

# Domain-Oriented Prefix-Tuning: Towards Efficient and Generalizable Fine-tuning for Zero-Shot Dialogue Summarization

Lulu Zhao<sup>1\*</sup>, Fujia Zheng<sup>1\*</sup>, Weihao Zeng<sup>1</sup>, Keqing He<sup>2</sup>, Weiran Xu<sup>1†</sup>,  
Huixing Jiang<sup>2</sup>, Wei Wu<sup>2</sup>, Yanan Wu<sup>1</sup>

<sup>1</sup>Pattern Recognition & Intelligent System Laboratory

<sup>1</sup>Beijing University of Posts and Telecommunications, Beijing, China

<sup>2</sup>Meituan Group, Beijing, China

{zhaoll, fujia\_zheng, ZengWH, xuweiran}@bupt.edu.cn

{kqin}@bupt.cn, {jhx\_bupt}@163.com, {wuwei19850318}@gmail.com

## Abstract

The most advanced abstractive dialogue summarizers lack generalization ability on new domains and the existing researches for domain adaptation in summarization generally rely on large-scale pre-trainings. To explore the lightweight fine-tuning methods for domain adaptation of dialogue summarization, in this paper, we propose an efficient and generalizable **Domain-Oriented Prefix-tuning** model, which utilizes a domain word initialized prefix module to alleviate domain entanglement and adopts discrete prompts to guide the model to focus on key contents of dialogues and enhance model generalization. We conduct zero-shot experiments and build domain adaptation benchmarks on two multi-domain dialogue summarization datasets, TODSum and QMSum. Adequate experiments and qualitative analysis prove the effectiveness of our methods.

## 1 Introduction

Abstractive dialogue summarization task aims to distill the most critical information in a conversation to produce a concise version, involving chit-chat (Gliwa et al., 2019; Chen and Yang, 2020), meeting (Zhong et al., 2021), customer service (Liu et al., 2019; Zou et al., 2021b), and task-oriented dialogue scenarios (Zhao et al., 2021b). Compared to the single-speaker texts, summarizing a dialogue presents a unique set of challenges, such as unstructured expressions and information sparsity issues. Recently, large-scale generative pre-trained models (Lewis et al., 2020; Liu and Lapata, 2019) have promoted the development of abstractive dialogue summarization but they all require extensive human-annotated golden summaries, which makes them not scalable to new domains where only few/no labeled data is available. Considering that real-world applications often face the problem

\* The first two authors contribute equally.

† Weiran Xu is the corresponding author.

<b>Dialogue</b> U: I am looking for a train. The train should arrive by 9:45 and should depart from Cambridge. ..... A: What is your destination? U: I need to go to Norwich. A: How about TR2534 that leaves at 5:36? .....	<b>Dialogue</b> Professor A: if you found that nine was better than six that would be OK. ..... Professor A: Yeah, just work with the models, yeah. PhD B: yeah, just look at the length of the models and just see what happens. .....
<b>Dialogue State</b> Inform (destination=Norwich; arrive by=9:45; departure=Cambridge)	<b>Query</b> What did the professor think about improving the back recognizer for the Aurora task?
<b>Summary</b> The user asks duration, arrive by and train id. The departure of the train is Cambridge and the destination is Norwich, and the train arrives at 9:45 on Friday.	<b>Summary</b> The Professor thought that the experiments were pretty simple, so computational power was not really a problem. He also thought that it would be okay to increase the states from six to nine.

Figure 1: An example of TODSum dataset (Zhao et al., 2021b) with dialogue state and an example of QMSum dataset (Zhong et al., 2021) with query. Note that dialogue state and query are existing characteristics of these two datasets, respectively.

of data in the new domain, it is vital to develop low-resource dialogue summarization models for the target domain by leveraging limited annotated data of source domains.

Therefore, we try to explore the efficient domain adaptation of dialogue summarization models from the source domain  $D_s$  to the target domain  $D_t$ , where  $D_s$  only has limited annotated summaries and  $D_t$  has few/no labeled data. There exist some domain adaptation approaches that focus on continual pre-trainings using some large domain/task-related corpora. Yu et al. (2021) added multiple pre-training stages both on source domains and target domains. Further, Zou et al. (2021c) decomposed the pre-training into three procedures, i.e., the pre-training of encoder, decoder, and the combined encoder-decoder model. Fabbri et al. (2021) constructed pseudo summaries based on external Wikipedia data to simulate characteristics of target dataset. Note that all these methods re-

quire time-consuming pre-trainings or large-scale external corpora. They only focus on the heavy pre-training stage rather than the lightweight fine-tuning, which makes it labor-expensive and environmentally unfriendly (Schwartz et al., 2020) to practical applications.

Different from existing works that adopt general pre-trainings on the large-scale external corpora, we focus on exploring efficient fine-tuning methods specifically targeted at domain adaptation for the dialogue summarization task. We consider the following key principles while designing our methods: (1) **Efficiency**: We do not use any external data or pre-training and aim at leveraging efficient fine-tuning mechanisms based on existing summarization models. (2) **Domain Disentanglement**: Traditional models often memorize excessive knowledge from  $D_s$  and generate wrong summaries containing specific domain words in  $D_s$ . We aim to disentangle shared domain knowledge from  $D_s$ . (3) **Generalization**: Models often learn specific features of the source domain, making it difficult to generalize in a new domain (Peng et al., 2019). For example, models learn the surface language style specific to  $D_s$  rather than adapting the way of saliency estimation and summary generation to  $D_t$ . We encourage the summarizer to only focus on generic key contents rather than domain-specific attributes.

To be consistent with above principles, we propose a lightweight and efficient **Domain-Oriented Prefix-tuning** method, DOP, for domain adaptation of dialogue summarization. For efficiency, we focus on fine-tuning summarization models instead of performing pre-trainings like existing works, which reduces expensive and time-consuming computation. For domain disentanglement, we design a domain-oriented prefix module, which contains a novel prompt initialization mechanism. Concretely, we use domain words extracted by unsupervised LDA (Hoffman et al., 2010) to initialize continuous prompt vectors and fit the outputs of MLP and pre-computed BART to obtain initial parameters and representations of the prefix module. We also add a domain-oriented prefix sequence of key-value pairs to augment the classical attention layer, which is independently applied to all Transformer layers of pre-trained models to elicit the knowledge interactively and achieve overall optimization. In this case, different domain words from  $D_s$  and  $D_t$  can induce relevant domain knowledge while adapting to a new domain. For generalization, we construct

discrete prompts using dialogue states or queries, as shown in Figure 1, to guide the model to focus on key contents in dialogues and enhance generalization capability on unseen domains. Considering there is no unified and practical benchmark for domain adaptation of dialogue summarization<sup>1</sup>, we build domain adaptation benchmarks based on two existing multi-domain summarization datasets TODSum (Zhao et al., 2021b) and QMSum (Zhong et al., 2021). Extensive experiments demonstrate the benefits of our methods both in zero-shot and few-shot settings for domain adaptation.

Our contributions are threefold: (1) To the best of our knowledge, we are the first to explore fine-tuning methods for domain adaptation of dialogue summarization task, and establish two practical and comprehensive benchmarks for TODSum and QMSum datasets. (2) We propose a lightweight and efficient Domain-Oriented Prefix-tuning model, with domain word initialized prefix and discrete prompts, to elicit knowledge from large-scale pre-trained models interactively. (3) We conduct sufficient experiments and qualitative analysis to prove the effectiveness of our methods and discuss current challenges of domain adaptation for dialogue summarization.

