyze the limitations of the proposed method in both efficiency and effectiveness aspects, and the proposed method has no obvious potential risks.

All the scientific artifacts used/created are properly cited/licensed, and the usage is consistent with their intended use. This paper does not collect new datasets, nor does the data used contain sensitive information.

References

- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Sam Coope, Tyler Farghly, Daniela Gerz, Ivan Vulic, and Matthew Henderson. 2020. Span-convert: Few-shot span extraction for dialog with pretrained conversational representations. In *Proc. of the ACL*, pages 107–121.
- Alice Coucke, Alaa Saade, Adrien Ball, Théodore Bluche, Alexandre Caulier, David Leroy, Clément Doumouro, Thibault Gisselbrecht, Francesco Calta-girone, Thibaut Lavril, Maël Primet, and Joseph Dureau. 2018. Snips voice platform: an embedded spoken language understanding system for private-by-design voice interfaces. *CoRR*, abs/1805.10190.
- Leyang Cui, Yu Wu, Jian Liu, Sen Yang, and Yue Zhang. 2021. [Template-based named entity recognition using BART](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 1835–1845, Online. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Alexander Fritzier, Varvara Logacheva, and Maksim Kretov. 2019. [Few-shot classification in named entity recognition task](#). *Proceedings of the 34th ACM/SIGAPP Symposium on Applied Computing*.
- Tianyu Gao, Adam Fisch, and Danqi Chen. 2021. [Making pre-trained language models better few-shot learners](#). 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)*, pages 3816–3830, Online. Association for Computational Linguistics.
- Xu Han, Weilin Zhao, Ning Ding, Zhiyuan Liu, and Maosong Sun. 2021. [Ptr: Prompt tuning with rules for text classification](#).
- Matthew Henderson and Ivan Vulic. 2021. [Convex: Data-efficient and few-shot slot labeling](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT*, pages 3375–3389.
- Yutai Hou, Wanxiang Che, Yongkui Lai, Zhihan Zhou, Yijia Liu, Han Liu, and Ting Liu. 2020. [Few-shot slot tagging with collapsed dependency transfer and label-enhanced task-adaptive projection network](#). In *Proc. of ACL*, pages 1381–1393. Association for Computational Linguistics.
- Yutai Hou, Yongkui Lai, Cheng Chen, Wanxiang Che, and Ting Liu. 2021. Learning to bridge metric spaces: Few-shot joint learning of intent detection and slot filling. In *Findings of the ACL*, pages 3190–3200.
- Jiaxin Huang, Chunyuan Li, Krishan Subudhi, Damien Jose, Shobana Balakrishnan, Weizhu Chen, Baolin Peng, Jianfeng Gao, and Jiawei Han. 2020. Few-shot named entity recognition: A comprehensive study. *arXiv preprint arXiv:2012.14978*.
- Zhengbao Jiang, Frank F Xu, Jun Araki, and Graham Neubig. 2020. How can we know what language models know? *Transactions of the Association for Computational Linguistics*, 8:423–438.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proc. of the ACL*, pages 7871–7880.
- Xiang Lisa Li and Percy Liang. 2021. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv preprint arXiv:2101.00190*.
- Jiachang Liu, Dinghan Shen, Yizhe Zhang, Bill Dolan, Lawrence Carin, and Weizhu Chen. 2021a. What makes good in-context examples for gpt-3? *arXiv preprint arXiv:2101.06804*.
- Jingjing Liu, Panupong Pasupat, Yining Wang, Scott Cyphers, and Jim Glass. 2013. Query understanding enhanced by hierarchical parsing structures. In *2013 IEEE Workshop on Automatic Speech Recognition and Understanding*, pages 72–77. IEEE.
- Xiao Liu, Yanan Zheng, Zhengxiao Du, Ming Ding, Yujie Qian, Zhilin Yang, and Jie Tang. 2021b. [GPT understands, too](#). *CoRR*, abs/2103.10385.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Jianqiang Ma, Zeyu Yan, Chang Li, and Yang Zhang.

