

Few-Shot Dialogue Summarization via Skeleton-Assisted Prompt Transfer in Prompt Tuning

Kaige Xie[♣] Tong Yu[†] Haoliang Wang^{†*} Junda Wu[◇]
Handong Zhao[†] Ruiyi Zhang[†] Kanak Mahadik[†] Ani Nenkova[†] Mark Riedl[♣]

[♣] School of Interactive Computing, Georgia Institute of Technology

[†] Adobe Research [◇] University of California San Diego

{kaigexie, riedl}@gatech.edu juw069@ucsd.edu

{tyu, hawang, hazhao, ruizhang, mahadik, nenkova}@adobe.com

Abstract

In real-world scenarios, labeled samples for dialogue summarization are usually limited (i.e., few-shot) due to high annotation costs for high-quality dialogue summaries. To efficiently learn from few-shot samples, previous works have utilized massive annotated data from other downstream tasks and then performed prompt transfer in prompt tuning so as to enable cross-task knowledge transfer. However, existing general-purpose prompt transfer techniques lack consideration for dialogue-specific information. In this paper, we focus on improving the prompt transfer from dialogue state tracking to dialogue summarization and propose Skeleton-Assisted Prompt Transfer (SAPT), which leverages skeleton generation as extra supervision that functions as a medium connecting the distinct source and target task and resulting in the model’s better consumption of dialogue state information. To automatically extract dialogue skeletons as supervised training data for skeleton generation, we design a novel approach with perturbation-based probes requiring neither annotation effort nor domain knowledge. Training the model on such skeletons can also help preserve model capability during prompt transfer. Our method significantly outperforms existing baselines. In-depth analyses demonstrate the effectiveness of our method in facilitating cross-task knowledge transfer in few-shot dialogue summarization.

1 Introduction

Automatic text summarization (Luhn, 1958) is one of the most important and challenging problems in NLP. Among the different forms the text to be summarized could take, dialogues have been serving as a critical part of human-human and human-machine interaction. There has been significant progress made in dialogue summarization these

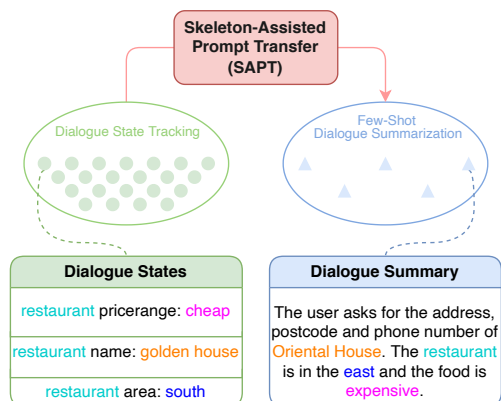


Figure 1: We study the problem of how to perform effective transfer learning from dialogue state tracking (DST) to few-shot dialogue summarization, in the scenario where there is a large set of dialogues with DST annotations, and another small set of dialogues with dialogue summarization annotations (i.e., few-shot learning for dialogue summarization).

days (Goo and Chen, 2018; Liu et al., 2019b; Chen and Yang, 2020). However, they generally rely on massive human-written golden dialogue summaries. In real-world scenarios, the availability of massive supervised data is not always guaranteed, as the data scarcity problem often occurs due to the high annotation cost that is normally required for acquiring large-scale high-quality dialogue summaries (Bražinskas et al., 2020).

In existing works, one common way to tackle the data scarcity problem is to perform transfer learning by leveraging off-the-shelf out-of-domain or out-of-task supervised data (Yang et al., 2020; Goodwin et al., 2020; Yu et al., 2021; Zou et al., 2021; Magooda et al., 2021). We observe that the supervised data of a relevant task called dialogue state tracking (DST) (Williams and Young, 2007) can bring conducive knowledge for the dialogue summarization task, as the semantic slots and values tracked by DST are expected to be covered in the dialogue summary (Shin et al., 2022). Besides the notable relevance between those two tasks, with

*Corresponding author

DST being a language understanding task as opposed to dialogue summarization being a language generation task, the annotations of DST should arguably be easier to get in practice than those of dialogue summarization.¹ These observations motivate us to herein focus on developing effective transfer learning techniques for the scenario where there are ample supervised data for DST whereas the annotations for dialogue summarization are limited, as depicted in Figure 1.

Among recent transfer learning techniques, prompt transfer (Vu et al., 2022) in prompt tuning (Li and Liang, 2021; Lester et al., 2021) has gained great popularity because of its parameter efficiency. Prompt tuning is a paradigm of utilizing pretrained language models (PLMs) for downstream tasks, in which a sequence of continuous trainable embeddings called “soft prompt” is prepended to the input sequence so as to provide PLMs with an adequate context. During training, only these embeddings can be updated while all the other parameters of PLMs will remain fixed. Prompt transfer realizes cross-task transfer learning under the prompt tuning paradigm by training soft prompts from source tasks and then using them as parameter initialization for the prompt tuning in target tasks. In general, prompt transfer works well in transfer learning between language understanding tasks while it can only provide relatively mediocre performance in language generation tasks (Su et al., 2022), indicating *the necessity to design task-specific prompt transfer approaches* for language generation tasks such as dialogue summarization.

How to improve prompt transfer *in a task-specific manner*? The existing general-purpose prompt transfer technique (Vu et al., 2022) relies solely on the source and target task supervision, suggesting the lack of an intermediate task-specific medium that could potentially better connect the distinct source and target task. Also, as the model capability of processing source task data is closely associated with the knowledge it has gained during the source task pretraining, it needs to be effectively preserved during the prompt transfer so as to facilitate the model in handling the target task.