## 2 Related Work

**Abstractive Dialogue Summarization** Dialogue summarization has drawn much attention recently. For chit-chat scenarios, researchers improved the performance on SAMSum dataset (Gliwa et al., 2019) via topic word information (Zhao et al., 2020; Liu et al., 2021), conversational structures (Chen and Yang, 2020, 2021), personal named entity planning (Liu and Chen, 2021), and semantic slots (Zhao et al., 2021a). Liu et al. (2019), Zou et al. (2021a,b), and Lin et al. (2021) proposed customer-service dialogue summarization datasets under diverse business scenarios. Besides, meeting transcripts, such as AMI (Carletta et al., 2005), ICSI (Janin et al., 2003), MediaSum (Zhu et al., 2021), and QMSum (Zhong et al., 2021), were also studied to promote dialogue summarization technologies. Zhao et al. (2021b) further proposed a task-oriented dialogue summariza-

<sup>1</sup>Existing work mostly takes news datasets like CNN/DailyMail (Hermann et al., 2015) as source domains and dialogue datasets like SAMSum (Gliwa et al., 2019) as target domains. We argue the setting doesn't fit in with industrial scenarios transferring from one business to another, like from attraction consultation to hotel reservation.

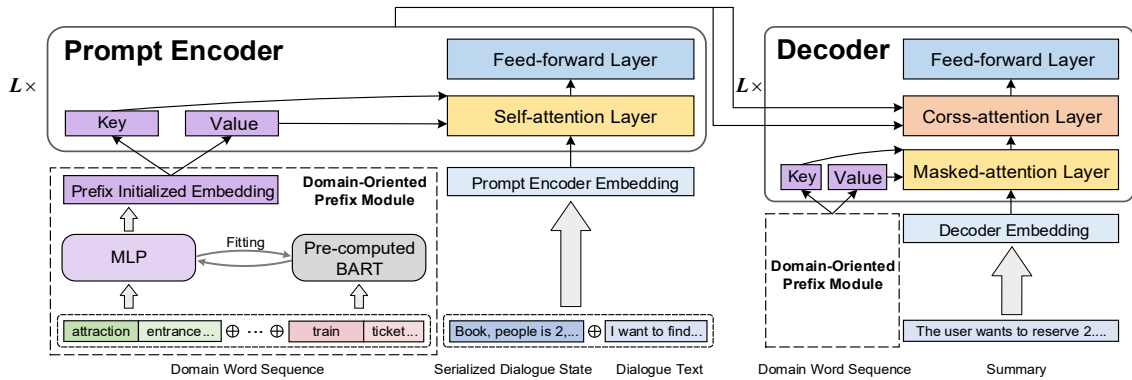


Figure 2: Overview of the Domain-Oriented Prefix-tuning model. The input sequence of domain-oriented prefix module includes domain words from four source domains, i.e., *attraction*, *hotel*, *taxi*, *train* and the domain words of *restaurant* domain are used as the prefix sequence of the target domain during the test. The input sequence of prompt encoder is the dialogue state from TODSum dataset and its original dialogue text. The input sequence of decoder is the golden summary. Here, we use the pre-computed outputs of BART as the target of the fitting process to initialize the parameters of MLP. Note that the grey pre-computed BART module does not participate in the training of the model.

tion dataset, TODSum, with dialogue state knowledge. Although great progress has been made in dialogue summarization, few people pay attention to the issue of domain adaptation in dialogue summarization. In this paper, we explore this issue in two multi-domain dialogue summarization datasets, i.e, TODSum and QMSum.

**Domain Adaptation in Summarization** Hua and Wang (2017) and Wang et al. (2019) adopted the document categories in news publications to build a multi-domain summarization dataset and investigated the domain shift for extractive summarization task. Yang et al. (2020), Zhang et al. (2020), Magooda and Litman (2020), and Fabri et al. (2021) regarded diverse summarization datasets as different domains and conducted the assessment of multi-domain settings. Furthermore, various stages of pre-trainings were added to narrow the gap between the pre-training in news domain and the fine-tuning in dialogue domain (Yu et al., 2021; Zou et al., 2021c). However, these methods focus on the heavy pre-training stage rather than the lightweight fine-tuning, which is time-consuming and relies on large-scale external corpora. Therefore, we try to explore the fine-tuning methods for domain adaptation of dialogue summarization task.

**Prompts in Summarization** With the arrival of GPT-3, prompt learning has become a nascent field, which conducts task-specific adaptation of large language models (LMs) via prepending an instruction. Schick and Schütze (2021) explored the fixed-prompt LM tuning for few-shot text summarization with manually crafted templates. Zhao

et al. (2021b) and Dou et al. (2021) further adopted the prompt+LM tuning strategy on text summarization task, where learnable prefix prompts are different types of guidance signals. Li and Liang (2021) investigated fixed-LM prompt tuning, where learnable prefix tokens are prepended to the input while parameters in pre-trained models are frozen. Following Li and Liang (2021), we design the domain information to initialize the continuous prefix module, and use discrete prompts and dialogue texts to optimize prefix parameters, which greatly reduces the size of parameters and is suitable for low-resource scenarios.

### 3 Problem Formulation

Domain adaptation of dialogue summarization aims to generate the summary  $y$  conditioned on the source dialogue  $x_d$ , where the training and test are in different domains. We add the domain word prefix  $x_{dw}$  and discrete prompt  $x_{dp}$  as additional input which we will describe details later. Note that we only update the prefix-related parameters and fix the parameters of BART. We train the model using the source data and test using the target data.

### 4 Methodology

As Figure 2 shows, our model is on the basis of the framework of BART, including a domain-oriented prefix module, a prompt encoder, and a decoder.

#### 4.1 Domain-oriented Prefix

To alleviate domain entanglement, we present a domain-oriented prefix module to obtain the shared

knowledge of the source domain  $D_s$  and the target domain  $D_t$ . It is designed as follows:

**Initialization** We extract some keywords from dialogue texts in each domain by LDA (Hoffman et al., 2010) and concatenate them all together as a domain word (prefix) sequence  $\mathbf{x}_{dw}$ <sup>2</sup>. Randomly initialized embeddings of the domain word sequence compose a learnable matrix  $M_\theta \in \mathbb{R}^{|\mathbf{x}_{dw}| \times d_m}$ .

**Parametrization** We use an MLP to encode the domain-oriented prefix module, which stably elicits knowledge from the large pre-trained model in the prefix-tuning process. Specifically, we first input the domain word sequence to the MLP and the pre-computed BART respectively, then re-train the MLP by fitting its outputs with the decoder hidden states of the pre-computed BART using MSE loss. In this fitting process, we only iteratively update MLP parameters  $\varphi \in \mathcal{R}^{d_m \times d_n}$  and keep the pre-computed BART fixed. Finally, we get the initialization parameters of MLP and use this pre-trained MLP to map the initialized embeddings of prefix representations for each Transformer layer both in prompt encoder and decoder:

$$M'_\theta[i, :] = \text{MLP}_\varphi(M_\theta[i, :]) \quad (1)$$

where  $i \in \mathbf{x}_{dw}$  and  $M'_\theta \in \mathbb{R}^{|\mathbf{x}_{dw}| \times d_n}$ . Note that this continuous prefix is applied for every layer of the large-scale pre-trained model independently.

## 4.2 Prompt Encoder

**Discrete Prompts** We utilize some discrete prompts to emphasize key elements in dialogues and enhance the model generalization to new domains. Here, discrete prompts are dialogue states of TODSum dataset or queries of QMSum. Considering that the original form of dialogue states is *book (people=5; day=Monday)* which is not compatible with BART encoder, we convert this structured information into a serialized sequence, i.e., *book, people is 5, day is Monday*, to improve the stability of training. Note that we do not make any changes to the query of QMSum dataset because it is already a serialized representation.