- 2021a. Frustratingly simple few-shot slot tagging. In *Findings of the ACL*, pages 1028–1033.
- Ruotian Ma, Xin Zhou, Tao Gui, Yiding Tan, Qi Zhang, and Xuanjing Huang. 2021b. Template-free prompt tuning for few-shot NER. *CoRR*, abs/2109.13532.
- Xuezhe Ma and Eduard Hovy. 2016. [End-to-end sequence labeling via bi-directional LSTM-CNNs-CRF](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics ACL*, pages 1064–1074. Proc. of the ACL.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. [Glove: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing EMNLP*, pages 1532–1543.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Timo Schick and Hinrich Schütze. 2021a. [Exploiting cloze-questions for few-shot text classification and natural language inference](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 255–269, Online. Association for Computational Linguistics.
- Timo Schick and Hinrich Schütze. 2021b. [It’s not just size that matters: Small language models are also few-shot learners](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2339–2352, Online. Association for Computational Linguistics.
- Mike Schuster and Kuldeep K. Paliwal. 1997. [Bidirectional recurrent neural networks](#). *IEEE Trans. Signal Process.*, 45(11):2673–2681.
- Taylor Shin, Yasaman Razeghi, Robert L Logan IV, Eric Wallace, and Sameer Singh. 2020. [Auto-prompt: Eliciting knowledge from language models with automatically generated prompts](#). *arXiv preprint arXiv:2010.15980*.
- Jake Snell, Kevin Swersky, and Richard Zemel. 2017. [Prototypical networks for few-shot learning](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Xin Tian, Liankai Huang, Yingzhan Lin, Siqi Bao, Huang He, Yunyi Yang, Hua Wu, Fan Wang, and Shuqi Sun. 2021. [Amendable generation for dialogue state tracking](#). *CoRR*, abs/2110.15659.
- Oriol Vinyals, Charles Blundell, Tim Lillicrap, Koray Kavukcuoglu, and Daan Wierstra. 2016. Matching networks for one shot learning. In *Proc. of NIPS*, pages 3630–3638.
- Yaqing Wang, Quanming Yao, James T. Kwok, and Lionel M. Ni. 2020. [Generalizing from a few examples: A survey on few-shot learning](#). *ACM Comput. Surv.*, 53(3):63:1–63:34.
- Sam Wiseman and Karl Stratos. 2019. [Label-agnostic sequence labeling by copying nearest neighbors](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics ACL*, pages 5363–5369.
- Hang Yan, Tao Gui, Junqi Dai, Qipeng Guo, Zheng Zhang, and Xipeng Qiu. 2021. [A unified generative framework for various NER subtasks](#). In *Proc. of the ACL/IJCNLP*, pages 5808–5822. Association for Computational Linguistics.
- Yi Yang and Arzoo Katiyar. 2020. [Simple and effective few-shot named entity recognition with structured nearest neighbor learning](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6365–6375, Online. Association for Computational Linguistics.
- Dian Yu, Luheng He, Yuan Zhang, Xinya Du, Panupong Pasupat, and Qi Li. 2021. Few-shot intent classification and slot filling with retrieved examples. In *Proc. of the NAACL*, pages 734–749.
- Tony Z Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. 2021. Calibrate before use: Improving few-shot performance of language models. *arXiv preprint arXiv:2102.09690*.
- Morteza Ziyadi, Yuting Sun, Abhishek Goswami, Jade Huang, and Weizhu Chen. 2020. [Example-based named entity recognition](#). *CoRR*, abs/2008.10570.
- Xu Zou, Da Yin, Qingyang Zhong, Hongxia Yang, Zhilin Yang, and Jie Tang. 2021. [Controllable generation from pre-trained language models via inverse prompting](#). In *KDD ’21: The 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Virtual Event, Singapore, August 14–18, 2021*, pages 2450–2460. ACM.