In this paper, we propose a dialogue-specific prompt transfer technique, named **Skeleton-Assisted Prompt Transfer (SAPT)**. SAPT provides the model with extra supervision during its prompt transfer by training it to perform skeleton gener-

ation along the way. This extra supervision can essentially function as an *intermediate task-specific medium* that is beneficial for the knowledge transfer between the distinct source and target task. To get the supervised training data for skeleton generation, we design a novel automatic skeleton extraction approach that requires neither annotation effort nor domain knowledge. Specifically, we observe the model’s output variation to perturbation-based probes and extract the dialogue turns to which the model displays the highest sensitivity as skeletons. Training the model on such skeletons can also help *preserve model capability during prompt transfer*. The idea behind this is that we try to prevent the model from forgetting the dialogue-state-related knowledge it has learned during its pretraining on supervised DST data, since the model sensitivity to perturbation-based probes in the DST task intrinsically reflects the capability of processing dialogue state information it has developed. Experimental results and in-depth analyses with BART (Lewis et al., 2020) on two dialogue summarization benchmarks (Zhao et al., 2021b; Yuan and Yu, 2019) demonstrate the effectiveness of our method.

In summary, our main contributions are:

- We focus on improving the prompt transfer in prompt tuning from dialogue state tracking to few-shot dialogue summarization. To the best of our knowledge, SAPT is the first effective dialogue-specific prompt transfer technique.
- By training the model to perform skeleton generation during prompt transfer, SAPT provides extra supervision that essentially functions as an intermediate task-specific medium between the distinct source and target task, allowing the model to better consume the dialogue state information from the source task.
- To preserve model capability during prompt transfer, we design a novel approach that employs perturbation-based probes to automatically extract dialogue skeletons as supervised training data for skeleton generation, requiring neither annotation effort nor domain knowledge.

2 Preliminaries

2.1 Problem Definition

Abstractive dialogue summarization is typically formulated as a sequence-to-sequence problem (Nallapati et al., 2016). Given a dialogue history x , a

¹In Appendix A, we validate it via a data annotation study.

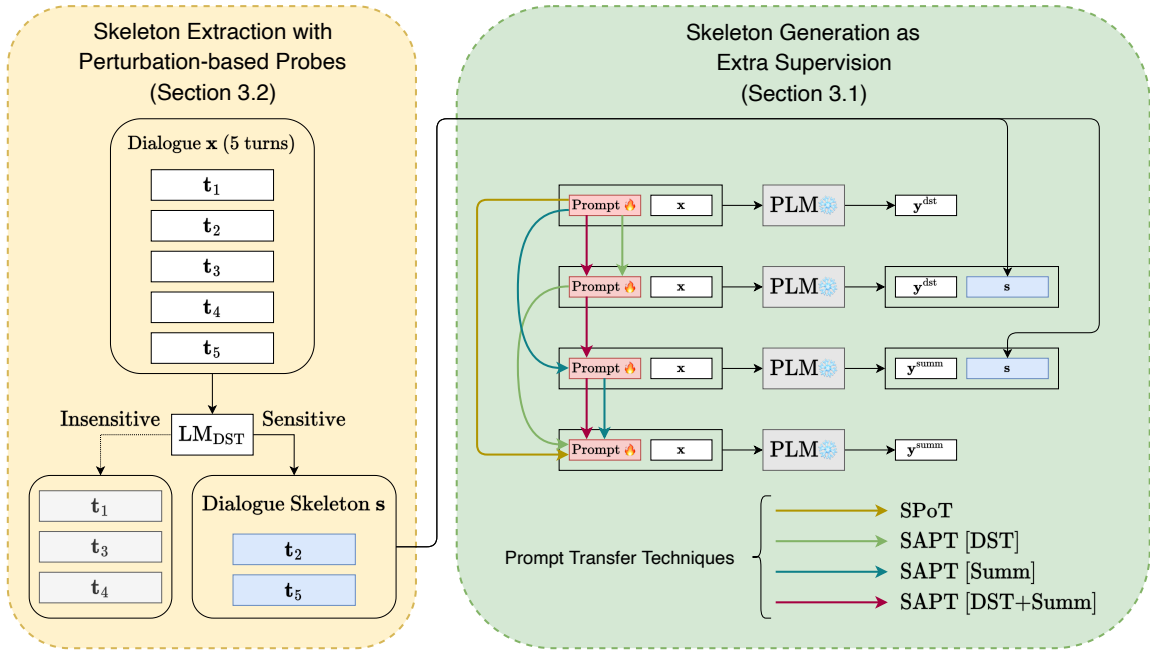


Figure 2: The overall workflow of Skeleton-Assisted Prompt Transfer (SAPT). Besides the original task supervision (\mathbf{y}^{dst} or \mathbf{y}^{summ}), SAPT uses skeleton generation as extra supervision (§3.1) by training on the dialogue skeletons \mathbf{s} extracted with perturbation-based probes (§3.2).

transformer-based encoder-decoder pretrained language model (PLM), $p_{\theta}(\mathbf{y}^{\text{summ}}|\mathbf{x})$, is trained to generate a summary \mathbf{y}^{summ} , where θ denotes the trainable parameters of the PLM. In this paper, we specifically study the dialogue summarization task in the few-shot setting (Bražinskas et al., 2020), meaning that there are only a limited number of annotated samples available for model training.