For prompt encoder, we firstly concatenate the discrete prompt sequence  $\mathbf{x}_{dp}$  and dialogue text sequence  $\mathbf{x}_d$  as the input sequence of encoder, i.e.,  $\mathbf{x}_{enc} = [\mathbf{x}_{dp}; \mathbf{x}_d]$ . Then, the  $\mathbf{x}_{enc}$  is fed into the

<sup>2</sup>We represent some of domain words for each domain in Appendix D.

prompt encoder based on the BART encoder, containing multiple Transformer layers. Note that we modify the self-attention mechanism by adding a domain-oriented prefix sequence of key-value pairs, which learns the knowledge from the pre-trained model through interactions with the dialogue text to carry out the overall task. For the typical  $l_e$ -th Transformer layer in encoder, the query ( $\mathbf{Q}_{l_e}$ ), key ( $\mathbf{K}_{l_e}$ ), and value ( $\mathbf{V}_{l_e}$ ) matrices are computed through linear transformations on the hidden states of  $\mathbf{x}_{enc}$ . Here, we further augment the  $\mathbf{K}_{l_e}$  and  $\mathbf{V}_{l_e}$ :

$$\mathbf{K}'_{l_e} = [\mathbf{P}_{l_e,k}; \mathbf{K}_{l_e}], \mathbf{V}'_{l_e} = [\mathbf{P}_{l_e,v}; \mathbf{V}_{l_e}] \quad (2)$$

where  $\mathbf{P}_{l_e,k}$ ,  $\mathbf{P}_{l_e,v}$  are computed through linear transformations on  $M'_\theta$ .  $\mathbf{K}'_{l_e}, \mathbf{V}'_{l_e} \in \mathbb{R}^{(|\mathbf{x}_{dw}|+|\mathbf{x}_{dp}|+|\mathbf{x}_d|) \times d_n}$  and  $l_e \in L$ . The augmented self-attention layer is finally calculated as follows:

$$\mathbf{A}_{self} = \text{softmax}(\mathbf{Q}_{l_e} \mathbf{K}'_{l_e}{}^T) \mathbf{V}'_{l_e} \quad (3)$$

## 4.3 Decoder

We also prepend the prefix module for decoder, where the cross-attention and masked-attention mechanisms are augmented in a similar way. The cross-attention between the  $l_e$ -th layer of prompt encoder and the  $l_d$ -th layer of the decoder is designed as:

$$\mathbf{A}_{cross} = \text{softmax}(\mathbf{Q}_{l_d} \mathbf{K}'_{l_e}{}^T) \mathbf{V}'_{l_e} \quad (4)$$

where  $\mathbf{Q}_{l_d}$  is computed through a linear transformation on the hidden states of the summary text  $\mathbf{x}_s$  and  $l_d \in L$ . Besides, the implementation of masked-attention layer is the same as the self-attention layer in the prompt encoder.

## 4.4 Training Strategy

In the domain-oriented prefix module, the parameter set of all linear transformations is symbolized as  $\alpha$ . For training strategy in DOP, we perform gradient updates on the following log-likelihood objective:

$$\max_{\alpha, \theta, \varphi} \log p_{\alpha, \theta, \varphi}(\mathbf{y}|\mathbf{x}) = \sum_{i \in \mathbf{y}} \log p_{\alpha, \theta, \varphi}(y_i | y_{<i}) \quad (5)$$

where the BART parameters  $\phi$  are fixed. The prefix parameters  $\alpha$ ,  $\theta$ , and  $\varphi$  are the only trainable parameters. During the training, we use the domain words from source domains as the prefix sequence. When the training is completed, we save all parameters of the domain-oriented prefix module and

Domains	Size	Dialog.len	Summ.len	DS.len
<b>Train</b>	345	120.67	24.93	18.29
<b>Taxi</b>	435	80.24	29.04	15.80
<b>Restaurant</b>	1,311	105.42	23.04	14.30
<b>Hotel</b>	636	145.16	30.06	21.38
<b>Attraction</b>	150	95.48	22.27	7.92
<b>All</b>	2,877	111.71	25.68	16.24

Table 1: Details of TODSum. "Dialog.len" denotes the average length of dialogues, "Summ.len" denotes the average length of summaries, and "DS.len" denotes the average length of serialized dialogue states.

Domains	Size	Dialog.len	Summ.len	QR.len
<b>Academic</b>	312	1,155.78	46.48	8.56
<b>Committee</b>	417	757.68	76.00	14.54
<b>Product</b>	847	971.65	63.96	13.36
<b>All</b>	1,576	951.49	63.68	12.73

Table 2: Details of QMSum. "QR.len" denotes the average length of queries.

drop the pre-computed BART module. During the test, the domain words from the target domain are mapped to the representations of prefixes only via the reserved  $\text{MLP}_\varphi(\cdot)$ , where the features of the source domain  $D_s$  can be transferred to the target domain  $D_t$ .

## 5 Experimental Setup

### 5.1 Datasets

We evaluate our model on two multi-domain dialogue summarization datasets and the details of statistics are shown in Table 1 and Table 2:

**TODSum:** This dataset is proposed by Zhao et al. (2021b), which is a task-oriented dialogue summarization dataset based on the classic dialogue dataset MultiWOZ (Budzianowski et al., 2018). According to the domain information, the dataset can be divided into five domains: *restaurant*, *hotel*, *attraction*, *taxi*, and *train*. Considering that parts of dialogues in TODSum contain multiple domains, in this paper, we firstly select all single-domain dialogues from TODSum as the dataset used for this experiment. Then we integrate all these samples in any four of five domains as source domains  $D_s$  and the other one is regarded as target domain  $D_t$ , where 200 samples extracted from  $D_s$  are as the validation set, the remaining as the training set, and samples of  $D_t$  are the test set.

**QMSum:** This dataset is proposed by Zhong et al. (2021), which contains hundreds of meeting transcriptions and includes three domains, i.e., *academic*, *committee*, *product*. In addition, each sample can be divided into several dialogue fragments according to some queries and the answer for the

corresponding query is its golden summary. Such (*query-dialogue-answer*) pairs are usually used for query-based meeting summarization tasks. In this paper, we separately integrate the training and validation sets, and test set in any two of three domains as the training data and validation data, i.e., the data of source domain  $D_s$ . All data of the other domain is used as the test data, i.e., the data of target domain  $D_t$ .

### 5.2 Baselines and Evaluation Metrics

We compare our methods with several baselines. The extractive baselines are included: (1) Lead-3; (2) Oracle; (3) BertExt (Liu and Lapata, 2019). Some abstractive methods are also added for comparison: (1) PGN (See et al., 2017); (2) Transformer (Vaswani et al., 2017); (3) BertAbs (Liu and Lapata, 2019); (4) BART (Lewis et al., 2020); (5) BART w. DS/QR (Zhao et al., 2021b); (6) Prefix-tuning (Li and Liang, 2021). For QMSum, we also introduce its benchmark (Zhong et al., 2021). Since this method feeds the extracted spans into BART, we integrate the results of this method with the results of BART. We use ROUGEs (Lin, 2004; Lin and Och, 2004) to quantitatively evaluate the performance of our methods. Our codes are publicly available<sup>3</sup>. We give the baselines and evaluation metrics in Appendix A.1 and Appendix A.2.

### 5.3 Training Details

Our implementation is based on the Hugging Face Transformer models<sup>4</sup>.  $\text{BART}_{\text{LARGE}}$  is used as a backbone and the source dialogue sequence is truncated to 1024 BPE tokens. For domain-oriented prefix module, the MLP maps 1024 dimension into 24576 dimension, which is calculated by  $2 * \text{number of decoder layers} * 1024$  and the numbers of domain words (prefix length) are set to 140 and 200 for TODSum and QMSum datasets. Following the settings in Li and Liang (2021), we use an AdamW optimizer (Loshchilov and Hutter, 2019) and a linear learning rate scheduler with initial rate of  $5 \cdot 10^{-5}$ , and the batch size is set to 5. Our model is trained on RTX 2080 Ti machines, taking only 5 minutes per epoch on TODSum dataset and 3 min per epoch on QSum dataset. However, BART w. DS takes 13 minutes and 8 minutes per epoch on TODSum and QMSum datasets. The reason for the shorter training time of our model is that the

<sup>3</sup><https://github.com/Zeng-WH/DOP-Tuning>.

<sup>4</sup><https://github.com/huggingface/transformers>

Models	Train			Taxi			Restaurant			Hotel			Attraction		
	2,332 / 200 / 345			2,242 / 200 / 435			1,366 / 200 / 1,311			2,041 / 200 / 636			2,527 / 200 / 150		
	R-1	R-2	R-L	R-1	R-2	R-L	R-1	R-2	R-L	R-1	R-2	R-L	R-1	R-2	R-L
<b>Lead-3</b>	20.36	2.78	16.07	24.20	7.34	20.75	28.27	6.10	23.49	23.86	4.58	18.80	22.76	5.28	19.66
<b>Oracle</b>	39.06	10.04	32.87	38.96	14.06	33.43	45.79	15.57	38.42	39.65	11.28	32.56	41.90	14.18	38.79
<b>BertExt</b>	39.19	9.71	33.24	38.49	13.57	33.36	40.64	12.34	34.43	35.96	9.71	30.10	36.25	11.19	31.41
<b>PGN</b>	32.50	10.47	29.33	32.48	7.79	29.82	33.63	10.78	31.47	32.18	9.36	30.93	32.66	9.95	30.29
<b>Transformer</b>	33.47	10.98	30.28	33.35	8.71	30.57	34.49	11.43	31.99	33.05	10.62	31.63	33.18	10.74	30.91
<b>BertAbs</b>	42.89	16.57	37.32	36.43	14.69	32.15	42.10	18.61	38.87	38.03	13.34	33.22	36.21	14.81	34.67
<b>BART</b>	46.82	18.42	42.06	39.98	15.79	34.41	47.02	22.62	44.93	40.84	14.20	36.83	43.67	20.23	41.44
<b>BART w. DS</b>	49.02	23.80	44.59	43.59	19.56	38.65	49.25	23.57	45.23	43.97	17.02	39.31	47.55	22.62	45.16
<b>Prefix-tuning</b>	45.92	22.70	41.06	41.89	19.47	39.62	47.19	24.20	42.99	43.41	18.75	36.75	44.48	22.43	40.94
<b>DOP (ours)</b>	<b>52.51</b>	<b>25.45</b>	<b>47.78</b>	<b>47.14</b>	<b>24.37</b>	<b>42.75</b>	<b>51.28</b>	<b>32.68</b>	<b>47.44</b>	<b>48.44</b>	<b>24.58</b>	<b>41.45</b>	<b>52.90</b>	<b>30.51</b>	<b>49.48</b>

Table 3: ROUGE scores of the zero-shot setting for TODSum. Results are averaged over three random runs. "DS" denotes the dialogue states. Values in the second row denote the size of train/valid/test set. ( $p < 0.05$  under t-test)

Models	Academic			Committee			Product		
	1,069 / 195 / 312			981 / 178 / 417			614 / 115 / 847		
	R-1	R-2	R-L	R-1	R-2	R-L	R-1	R-2	R-L
<b>Lead-3</b>	14.86	2.68	12.62	23.18	7.71	19.33	18.61	4.20	15.55
<b>Oracle</b>	32.44	8.78	27.92	54.92	35.03	51.85	40.23	15.29	35.83
<b>BertExt</b>	19.36	2.72	16.90	31.08	11.32	27.46	23.43	4.93	20.63
<b>PGN</b>	19.38	2.65	17.06	20.14	3.30	18.22	17.59	2.07	15.21
<b>Transformer</b>	18.72	2.53	16.97	20.01	3.16	17.88	17.14	2.02	15.14
<b>BertAbs</b>	20.36	2.45	17.84	20.93	3.23	18.26	17.32	1.86	15.30
<b>BART</b>	20.53	3.63	18.09	26.65	8.28	24.84	23.14	5.12	20.87
<b>BART w. QR</b>	22.45	4.11	20.14	27.93	9.73	25.87	24.89	5.93	22.39
<b>Prefix-tuning</b>	21.90	4.01	18.71	27.70	9.37	25.62	25.77	7.26	23.06
<b>DOP</b>	<b>24.43</b>	<b>5.84</b>	<b>21.49</b>	<b>30.28</b>	<b>11.63</b>	<b>27.33</b>	<b>29.85</b>	<b>9.41</b>	<b>26.88</b>

Table 4: ROUGE scores of the zero-shot setting for QMSum. "QR" denotes the queries. Same as TODSum, values in the second row denote the size of train/valid/test set. All results are averaged over three random runs. ( $p < 0.05$  under t-test)

trainable number of parameters of the DOP model is only 15% of the BART w. DS. For all experiments, we set the number of training epochs to 30. At the decoding phase, we use a beam size of 6 and max generated tokens of 125. The decoding takes 1.3 seconds per batch with the size of 2.

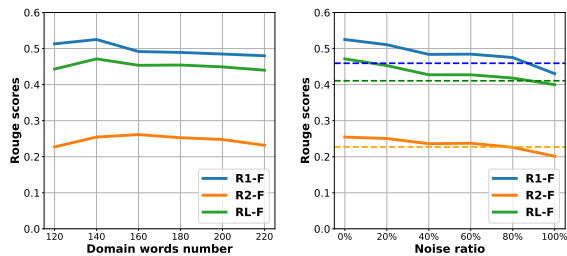
## 5.4 Main Results

**Results on TODSum** Table 3 presents the results of the zero-shot setting for TODSum dataset, where each of the five domains is regarded as the target domain respectively. The division of dataset in the second row intuitively shows that the amount of data in  $D_s$  is small and limited. We conduct experiments based on some common extractive models and some strong abstractive baselines. We also add a lightweight fine-tuning summarizer for comparison. As observed, for most ROUGEs, Prefix-tuning performs worse than BART and BART w. DS. It is because the dialogue text is long and complex, and using only 20% parameters of fine-tuning can not well understand the domain knowledge and identify the key contents in dialogues. Compared to Prefix-tuning of the same magnitude parameters, our model improves by 7%, 3%, 7% for *train* domain, 5%, 5%, 3% for *taxi* domain, 4%, 8%,

4% for *restaurant* domain, 5%, 6%, 5% for *hotel* domain, and 8%, 8%, 9% for *attraction* domain. This shows that the prefix module initialized by domain words and the discrete prompts composed of dialogue states play important roles. Besides, our model still surpasses BART w. DS, a full-parameter fine-tuning based model, which further illustrates that our model efficiently disentangles the knowledge of the source and target domains. Note that *attraction* domain gets the highest ROUGEs and the increased margins are also the largest. This may be due to the high overlaps between the *attraction* and source domains. All the results suggest that with limited data, the performance of our model still reaches state-of-the-art.

**Results on QMSum** Table 4 displays the results on zero-shot out-domain tests in three domains of QMSum dataset. As seen from the second row, in addition to the limited data, the source domain size may be even less than the target domain size, i.e., *product* domain. The trend of overall performance is consistent with that of TODSum dataset<sup>5</sup>, where the improvement in *product* domain is the most obvious and there are 5%, 4% and 5% increased for R-1, R-2, and R-L, respectively. However, all the ROUGEs are low, which is because there are no obvious domain words, leading to serious domain entanglement. Besides, due to the longer meeting text, it is hard to capture the key contents in dialogues, so as to the poor generalization in the target domain. Generally, these results show that the multi-domain setting in meeting summarization task is apparently necessary and meaningful. Meeting transcripts cover various domains, making the

<sup>5</sup>Note that ROUGEs of Oracle are very high in QMSum, which is because most parts of the golden summaries are directly copied from the original dialogue. It is determined by the characteristics of the QMSum and the results are consistent with Zhong et al. (2021).



(a) Number of domain words (b) Quality of domain words

Figure 3: Effect of different numbers and different qualities of domain words for *train* domain in TODSum. The dotted lines in (b) respectively represent the F1 sources of ROUGE-1, ROUGE-2, and ROUGE-L of the Prefix-tuning model.

adaptation of models particularly difficult.