To mitigate the data scarcity problem, it is common to turn to transfer learning by leveraging massive supervised data from other related domains or tasks that could potentially provide useful knowledge. Dialogue state tracking (DST), a related task to dialogue summarization, aims to correctly infer the speaker’s goal in the form of semantic slot-value pairs ([slot, value]) as a dialogue progresses, such as [food, Italian] and [pricerange, high]. We thus notice that the supervised data of the DST task should be able to bring conducive knowledge for the dialogue summarization task, as the semantic slots and values tracked by DST are expected to be covered in the dialogue summary. Besides the notable relevance between those two tasks, with DST being a language understanding task as opposed to dialogue summarization being a language generation task, the annotations of DST should arguably be easier to get in practice, compared to those of dialogue summarization. Therefore, we herein focus on how

to perform effective transfer learning with ample supervised DST data to benefit the few-shot dialogue summarization.

Although the DST task is traditionally formulated as a classification problem, recent work (Lin et al., 2021; Zhao et al., 2021a) has shown the possibility of achieving competitive DST performance by treating DST as a sequence-to-sequence generation task. Specifically, conditioned on the dialogue history \mathbf{x} , the encoder-decoder model is trained to generate a sequence of tokens, in the format of “slot1 is value1, slot2 is value2, ...”, denoted as \mathbf{y}^{dst} . We thereby adopt this formulation for DST throughout our work so as to allow the generative encoder-decoder model’s knowledge transfer (from DST to dialogue summarization) to happen. With the unified generative sequence-to-sequence-based DST and dialogue summarization, the conditional generation task can be formulated as follows (\mathbf{y} can be either \mathbf{y}^{summ} or \mathbf{y}^{dst}):

$$P(\mathbf{y}|\mathbf{x}) = \prod_{i=1}^{|\mathbf{y}|} p_{\theta}(y_i|\mathbf{x}, y_{<i}).$$

2.2 Prompt Transfer in Prompt Tuning

Among recent transfer learning techniques, prompt transfer (Vu et al., 2022) in prompt tuning (Li and Liang, 2021; Lester et al., 2021) has gained great popularity because of its parameter efficiency. We

thus adopt it as our starting point for transfer learning from DST to dialogue summarization, and further improve it in section 3.

Prompt tuning is a new paradigm of utilizing PLMs for downstream tasks. It is motivated by the intuition that PLMs can be steered with a proper context, without the need for any model parameter updates. In prompt tuning, a sequence of continuous trainable embeddings called “soft prompt”, parameterized by ϕ , is prepended to the input sequence. During training, all parameters of the PLM (θ) are frozen, but unlike prompt design (Brown et al., 2020) which searches for actual tokens in the discrete space, prompt tuning optimizes the “soft prompt” (ϕ) directly in the continuous space, allowing it to be more expressive. The log-likelihood training objective can be formulated as follows:

$$\max_{\phi} \log p_{\theta, \phi}(\mathbf{y}|\mathbf{x}) = \sum_{i=1}^{|\mathbf{y}|} \log p_{\theta, \phi}(y_i|\mathbf{x}, y_{<i}).$$

Prompt transfer is a simple yet effective transfer learning technique designed for prompt tuning. The soft prompt is first trained in the source task and then used as parameter initialization for the prompt tuning in the target task. Prompt transfer inherits the advantage of prompt tuning in terms of parameter efficiency, as its transfer learning process likewise relies merely on the lightweight soft prompt. Su et al. (2022) show that prompt transfer generally works well in the transfer learning between language understanding tasks while it can only provide relatively mediocre performance in language generation tasks. This indicates the necessity to design task-specific prompt transfer approaches in language generation tasks such as dialogue summarization, which is exactly the central problem we focus on in this paper (detailed in section 3).

3 Method: Skeleton-Assisted Prompt Transfer (SAPT)

The existing non-task-specific general-purpose prompt transfer technique (Vu et al., 2022) relies solely on the source and target task supervision to train the soft prompt, without the help of any intermediate task-specific medium. Even though DST and dialogue summarization are closely related tasks, the intrinsic domain shift between them should still not be ignored. Therefore, having an intermediate task-specific medium should conceivably be helpful for better connecting the distinct source and target task. Such a medium can take

Algorithm 1 Skeleton Extraction with Perturbation-based Probes

Input:

a collection of dialogues \mathcal{X} containing N dialogues: $\mathcal{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\}$, where dialogue \mathbf{x}_i contains p_i dialogue turns: $\mathbf{x}_i = [\mathbf{t}_{i1}, \mathbf{t}_{i2}, \dots, \mathbf{t}_{ip_i}]$, $1 \leq i \leq N$;
a trained DST model LM_{DST} ;
a textual similarity metric $\text{Sim}(\cdot, \cdot)$ (higher means more similar).