## 6 Qualitative Analysis

### 6.1 Effect of Domain Words

**Number of Domain Words** We set different numbers of domain words, i.e., prefix length, to test the performance of our model DOP for *train* domain in TODSum dataset. As shown in Figure 3 (a), among these setting candidates, there is a threshold (140) that allows the ROUGEs to reach the peak. When choosing a fewer setting, the model does not perform well due to insufficient number of parameters, which is improved as the number increases. When being more than this threshold, a drop in performance occurs. One reason is that too long a sequence adds a large burden to BART, and the other one is that it introduces excessive noise. However, the change in the number of the domain words does not have a great impact on the performance (only 2~3% fluctuation), which also reflects the effectiveness of domain-oriented prefix module and the robustness of our model.

**Quality of Domain Words** For *train* domain in TODSum, we randomly replace a certain percentage of the domain words with words that are not related to the source domain. As Figure 3 (b) shows, when more noise is introduced, the model suffers more interference and its performance decreases. However, it performs better than Prefix-tuning. Only when the proportion of noise reaches 100%, the performance of our model is even worse than that of Prefix-tuning. This is because we especially use completely irrelevant words for initialization and fitting, which introduces more noise than simple random initialization and affects the performance of DOP. From this point of view, introducing high-quality domain words is good for

Model	R-1	R-2	R-L
<b>DOP (ours)</b>	52.51	25.45	47.78
<b>w/o DW</b>	48.87	23.81	44.52
<b>w/o DS</b>	47.59	23.25	43.41
<b>w/o DW &amp; DS</b>	45.92	22.70	41.06

Table 5: F1 scores of ablation study on *train* domain of TODSum dataset. "DW" denotes domain words and "DS" denotes dialogue states.

Model	R-1	R-2	R-L
<b>DOP (ours)</b>	30.28	11.63	27.33
<b>w/o DW</b>	29.10	10.59	26.53
<b>w/o QR</b>	28.47	9.60	26.38
<b>w/o DW &amp; QR</b>	27.70	9.37	25.62

Table 6: F1 scores of ablation study on *committee* domain of QMSum dataset. "DW" denotes domain words and "QR" denotes queries.

domain disentanglement and the quality of domain words is critical to summary generation.

### 6.2 Ablation Study

We perform ablation experiments on *train* domain of TODSum dataset and *committee* domain of QMSum dataset, as shown in Tables 5 and 6. We can observe that the removal of domain-oriented initialization in the prefix module will make the ROUGEs decrease significantly. Especially for TODSum, R-1, R-2, and R-L drop by 4%, 2%, and 3%, which shows the importance of domain word information for inducing the relevant knowledge while adapting to a new domain. In addition, after we remove the discrete prompts, i.e. dialogue state and query, the performance of the models becomes worse, but still outperforms the results of Prefix-tuning. It demonstrates that discrete prompts help the model pay attention to the key elements in the dialogue and improve the generalization of the model. Notably, our model achieves summary generation only by optimizing the domain-oriented prefix module, where domain words are available in all datasets. Since the DS and QR features happen to exist in the two datasets, we take advantage of them together with dialogue texts. When removing both DW and DS/QR at the same time, the model is equivalent to Prefix-tuning and the results are consistent.

### 6.3 Effect of Prefix Module in Encoder and Decoder

Since both the encoder and the decoder in our DOP introduce the prefix module, we verify their effects in the *train* and *committee* domains respectively.

Domain	Model	R-1	R-2	R-L
Train	DOP	52.51	25.45	47.14
	w/o enc.prefix	50.69	23.58	45.98
	w/o dec.prefix	45.51	22.15	40.67
Committee	DOP	30.28	11.63	27.33
	w/o enc.prefix	29.20	11.34	26.37
	w/o dec.prefix	29.11	11.29	26.30

Table 7: Effects of prefix module in encoder and decoder. We removed the prefix module from the encoder and decoder respectively to verify its effectiveness.




Figure 4: Effect of different few-shot data size for *train* domain in TODSum dataset. "0" denotes the zero-shot setting.

As shown in Table 7, when the encoder prefix module or decoder prefix module is removed, the performance of the model decreases, which shows that both are necessary and effective. In addition, we find that it is interesting that removing the prefix on the encoder side has a smaller impact on the model than removing the decoder side, especially in TODSum (about 5% on R-1). A reasonable explanation is that the prefix modules in encoder and decoder are responsible for different tasks. The prefix module on the encoder side assists the model to understand the dialogue, while the prefix module on the decoder side assists in model generation. Therefore, for summary generation, the prefix module in decoder is more helpful to the model.

#### 6.4 Effect of Training Data

**Performance in Few-shot Settings** For TODSum, we fix the size of source domain data, adding a certain amount of target (*train*) domain data for training, as shown in Figure 4. As the size of target domain data increases, the performance of both BART w. DS and DOP present a steady improvement trend and that of our DOP model is consistently better than BART w. DS, which is as expected. Besides, there is a substantial improvement from 50 to 100. This phenomenon shows that adding target knowledge can help the model learn about information of target domain and after adding a certain amount will help the model more efficiently.

**Effect of Source Domain Data Size** We keep




Figure 5: Effect of different source domain data size for *train* domain in TODSum dataset. "origin (2332)" denotes the same setting as the main results.

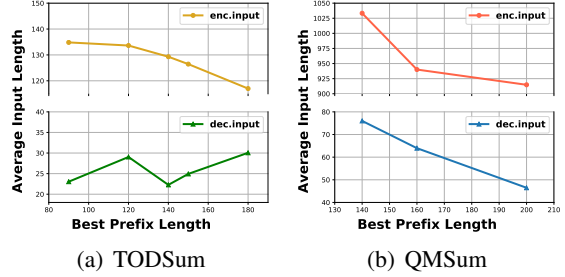


Figure 6: Relationships between length of encoder/decoder input and the best prefix length.

the zero-shot setting unchanged and adjust the size of source domain data for training to observe changes in the performance of the two models for *train* domain in TODSum. As shown in Figure 5, the smaller of data size, the greater the difference between the performance of the DOP and BART w. DS, that is, the performance of BART w. DS is getting worse, while the DOP maintains excellent performance steadily. This demonstrates that our DOP model is insensitive to the data scale and robustness to a certain extent. This also confirms that in the main experiment, our model can be outstanding in very limited and uneven data.

#### 6.5 Prefix Length vs. Input Length

Through experiments, we explore an interesting thing, that is, the prefix length (number of domain words) that makes the model perform best may be related to the input length. Based on this assumption, we collect the source input length, target input length, and their corresponding optimal prefix length from two datasets, as shown in Figure 6. We conclude a general rule that the longer inputs may prefer the shorter prefix. This phenomenon may serve as a research point in the future.

## 7 Discussion

### 7.1 Human Evaluation

We further conduct a manual evaluation to assess the models. We randomly select 50 samples from



Model	Flu.	Inf.	Fac.	Dom.	Red.
Ground Truth	4.95	4.56	4.28	4.71	4.33
BART	4.19	4.21	3.55	3.19	3.53
BART w. DS	4.36	4.34	4.09	3.35	3.62
Prefix-tuning	4.23	4.29	3.67	3.28	4.09
DOP	4.68	4.42	4.10	4.07	4.13

Table 8: Human evaluation on Fluency (Flu.), Informativeness (Inf.), Factual Consistency (Fac.), Domain Relevance (Dom.), and Redundancy (Red.) for TODSum dataset.

each target domain dataset of TODSum (Zhao et al., 2021b) and ask 10 professional linguistic evaluators to score the ground truth and summaries generated by 4 models according to 5 metrics: fluency, informativeness, factual consistency, domain relevance, and redundancy. Each metric is rated by 3 evaluators from 1 (worst) to 5 (best) and the scores for each summary are averaged. Note that the intra-class agreement score is 0.605.