Output:

a collection of dialogue skeletons: $\mathcal{S} = \{\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_N\}$, a subset $\text{set}(\mathbf{s}_i) \subseteq \text{set}(\mathbf{x}_i)$ for each dialogue $\mathbf{x}_i \in \mathcal{X}$, $1 \leq i \leq N$.

```

1:  $\mathcal{M} = \{\}$ 
2: for  $i = 1, 2, \dots, N$  do
3:    $\mathbf{o}_i = \text{LM}_{\text{DST}}(\mathbf{x}_i)$ 
4:   for  $j = 1, 2, \dots, p_i$  do
5:      $\mathbf{o}_{ij} = \text{LM}_{\text{DST}}(\mathbf{x}_i \setminus [\mathbf{t}_{ij}])$ 
6:      $m_{ij} = \text{Sim}(\mathbf{o}_i, \mathbf{o}_{ij})$ 
7:     add  $m_{ij}$  to  $\mathcal{M}$ 
8:  $\mathcal{S} = \{\}$ 
9:  $m_{\text{median}} = \text{Median}(\mathcal{M})$ 
10: for  $i = 1, 2, \dots, N$  do
11:    $\mathbf{s}_i = []$ 
12:   for  $j = 1, 2, \dots, p_i$  do
13:     if  $m_{ij} < m_{\text{median}}$  then
14:       append  $\mathbf{t}_{ij}$  to  $\mathbf{s}_i$ 
15:   add  $\mathbf{s}_i$  to  $\mathcal{S}$ 
16: return  $\mathcal{S}$ 

```

the form of extra task supervision separately incorporated into both the source and target task supervision, since in this way the updated source and target task have more overlap and get semantically closer to each other.

Also, as the model capability of processing source task data is closely associated with the knowledge it has gained during the source task pretraining, it needs to be effectively preserved during the prompt transfer to facilitate the target task. Nonetheless, the capability per se is admittedly a bit abstract and thus hard to concretely model in practice. Inspired by recent advances in interpretable NLP, we argue that the model sensitivity to perturbation-based probes should arguably be a concretization of model capability (Talmor et al., 2020). Thus, maintaining model sensitivity during the prompt transfer should logically benefit the preservation of model capability. And notably, the aforementioned extra task supervision can ex-

actively create conditions for (source-task) model-sensitivity information to be explicitly passed to the (target-task) model during the prompt transfer.

To these ends, we propose Skeleton-Assisted Prompt Transfer (SAPT), a dialogue-specific prompt transfer technique. SAPT provides the model with extra supervision during its prompt transfer by training it to perform skeleton generation along the way (detailed in subsection 3.1). This extra supervision (i.e. skeleton generation) is separately incorporated into both the source and target task supervision, and thus can essentially function as an *intermediate task-specific medium* (because of the increased overlap between the updated source and target task) that is beneficial for the cross-task knowledge transfer.

To get the supervised training data for skeleton generation, we design a novel automatic skeleton extraction approach that requires neither annotation effort nor domain knowledge (detailed in subsection 3.2). Specifically, we observe the model’s output variation to perturbation-based probes and extract the dialogue turns to which the model displays the highest sensitivity as skeletons. Training the model on such skeletons can also help *preserve model capability during prompt transfer*. This is because those skeletons (extracted with perturbation-based probes) embody the model sensitivity to perturbation-based probes which is a concretization of model capability.

On the whole, SAPT creates an intermediate task-specific medium using skeleton generation as extra supervision (§3.1), and preserves model capability during prompt transfer by training the model on the skeletons extracted with perturbation-based probes (§3.2). As a result, the distinct source and target task is able to be better connected because they have got semantically closer to each other, and the target task is able to be facilitated because the model has been discouraged from forgetting the knowledge it has gained during the source task pre-training. §3.3 describes SAPT’s overall workflow.

3.1 Skeleton Generation as Extra Supervision

In SAPT, the skeleton generation task is incorporated into the original task (either the source or the target task, or both) as extra supervision. We denote a supervised sample of the original task as (\mathbf{x}, \mathbf{y}) , where \mathbf{x} represents the dialogue history and \mathbf{y} represents the original task supervision that could be either the sequence-to-sequence-based dialogue state ground-truth or the dialogue summary

ground-truth. For each sample (\mathbf{x}, \mathbf{y}) , We also have a dialogue skeleton, denoted as \mathbf{s} , extracted from the dialogue history \mathbf{x} (the skeleton extraction algorithm is detailed in subsection 3.2). Such a dialogue skeleton is essentially an ordered collection of dialogue turns. For instance, if a dialogue history \mathbf{x} contains p dialogue turns, i.e. $\mathbf{x} = [\mathbf{t}_1, \mathbf{t}_2, \dots, \mathbf{t}_p]$, its dialogue skeleton \mathbf{s} will contain q dialogue turns ($q \leq p$), denoted as $\mathbf{s} = [\mathbf{t}_1^s, \mathbf{t}_2^s, \dots, \mathbf{t}_q^s]$, and thus $\text{set}(\mathbf{s}) \subseteq \text{set}(\mathbf{x})$. The dialogue skeleton \mathbf{s} is appended to the original task supervision \mathbf{y} as extra supervision, and the model is trained to perform the original task and then skeleton generation. The new log-likelihood training objective is:

$$\begin{aligned} & \max_{\theta, \phi} \log p_{\theta, \phi}(\mathbf{y} \oplus \mathbf{s} \mid \mathbf{x}) \\ &= \log p_{\theta, \phi}(\mathbf{y} \mid \mathbf{x}) + \log p_{\theta, \phi}(\mathbf{s} \mid \mathbf{x}, \mathbf{y}) \\ &= \log p_{\theta, \phi}(\mathbf{y} \mid \mathbf{x}) + \sum_{i=1}^q \log p_{\theta, \phi}(\mathbf{t}_i^s \mid \mathbf{x}, \mathbf{y}, \mathbf{t}_{<i}^s). \end{aligned}$$