As shown in Table 8, the fluency scores of all models are high, which is because that abstractive models fine-tuned on contextualized language backbones can generate fluent sentences (Lewis et al., 2020). For factual consistency, both DOP and BART w. DS achieve better performance than Prefix-tuning, which suggests that the dialogue state information guides the model to focus more on the key information, such as slots and intents. Besides, the DOP outperforms all baselines in the domain relevance metric. This demonstrates that the domain-oriented prefix module plays a crucial role in enhancing the ability of the model to identify domain-related features and disentangle the knowledge of the source and target domains. Surprisingly, the scores about the redundancy of Prefix-tuning and DOP are higher than that of BART and BART w. DS. This is because the model can efficiently extract key contents from a limited amount of data without relying on large-scale pre-trainings.

## 8 Challenges

Through the analysis of cases in Appendix C, we summarize two challenges of low-resource domain adaption for abstractive dialogue summarization task:

**1. Confusion between domains with high similarity:** We found that in domains with high-overlap vocabularies, i.e., *restaurant* and *hotel*, *train* and *taxi*, the model generates some domain-confusing sentences. Taken *hotel-restaurant* pair as an example, when *restaurant* is as the target domain,

a sentence like "book a restaurant room that can accommodate 3 people, ..." is generated, which is more likely to exist in the *hotel* domain. Note that this challenge does not affect the accuracy of key elements, but the language style is not appropriate.

**2. Information dispersion:** Because of the long input sequence, it is difficult for the models to pay attention to all aspects of the long dialogue and there will be problems with attention deviations on the key elements of dialogues, especially for this lightweight and small parameter training paradigm.

## 9 Conclusion

In this paper, we present a domain-oriented prefix-tuning model to handle the domain adaptation for dialogue summarization based on an efficient and generalizable fine-tuning method. The domain word initialized prefix module disentangles the target domain knowledge from the source domain and the discrete prompts enhance the generalization ability of the model. The experiments in zero-shot and few-shot settings show that our methods have made great progress on two datasets.

## Acknowledgement

We thank all anonymous reviewers. This work was partially supported by National Key R&D Program of China No. 2019YFF0303300 and Subject II No. 2019YFF0303302, DOCOMO Beijing Communications Laboratories Co., Ltd, MoE-CMCC "Artificial Intelligence" Project No. MCM20190701, BUPT Excellent Ph.D. Students Foundation CX2021204.

## References

- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gašić. 2018. [MultiWOZ - a large-scale multi-domain Wizard-of-Oz dataset for task-oriented dialogue modelling](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5016–5026, Brussels, Belgium. Association for Computational Linguistics.
- J. Carletta, S. Ashby, S. Bourban, Mike Flynn, Maël Guillemot, Thomas Hain, J. Kadlec, Vasilis Karaiskos, Wessel Kraaij, Melissa Kronenthal, G. Lathoud, M. Lincoln, Agnes Lisowska Masson, I. McCowan, W. Post, D. Reidsma, and P. Wellner. 2005. The ami meeting corpus: A pre-announcement. In *MLMI*.
- Jiaao Chen and Diyi Yang. 2020. [Multi-view sequence-to-sequence models with conversational structure for](#)

- abstractive dialogue summarization. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4106–4118, Online. Association for Computational Linguistics.
- Jiaao Chen and Diyi Yang. 2021. Structure-aware abstractive conversation summarization via discourse and action graphs. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Online. Association for Computational Linguistics.
- Zi-Yi Dou, Pengfei Liu, Hiroaki Hayashi, Zhengbao Jiang, and Graham Neubig. 2021. **GSum: A general framework for guided neural abstractive summarization**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4830–4842, Online. Association for Computational Linguistics.
- Alexander Fabbri, Simeng Han, Haoyuan Li, Haoran Li, Marjan Ghazvininejad, Shafiq Joty, Dragomir Radev, and Yashar Mehdad. 2021. **Improving zero and few-shot abstractive summarization with intermediate fine-tuning and data augmentation**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 704–717, Online. Association for Computational Linguistics.
- Bogdan Gliwa, Iwona Mochol, Maciej Biesek, and Aleksander Wawer. 2019. **SAMSum corpus: A human-annotated dialogue dataset for abstractive summarization**. In *Proceedings of the 2nd Workshop on New Frontiers in Summarization*, pages 70–79, Hong Kong, China. Association for Computational Linguistics.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. In *NIPS*.
- Matthew Hoffman, Francis R. Bach, and David M. Blei. 2010. **Online learning for latent dirichlet allocation**. In J. D. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. S. Zemel, and A. Culotta, editors, *Advances in Neural Information Processing Systems 23*, pages 856–864. Curran Associates, Inc.
- Xinyu Hua and Lu Wang. 2017. **A pilot study of domain adaptation effect for neural abstractive summarization**. In *Proceedings of the Workshop on New Frontiers in Summarization*, pages 100–106, Copenhagen, Denmark. Association for Computational Linguistics.
- A. Janin, D. Baron, J. Edwards, D. Ellis, D. Gelbart, N. Morgan, B. Peskin, T. Pfau, E. Shriberg, A. Stolcke, and C. Wooters. 2003. **The icsi meeting corpus**. In *2003 IEEE International Conference on Acoustics, Speech, and Signal Processing, 2003. Proceedings. (ICASSP '03)*, volume 1, pages I–I.
- M. 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 *ACL*.
- Xiang Lisa Li and Percy Liang. 2021. **Prefix-tuning: Optimizing continuous prompts for generation**. *arXiv preprint arXiv:2101.00190*.
- Chin-Yew Lin. 2004. **Rouge: A package for automatic evaluation of summaries**. In *Text summarization branches out*, pages 74–81.
- Chin-Yew Lin and Franz Josef Och. 2004. **Automatic evaluation of machine translation quality using longest common subsequence and skip-bigram statistics**. In *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics (ACL-04)*, pages 605–612.
- Haitao Lin, Liqun Ma, Junnan Zhu, Lu Xiang, Yu Zhou, Jiajun Zhang, and Chengqing Zong. 2021. **Csds: A fine-grained chinese dataset for customer service dialogue summarization**.
- Chunyi Liu, P. Wang, Jiang Xu, Zang Li, and Jieping Ye. 2019. **Automatic dialogue summary generation for customer service**. *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*.
- Junpeng Liu, Yanyan Zou, Hainan Zhang, Hongshen Chen, Zhuoye Ding, Caixia Yuan, and Xiaojie Wang. 2021. **Topic-aware contrastive learning for abstractive dialogue summarization**. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 1229–1243, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Yang Liu and Mirella Lapata. 2019. **Text summarization with pretrained encoders**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3730–3740, Hong Kong, China. Association for Computational Linguistics.
- Zhengyuan Liu and Nancy Chen. 2021. **Controllable neural dialogue summarization with personal named entity planning**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 92–106, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Ilya Loshchilov and Frank Hutter. 2019. **Decoupled weight decay regularization**. In *International Conference on Learning Representations*.
- Ahmed Magooda and Diane Litman. 2020. **Abstractive summarization for low resource data using domain transfer and data synthesis**.