3.2 Skeleton Extraction with Perturbation-based Probes

We extract dialogue skeletons (used as supervised training data for skeleton generation in subsection 3.1) with perturbation-based probes. Given a dialogue in a collection of dialogues, $\mathbf{x}_i \in \mathcal{X}$, we first construct the perturbation-based probes by deleting a dialogue turn from \mathbf{x}_i at a time. The resultant perturbation-based probes can be expressed as $\mathbf{x}_i \setminus [\mathbf{t}_{ij}], 1 \leq j \leq p_i$ (\mathbf{x}_i contains p_i dialogue turns). We then feed those perturbation-based probes individually into the trained source-task (DST) model, LM_{DST} , and get the model output \mathbf{o}_{ij} corresponding to each deleted dialogue turn \mathbf{t}_{ij} . In the meantime, we also feed the whole dialogue history \mathbf{x}_i into LM_{DST} and get the model output \mathbf{o}_i . Next, we compute the textual similarity score m_{ij} between \mathbf{o}_i and \mathbf{o}_{ij} using a textual similarity metric $\text{Sim}(\cdot, \cdot)$ (higher means more similar). We execute the aforementioned procedure for each dialogue in \mathcal{X} . After that, we group together all the similarity scores we compute along the way and find the median of them. Finally, we extract those dialogue turns, whose corresponding similarity scores are less than the median, as the dialogue skeletons. Algorithm 1 presents the process of extracting a dialogue skeleton \mathbf{s}_i for each dialogue $\mathbf{x}_i \in \mathcal{X}$.

3.3 Overall Workflow

Built on top of SPOT (Vu et al., 2022) while following the new training objective derived in subsection 3.1, SAPT uses skeleton generation as extra supervision by training on the dialogue skeletons extracted in subsection 3.2. Skeleton generation (as extra supervision) can be separately incorporated into either the source (DST) or the target (dialogue summarization) task supervision, or both. We thereby propose three SAPT variants: SAPT [DST], SAPT [SUMM], and SAPT [DST+SUMM].

As depicted in Figure 2, SAPT [DST+SUMM] includes four steps:

1. perform prompt tuning on the DST (source task) supervision;
2. perform prompt transfer from the previous step, and then perform prompt tuning on the DST (source task) & skeleton generation supervision;
3. perform prompt transfer from the previous step, and then perform prompt tuning on the (few-shot) dialogue summarization (target task) & skeleton generation supervision;
4. perform prompt transfer from the previous step, and then perform prompt tuning on the (few-shot) dialogue summarization (target task) supervision.

Compared to SAPT [DST+SUMM], SAPT [DST] omits step #3 while SAPT [SUMM] omits step #2; SPOT (Vu et al., 2022) omits both step #2 and step #3.

4 Experiment

4.1 Dataset and Baseline

To study the cross-task prompt transfer from dialogue state tracking (DST) to few-shot dialogue summarization, we perform experiments on a DST dataset: MultiWOZ 2.2 (Zang et al., 2020), and on two task-oriented dialogue summarization datasets: TODSUM (Zhao et al., 2021b) and SPNET (Yuan and Yu, 2019). MultiWOZ 2.2 is an error-fixed version of MultiWOZ (Budzianowski et al., 2018), which is a classic task-oriented multi-domain dialogue dataset containing over 10,000 annotated dialogues and has been extensively used for studying DST. TODSUM and SPNET are both constructed using the dialogues from MultiWOZ, and differ mainly in terms of summary style and length. On average, the summaries in SPNET are roughly two

times longer than those in TODSUM (96.4 vs. 45.4 words). To evaluate our method under the few-shot setting, on each dialogue summarization dataset we randomly choose 100 samples from the training set for model training and test on the full test set.

We use BART-large² (Lewis et al., 2020) as the backbone throughout the experiments. We focus on the comparison between prompt-tuning-based methods, as they have been proven to be able to maintain as comparable performance as the adapter-based methods while being much more parameter-efficient (Li and Liang, 2021; Vu et al., 2022). We choose SPOT (Vu et al., 2022) as the baseline method, which has been commonly used as a parameter-efficient transfer learning technique. Appendix B presents the implementation details.

4.2 Automatic Evaluation

We use the widely-used ROUGE metrics (Lin, 2004) as automatic evaluation metrics, including ROUGE-1 (R-1), ROUGE-2 (R-2), and ROUGE-L (R-L) F1 scores with rouge-score python package³. Few-shot (100-shot) results are presented in Table 1, where we also attach the results of PROMPT TUNING (Li and Liang, 2021). Unsurprisingly, PROMPT TUNING performs badly without any knowledge transfer, which indicates the necessity of conducting prompt transfer from DST to few-shot dialogue summarization. Among different prompt transfer techniques, all three SAPT variants outperform the baseline method SPOT on both datasets, suggesting the effectiveness of the proposed SAPT method. It is also observed that SAPT[DST] consistently outperforms SAPT[SUMM]. We attribute this to the fact that there are much more dialogue samples (along with their extracted dialogue skeletons) that are used for the preservation of model capability during step #2 than step #3, as in step #3 we only make use of 100 dialogue samples dedicated for few-shot dialogue summarization. Notably, when both step #2 and step #3 are executed, SAPT[DST+SUMM] is able to further improve the performance by a significant margin, compared to SAPT[DST] and SAPT[SUMM]. This demonstrates the effectiveness of creating an intermediate task-specific medium between the source DST task and the target few-shot dialogue summarization task (by incorporating the skeleton generation task into both

²https://huggingface.co/docs/transformers/model_doc/bart

³<https://pypi.org/project/rouge-score/>

Models	TODSUM			SPNET		
	R-1	R-2	R-L	R-1	R-2	R-L
PROMPT TUNING (Lester et al., 2021)	18.67	2.85	13.33	33.29	11.24	19.32
SPoT (Vu et al., 2022)	56.96	30.26	38.40	45.46	33.27	39.49
SAPT [DST]	62.00	36.95	43.13	53.43	40.07	44.92
SAPT [SUMM]	57.39	34.60	42.50	49.65	37.30	42.57
SAPT [DST+SUMM]	62.25	40.75	48.30	56.49	41.93	47.46

Table 1: Few-shot (100-shot) results on the full TODSUM (Zhao et al., 2021b) and SPNET (Yuan and Yu, 2019) test set. All three SAPT variants outperform the baseline model on both datasets, SPoT (Vu et al., 2022). SAPT [DST+SUMM] achieves the highest ROUGE scores with **significant** performance improvements.

	Informativeness	Faithfulness	Fluency	Redundancy
Ground Truth	1.92	1.90	1.95	1.97
SPoT (Vu et al., 2022)	1.77	1.70	1.73	1.71
SAPT [DST]	1.82	1.77	1.85	1.79
SAPT [SUMM]	1.80	1.76	1.85	1.73
SAPT [DST+SUMM]	1.86	1.82	1.90	1.81

Table 2: Human evaluation results in terms of the informativeness, faithfulness, fluency, and redundancy of the generated summaries on TODSUM test set. SAPT [DST+SUMM] consistently performs the best across all metrics.

skeleton type	R-1	R-2	R-L	decoding order	R-1	R-2	R-L	source & target task supervision	R-1	R-2	R-L
SAPT [DST]											
random skeleton	57.89	37.04	42.47	prepended skeleton	58.85	34.17	42.80	w/o source task supervision	59.61	32.26	41.81
our skeleton	62.00	36.95	43.13	appended skeleton	62.00	36.95	43.13	w/ source task supervision	62.00	36.95	43.13
SAPT [SUMM]											
random skeleton	55.14	33.07	41.70	prepended skeleton	53.33	31.16	39.34	w/o target task supervision	57.81	32.03	41.02
our skeleton	57.39	34.60	42.50	appended skeleton	57.39	34.60	42.50	w/ target task supervision	57.39	34.60	42.50
SAPT [DST+SUMM]											
random skeleton	58.34	38.98	43.11	prepended skeleton	62.04	40.92	46.74	w/o source & target task supervision	59.85	32.46	41.62
our skeleton	62.25	40.75	48.30	appended skeleton	62.25	40.75	48.30	w/ source & target task supervision	62.25	40.75	48.30

Table 3: Results of ablation studies on the effect of skeleton type, decoding order, and source & target task supervision for all three SAPT variants on TODSUM test set.

of them).

4.3 Human Evaluation

To further evaluate the generated summaries, we perform a human evaluation via crowdsourcing. We randomly select 100 samples from TODSUM test set and run different models on them to generate summaries. We recruit human participants on Prolific⁴, a crowdsourcing platform, to rate the generated summaries (and also the ground-truth summaries) from 0 to 2 in terms of four evaluation metrics: informativeness, faithfulness, fluency, and redundancy⁵. Each summary instance is evaluated by 5 different human participants, and the inter-annotator agreement (IAA) score for each metric is 0.577, 0.635, 0.649, 0.591, with an average IAA of 0.613. Results shown by the average scores in Ta-

ble 2 are consistent with the automatic evaluation results: all three SAPT variants outperform the baseline method SPoT, and SAPT[DST+SUMM] consistently performs the best across all metrics. Meanwhile, all generated summaries are deemed to be worse than the ground-truth summaries, meaning that there is still room for these summarization models to be improved. We also conduct a case study by ourselves, detailed in Appendix D.

4.4 Ablation Study

To fully investigate the effectiveness of SAPT, we study the impact of skeleton type, decoding order, and source & target task supervision. Table 3 shows the results of ablation studies.

Skeleton Type. We replace the extracted skeletons (§3.2) with randomly-extracted skeletons. We make sure that in total half of the dialogue turns are selected as skeletons to align with our usage

⁴<https://www.prolific.co/>

⁵Details of the metrics can be found in Appendix C.

of Median(), and that there is at least one dialogue turn selected for each dialogue. The observed performance drop demonstrates the effectiveness of our skeletons extracted with perturbation-based probes. Models with random skeleton still outperform SPOT in general, and we attribute this to the possible match between random skeleton and our skeleton, and also the imperfect intermediate task-specific medium which persists in the workflow.

Decoding Order. We prepend the skeletons instead of appending them. The observed performance drop demonstrates that the original task supervision needs to be prioritized, and prepending makes it more difficult for models to learn the cross-task knowledge.

Source & Target Task Supervision. We remove all the original task supervision along the way. The observed performance drop is as expected, but the superior performance against SPOT demonstrates the benefit our skeletons bring for cross-task knowledge transfer.

5 Related Work

Parameter-Efficient Transfer Learning. To efficiently make use of pretrained language models (PLMs) (Devlin et al., 2019; Raffel et al., 2020; Lewis et al., 2020; Brown et al., 2020), Li and Liang (2021) propose to prepend continuous trainable task-specific embeddings to the input sequence while keeping the entire PLM frozen. Lester et al. (2021) provide a simplified approach, named prompt tuning, which becomes more competitive with model fine-tuning as scale increases. To enable cross-task knowledge transfer (Ruder, 2017; Liu et al., 2019a) under the prompt tuning paradigm, Vu et al. (2022) propose SPoT, which learns soft prompts from source tasks as initialization for target tasks. Su et al. (2022) further explore the transferability of soft prompts across different downstream tasks. Built on top of Vu et al. (2022), our method is able to improve the effectiveness of cross-task prompt transfer in few-shot dialogue summarization.