- Ramesh Nallapati, Feifei Zhai, and Bowen Zhou. 2017. Summarunner: A recurrent neural network based sequence model for extractive summarization of documents. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- Xingchao Peng, Zijun Huang, Ximeng Sun, and Kate Saenko. 2019. Domain agnostic learning with disentangled representations. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 5102–5112. PMLR.
- Timo Schick and Hinrich Schütze. 2021. Few-shot text generation with pattern-exploiting training.
- Roy Schwartz, Jesse Dodge, Noah Smith, and Oren Etzioni. 2020. Green ai. *Communications of the ACM*, 63:54 – 63.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. Get to the point: Summarization with pointer-generator networks. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1073–1083, Vancouver, Canada. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, page 6000–6010, Red Hook, NY, USA. Curran Associates Inc.
- Danqing Wang, Pengfei Liu, Ming Zhong, Jie Fu, Xipeng Qiu, and Xuanjing Huang. 2019. Exploring domain shift in extractive text summarization.
- Ziyi Yang, Chenguang Zhu, Robert Gmyr, Michael Zeng, Xuedong Huang, and Eric Darve. 2020. TED: A pretrained unsupervised summarization model with theme modeling and denoising. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1865–1874, Online. Association for Computational Linguistics.
- Tiezheng Yu, Zihan Liu, and Pascale Fung. 2021. AdaptSum: Towards low-resource domain adaptation for abstractive summarization. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5892–5904, Online. Association for Computational Linguistics.
- Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter Liu. 2020. Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. In *International Conference on Machine Learning*, pages 11328–11339, Online.
- Lulu Zhao, Weiran Xu, and Jun Guo. 2020. Improving abstractive dialogue summarization with graph structures and topic words. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 437–449, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Lulu Zhao, Weihao Zeng, Weiran Xu, and Jun Guo. 2021a. Give the truth: Incorporate semantic slot into abstractive dialogue summarization. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2435–2446, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Lulu Zhao, Fujia Zheng, Keqing He, Weihao Zeng, Yuejie Lei, Huixing Jiang, Wei Wu, Weiran Xu, Jun Guo, and Fanyu Meng. 2021b. Todsum: Task-oriented dialogue summarization with state tracking. *arXiv preprint arXiv:2110.12680*.
- Ming Zhong, Da Yin, Tao Yu, Ahmad Zaidi, Mutethia Mutuma, Rahul Jha, Ahmed Hassan Awadallah, Asli Celikyilmaz, Yang Liu, Xipeng Qiu, and Dragomir Radev. 2021. QMSum: A new benchmark for query-based multi-domain meeting summarization. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5905–5921, Online. Association for Computational Linguistics.
- Chenguang Zhu, Yang Liu, Jie Mei, and Michael Zeng. 2021. MediaSum: A large-scale media interview dataset for dialogue summarization. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5927–5934, Online. Association for Computational Linguistics.
- Yicheng Zou, Jun Lin, Lujun Zhao, Yangyang Kang, Zhuoren Jiang, Changlong Sun, Qi Zhang, Xuanjing Huang, and Xiaozhong Liu. 2021a. Unsupervised summarization for chat logs with topic-oriented ranking and context-aware auto-encoders. In *AAAI*.
- Yicheng Zou, Lujun Zhao, Yangyang Kang, Jun Lin, Minlong Peng, Zhuoren Jiang, Changlong Sun, Qi Zhang, Xuanjing Huang, and Xiaozhong Liu. 2021b. Topic-oriented spoken dialogue summarization for customer service with saliency-aware topic modeling. In *AAAI*.
- Yicheng Zou, Bolin Zhu, Xingwu Hu, Tao Gui, and Qi Zhang. 2021c. Low-resource dialogue summarization with domain-agnostic multi-source pretraining. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 80–91, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

## A Experiment details

### A.1 Baselines

We describe baselines in detail as follows.

**Lead-3:** This method is commonly used in news summarization task, which treats the first three sentences of the document as the summary.

**Oracle:** The method is used to obtain an oracle through a greedy way similar to Nallapati et al. (2017), which treats the sentences that maximize the ROUGE-2 as the summary.

**BertExt:** Proposed by Liu and Lapata (2019), this model is extractive and its parameters are initialized with BERT.

**PGN:** Proposed by See et al. (2017), this model adopts the pointer mechanism to deal with the issue of Out-Of-Vocabulary in the summary generation process.

**Transformer:** Proposed by Vaswani et al. (2017), this model captures long-distance information through the self-attention mechanism.

**BertAbs:** Proposed by Liu and Lapata (2019), this model is abstractive, which encoder is initialized with BERT and is trained with a different optimizer than decoder.

**BART:** Proposed by Lewis et al. (2020), the model is a state-of-the-art abstractive summarization model pre-trained with a denoising autoencoding objective.

**BART w. DS/QR:** Proposed by Zhao et al. (2021b), the model is a general summarization framework, with two encoders that share the underlying parameters and a decoder, which can fuse the input text and dialogue state/query.

**Prefix-tuning:** Proposed by Li and Liang (2021), this model introduces a prefix matrix on the basis of fixed pre-training BART parameters and allows the prefix matrix to learn task information through training, which optimizes the summarization performance in the small parameters and few-shot scenarios.

**QMSum:** This model is proposed by Zhong et al. (2021), which is a two-stage locate-then-summarize solution on query-based meeting summarization task.

## A.2 Evaluation Metrics

We use the ROUGE (Lin, 2004; Lin and Och, 2004)<sup>6</sup> metrics to quantitatively evaluate the performance of our model. Rouge (Recall-Oriented Understudy for Gisting Evaluation) is metrics for evaluating summarization. It calculates by comparing the generated summary with references to

<sup>6</sup><https://pypi.org/project/rouge/>

Models	Trainable Parameters
BertExt	120.51M
PGN	27.64M
Transformer	257M
BertAbs	180.22M
BART-large	400M
BART w. DS/QR	406M
Prefix-tuning	81.82M
DOP-tuning	61.52M

Table 9: Trainable parameter scales of different models, where "DS" means the dialogue state in TODSum and "QR" means the query in QMSum.

obtain the corresponding score to measure the similarity between them.

## B Parameter Scale of Models

We show the amount of **trainable** parameters for our DOP model and other baseline models in Table 9. Among the full-parameter fine-tuning methods, except for the relatively simple PGN model, other models have reached the scale of hundreds of megabytes, which will take up a lot of time and space in model training and storage. Prefix-tuning and DOP-tuning greatly reduce the storage space of the model, improving the efficiency of the model. Compared with prefix-tuning, our method achieves better results with fewer parameters.

## C Case Study

Figure 7 shows two examples from the TODSum and QMSum respectively. For example one of *train* domain in TODSum, BART w. DS generates some incorrect and redundant information related to the *taxi* domain and *hotel* domain. To make matters worse, for *train* domain, it loses the intent of the user about booking tickets and wrongly generates the key information, i.e., the departure location. The Prefix-tuning still confuses the knowledge of *train* and *taxi* domains, that is, the booking intent in the *train* domain is wrongly predicted as the user wants to know something about "cars". Moreover, the quality of the generated key information is not high, i.e., the wrong departure location "Seachage" and the missing time "Monday". For example two of *academic* domain in QMSum, both BART w. QR and Prefix-tuning predict too many details of the dialogue, which makes the summary redundant. Besides, Prefix-tuning generates the wrong speaker "Professor B", which leads to the summary being inconsistent.

Compared to the above two models, our method

solves some difficult issues in low-resource domain adaptation for dialogue summarization. By initializing prefix matrix with domain words, our model achieves domain disentanglement and the prediction of domain-related information is basically accurate. Through discrete prompts, our model has the ability to generalize to new domains and the accuracy of prediction about domain-independent key information is greatly improved.

## D Domain Words

We present some domain words for each domain in Figure 8. In order to facilitate reading, we only show some of the domain words, that is, we select the first 20 words for each domain of TODSum, and the first 30 words for each field of QMSum as display.

For TODSum, we can see that there are relatively many common domain words, which are more concentrated on location words, or some information, such as "phone", "postcode" and so on. In addition, there are many common domain words that only appear in *restaurant & hotel* or *train & taxi*. For example, price-related descriptions are usually mentioned when booking a restaurant or hotel, and "destination", "depart", "from", "to" are usually mentioned when booking train tickets or taking a taxi. Special domain words can better distinguish different domains. There will be more special domain words in *attraction*, such as "entrance", "college", "nightclub", etc., which will not appear in other domains. Besides, users will mention the star rating when booking a hotel, and want to know the food type when booking a restaurant seat. When booking train tickets, they usually plan to travel, and when taking a taxi, they want to know the color of the car.

Compared with TODSum, QMSum has many more special field words, because the three fields contained in QMSum are more different. For *product*, participants will discuss various features of products such as TVs, LCDs, etc., such as screens, buttons, colors, and functions. For *academic*, participants generally discuss models, experimental data, or some errors. And for *committee*, participants generally discuss student education or national government issues. Common domain words have only some general words, such as "different", "system", etc., and only a few special common domain words exist in *product & academic*, such as "design", "bit" and other technology-related words.