Low-Resource Abstractive Summarization. Multiple lines of approaches have been proposed to mitigate the data scarcity problem in abstractive summarization, such as reinforcement learning (Kohita et al., 2020; Hyun et al., 2022), self-supervised learning (Fu et al., 2021; Wang and Wan, 2021; Zhuang et al., 2022), data augmentation (Amplayo and Lapata, 2020; Laskar et al.,

2020; Fabbri et al., 2021; Chen and Yang, 2021), model pretraining or fine-tuning with in-domain unlabeled data or out-of-domain labeled data (Yang et al., 2020; Goodwin et al., 2020; Yu et al., 2021; Zou et al., 2021; Magooda et al., 2021), and few-shot learning via adapters (Brazinskas et al., 2020; Brazinskas et al., 2022) or prompt tuning (Zhao et al., 2022; Liu et al., 2022; Yuan et al., 2022). In this paper, we focus on the few-shot dialogue summarization and improve it by ameliorating cross-task prompt transfer in prompt tuning with cross-task labeled data.

Perturbation-based Probes. In interpretable NLP, while probes sometimes refer to algorithms or models aiming to extract information from continuous embeddings (Adi et al., 2017), they can also refer to textual inputs designed for acquiring model outputs that are either useful for downstream tasks (Petroni et al., 2019; Zhong et al., 2021) or informative for model interpretability (Goldberg, 2019; Bacon and Regier, 2019; Xie et al., 2022). Perturbation-based probes, which fall into the latter category, have gained popularity because of their simplicity and cost-efficiency. For instance, Sankar et al. (2019); Abdou et al. (2020); Ettinger (2020); Clouatre et al. (2022) investigate the sensitivity of neural language models to input perturbation; Richardson and Sabharwal (2020); Talmor et al. (2020); Bitton et al. (2021); Gupta et al. (2022) utilize perturbation to construct better NLP testbeds; In contrast, we leverage perturbation-based probes to automatically extract skeletons from dialogues.

6 Conclusion

We focus on improving the prompt transfer in prompt tuning from dialogue state tracking to few-shot dialogue summarization, and propose SAPT, a dialogue-specific prompt transfer technique, which uses skeleton generation as extra supervision by training the model on the dialogue skeletons extracted with perturbation-based probes. In this way, a beneficial intermediate task-specific medium is created between the source and target task, and the model capability is able to be better preserved during the prompt transfer, resulting in the model’s better consumption of dialogue state information from the source task. Significantly stronger empirical performance and in-depth analyses on two dialogue summarization benchmarks demonstrate the effectiveness of our method in few-shot dialogue summarization.

7 Limitations

Despite the strong performance achieved by SAPT, we use the pre-trained language model (PLM) as the backbone of our method. Therefore, we cannot go beyond the limitation of the maximum sequence length of the PLM. In fact, long-form language understanding and generation have been widely acknowledged as an open research question that needs much further investigation, which is beyond the scope of our paper.

8 Ethics & Broader Impacts

All datasets used in this work are public. We did not collect any personal information from our human participants nor did we present them with any harmful model outputs. Our dialogue summarization models face the same potential pitfalls as other contemporary language learning systems do, e.g. being prone to echoing the biases present in the dataset (Sheng et al., 2019).