Example one: <i>Train</i> domain in TODSum	Example two: <i>Academic</i> domain in QMSum
<p>Dialogue:  U: I am looking to find a train. It needs to go to Cambridge and it has to leave on Monday.  A: I have 202 trains travelling to Cambridge on Monday, could you be more specific on when you'd like to depart and from where?  U: I'd like to depart from Stevenage after 19:45.  A: I have one train departing at 19:54, one at 21:54, and the last one at 23:54. Will any of those work?  U: Yes, 19:54 will be fine. I'd like tickets for 5 people.  A: Congratulations! Your booking was successful with a reference number of DMEL46A6 and a fee of 64 GBP. Can I help you with anything else?  U: That is all. Thank you.  A: You're welcome. Enjoy your day.</p>	<p>Dialogue:  PhD E: Mmhm . So is this, uh? Guenter was putting a bunch of Wall Street Journal data on our disks.  Professor B: That's it.  PhD E: So that 's the data that we 'll be running on?  Professor B: Yeah.  PhD E: I see. OK.  Professor B: Yeah. So we have the data, just not the recognizer. OK.  PhD E: So this test may take quite a while to run, then. May judging by the amount of data that he was putting on.  Professor B: Uh, well there 's training and test, right?  PhD E: I guess, I'm not sure.  Professor B: No, I mean, if it 's like the other things, there 's data for training the HMMs and data for testing it.  PhD E: OK. So there is.  Professor B: So, training the recognizer. But I think it 's trained on clean.  PhD D: The Wall Street?  Professor B: Yeah!</p>
<p>Dialogue State:  train book (people=5)  train inform (leaveAt=19:45 ; destination=cambridge ; day=monday ; departure=stevenage)</p>	<p>Query:  What did the team discuss about the Wall Street Journal data?</p>
<p><input type="checkbox"/> correct domain word    <input type="checkbox"/> wrong domain word</p>	<p><input type="checkbox"/> correct key word    <input type="checkbox"/> wrong key word    <input type="checkbox"/> redundancy</p>
<p>Ground Truth  The user wants to book <u>5 train tickets</u>. This <u>train</u> leaves at <u>19:45</u> on <u>Monday</u>, from <u>Stevenage</u> to <u>Cambridge</u>.</p>	<p>Ground Truth  <u>PhD E</u> informed the team that <u>Guenter</u> was putting the <u>Wall Street Journal data</u> on the team's <u>disks</u>. There was a lot of <u>data</u> so it would take <u>some time</u> to run the models.</p>
<p>BART w. DS  <del>The user wants to know the car type and the driver's phone number. The train leaves at 19:45, departure is from Stevenage to Cambridge, and destination is Cambridge. The user asks if there is a hotel that can accommodate 5 people for 5 days, and he plans to stay on Monday. The hotel is in the centre, and the food is moderate.</del></p>	<p>BART w. QR  <del>PhD E suggested that Guenter was putting a bunch of Wall Street Journal data on our disks, which would be the data for training the recognizer and for testing it. However, Professor B was not sure if there would be enough data for both the training and the test. He thought it would take quite a while to run the test, since the amount of data that he was putting on the disks.</del></p>
<p>Prefix-tuning  <del>The user asks the agent what the driver's phone number is, and the type of the car. The train leaves at 19:45, from Seachage to Cambridge.</del></p>	<p>Prefix-tuning  According to Professor B, this test may take quite a while to run, judging by the amount of data that Guenter was putting on. <del>There would be data for training the HMMs and data for testing it.</del></p>
<p>DOP  The user wants to book the <u>train</u> seat for <u>5</u> people at <u>19:45</u> on <u>Monday</u>. The <u>train</u> leaves from <u>Stevenage</u> to <u>Cambridge</u>.</p>	<p>DOP  According to <u>PhD E</u>, <u>Guenter</u> was putting a bunch of <u>Wall Street Journal data</u> on our <u>disks</u>, so that's the <u>data</u> that they 'll be running on. The amount of <u>data</u> is large and it would take <u>quite a while</u> to run.</p>

Figure 7: Case study for two examples from TODSum and QMSum datasets. We present the dialogue, its corresponding dialogue state/query, ground truth, BART w. DS prediction, Prefix-tuning prediction, and prediction of our DOP model.

orange	special domain words	green	common domain words	<span style="border: 1px solid black; padding: 0 2px;"> </span>	common domain words only in <i>product</i> and <i>academic</i>
<span style="border: 1px solid black; padding: 0 2px;"> </span>	common domain words only in <i>restaurant</i> and <i>hotel</i>	<span style="border: 1px solid black; padding: 0 2px;"> </span>	common domain words only in <i>train</i> and <i>taxi</i>		
<b>TODSum</b>					
<b>attraction</b>					
entrance, fee, postcode, museum, attraction, phone, college, visit, town, holiday, centre, entertainment, colleges, corner, church, architecture, pool, nightclub, address, art					
<b>hotel</b>					
hotel, book, free, stay, people, parking, price, <span style="border: 1px solid black; padding: 0 2px;">range</span> , table, centre, area, star, nights, reservation, guesthouse, wifi, <span style="border: 1px solid black; padding: 0 2px;">moderate</span> , <span style="border: 1px solid black; padding: 0 2px;">expensive</span> , <span style="border: 1px solid black; padding: 0 2px;">cheap</span> , north					
<b>restaurant</b>					
food, phone, address, serve, restaurant, price, <span style="border: 1px solid black; padding: 0 2px;">range</span> , south, <span style="border: 1px solid black; padding: 0 2px;">expensive</span> , centre, italian, <span style="border: 1px solid black; padding: 0 2px;">cheap</span> , postcode, indian, east, north, <span style="border: 1px solid black; padding: 0 2px;">moderate</span> , chinese, area, type					
<b>train</b>					
train, book, <span style="border: 1px solid black; padding: 0 2px;">destination</span> , table, time, cambridge, day, people, <span style="border: 1px solid black; padding: 0 2px;">depart</span> , reservation, <span style="border: 1px solid black; padding: 0 2px;">from</span> , <span style="border: 1px solid black; padding: 0 2px;">to</span> , looking, <span style="border: 1px solid black; padding: 0 2px;">leave</span> , centre, <span style="border: 1px solid black; padding: 0 2px;">by</span> travel, minutes, Stevenage, price					
<b>taxi</b>					
taxi, contact, <span style="border: 1px solid black; padding: 0 2px;">from</span> , <span style="border: 1px solid black; padding: 0 2px;">leave</span> <span style="border: 1px solid black; padding: 0 2px;">by</span> <span style="border: 1px solid black; padding: 0 2px;">arrive</span> , car, time, book, <span style="border: 1px solid black; padding: 0 2px;">destination</span> , <span style="border: 1px solid black; padding: 0 2px;">depart</span> , pick, when, thank, day, grey, type, white, black, yellow					
<b>QMSum</b>					
<b>product</b>					
remote, gap, control, button, <span style="border: 1px solid black; padding: 0 2px;">design</span> , TV, LCD, channel, screen, recognition, idea, different, television, scroll, speech, easy, point, market, battery, <span style="border: 1px solid black; padding: 0 2px;">bit</span> , new, functions, volume, colour, rubber, voice, product, chip, time, system					
<b>academic</b>					
phd, professor, pause, data, digits, system, train, noise, net, level, filter, language, neural, test, features, result, <span style="border: 1px solid black; padding: 0 2px;">bit</span> line, model, <span style="border: 1px solid black; padding: 0 2px;">design</span> , discourse, error, frames, server, run, file, subtraction, spectral, disk, bunch					
<b>committee</b>					
children, chair, government, Welsh, minister, school, support, wales, students, work, canada, health, services, local, education, sector, institutions, question, universities, time, care, funding, parents, authorities, across, staff, finch, public, curriculum, system					

Figure 8: Domain words in each domain of two datasets. We use different signs to mark the special domain words, common domain words, and common domain words in specific domains separately.