References

- Mostafa Abdou, Vinit Ravishankar, Maria Barrett, Yonatan Belinkov, Desmond Elliott, and Anders Søgaard. 2020. [The sensitivity of language models and humans to Winograd schema perturbations](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7590–7604, Online. Association for Computational Linguistics.
- Yossi Adi, Einat Kermany, Yonatan Belinkov, Ofer Lavi, and Yoav Goldberg. 2017. [Fine-grained analysis of sentence embeddings using auxiliary prediction tasks](#). In *International Conference on Learning Representations*.
- Reinald Kim Amplayo and Mirella Lapata. 2020. [Unsupervised opinion summarization with noising and denoising](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1934–1945, Online. Association for Computational Linguistics.
- Geoff Bacon and Terry Regier. 2019. [Does bert agree? evaluating knowledge of structure dependence through agreement relations](#). *arXiv preprint arXiv:1908.09892*.
- Yonatan Bitton, Gabriel Stanovsky, Roy Schwartz, and Michael Elhadad. 2021. [Automatic generation of contrast sets from scene graphs: Probing the compositional consistency of GQA](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 94–105, Online. Association for Computational Linguistics.
- Arthur Bražiņskas, Mirella Lapata, and Ivan Titov. 2020. [Few-shot learning for opinion summarization](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4119–4135, Online. Association for Computational Linguistics.
- Arthur Bražiņskas, Ramesh Nallapati, Mohit Bansal, and Markus Dreyer. 2022. [Efficient few-shot fine-tuning for opinion summarization](#). In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 1509–1523, Seattle, United States. Association for Computational Linguistics.
- 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.
- 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.
- 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. [Simple conversational data augmentation for semi-supervised abstractive dialogue summarization](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6605–6616, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Louis Clouatre, Prasanna Parthasarathi, Amal Zouaq, and Sarath Chandar. 2022. [Detecting languages unintelligible to multilingual models through local structure probes](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 5375–5396, Abu Dhabi, United Arab Emirates. 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.
- Allyson Ettinger. 2020. What BERT is not: Lessons from a new suite of psycholinguistic diagnostics for language models. *Transactions of the Association for Computational Linguistics*, 8:34–48.
- 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.
- Xiyan Fu, Yating Zhang, Tianyi Wang, Xiaozhong Liu, Changlong Sun, and Zhenglu Yang. 2021. RepSum: Unsupervised dialogue summarization based on replacement strategy. 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 6042–6051, Online. Association for Computational Linguistics.
- Yoav Goldberg. 2019. Assessing bert’s syntactic abilities. *arXiv preprint arXiv:1901.05287*.
- Ana Valeria Gonzalez and Anders Søgaard. 2020. The reverse turing test for evaluating interpretability methods on unknown tasks. In *NeurIPS Workshop on Human And Machine in-the-Loop Evaluation and Learning Strategies*, volume 61, page 62.
- Chih-Wen Goo and Yun-Nung Chen. 2018. Abstractive dialogue summarization with sentence-gated modeling optimized by dialogue acts. In *2018 IEEE Spoken Language Technology Workshop (SLT)*, pages 735–742. IEEE.
- Travis Goodwin, Max Savery, and Dina Demner-Fushman. 2020. Towards Zero-Shot Conditional Summarization with Adaptive Multi-Task Fine-Tuning. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3215–3226, Online. Association for Computational Linguistics.
- Vivek Gupta, Riyaz A. Bhat, Atreya Ghosal, Manish Shrivastava, Maneesh Singh, and Vivek Srikumar. 2022. Is my model using the right evidence? systematic probes for examining evidence-based tabular reasoning. *Transactions of the Association for Computational Linguistics*, 10:659–679.
- Dongmin Hyun, Xiting Wang, Chayoung Park, Xing Xie, and Hwanjo Yu. 2022. Generating multiple-length summaries via reinforcement learning for unsupervised sentence summarization. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 2939–2951, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Ryosuke Kohita, Akifumi Wachi, Yang Zhao, and Ryuki Tachibana. 2020. Q-learning with language model for edit-based unsupervised summarization. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 470–484, Online. Association for Computational Linguistics.
- Md Tahmid Rahman Laskar, Enamul Hoque, and Jimmy Xiangji Huang. 2020. WSL-DS: Weakly supervised learning with distant supervision for query focused multi-document abstractive summarization. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5647–5654, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. The power of scale for parameter-efficient prompt tuning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3045–3059, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- 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 *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Xiang Lisa Li and Percy Liang. 2021. Prefix-tuning: Optimizing continuous prompts for generation. 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 4582–4597, Online. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Zhaojiang Lin, Bing Liu, Seungwhan Moon, Paul Crook, Zhenpeng Zhou, Zhiguang Wang, Zhou Yu, Andrea Madotto, Eunjoon Cho, and Rajen Subba. 2021. Leveraging slot descriptions for zero-shot cross-domain dialogue StateTracking. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5640–5648, Online. Association for Computational Linguistics.
- Xiaochen Liu, Yang Gao, Yu Bai, Jiawei Li, Yanan Hu, Heyan Huang, and Boxing Chen. 2022. PSP:

- Pre-trained soft prompts for few-shot abstractive summarization. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 6355–6368, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Xiaodong Liu, Pengcheng He, Weizhu Chen, and Jianfeng Gao. 2019a. Multi-task deep neural networks for natural language understanding. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4487–4496, Florence, Italy. Association for Computational Linguistics.
- Zhengyuan Liu, Angela Ng, Sheldon Lee, Ai Ti Aw, and Nancy F Chen. 2019b. Topic-aware pointer-generator networks for summarizing spoken conversations. In *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 814–821. IEEE.
- Ilya Loshchilov and Frank Hutter. 2019. Decoupled weight decay regularization. In *International Conference on Learning Representations*.
- Hans Peter Luhn. 1958. The automatic creation of literature abstracts. *IBM Journal of research and development*, 2(2):159–165.
- Ahmed Magooda, Diane Litman, and Mohamed Elaraby. 2021. Exploring multitask learning for low-resource abstractive summarization. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 1652–1661, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Ramesh Nallapati, Bowen Zhou, Cicero dos Santos, Caglar Gulcehre, and Bing Xiang. 2016. Abstractive text summarization using sequence-to-sequence RNNs and beyond. In *Proceedings of the 20th SIGNLL Conference on Computational Natural Language Learning*, pages 280–290, Berlin, Germany. Association for Computational Linguistics.
- Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. Language models as knowledge bases? 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 2463–2473, Hong Kong, China. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140):1–67.
- Kyle Richardson and Ashish Sabharwal. 2020. What does my QA model know? devising controlled probes using expert knowledge. *Transactions of the Association for Computational Linguistics*, 8:572–588.
- Sebastian Ruder. 2017. An overview of multi-task learning in deep neural networks. *arXiv preprint arXiv:1706.05098*.
- Chinnadhurai Sankar, Sandeep Subramanian, Chris Pal, Sarath Chandar, and Yoshua Bengio. 2019. Do neural dialog systems use the conversation history effectively? an empirical study. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 32–37, Florence, Italy. Association for Computational Linguistics.
- Emily Sheng, Kai-Wei Chang, Premkumar Natarajan, and Nanyun Peng. 2019. The woman worked as a babysitter: On biases in language generation. 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 3407–3412, Hong Kong, China. Association for Computational Linguistics.
- Jamin Shin, Hangeul Yu, Hyeongdon Moon, Andrea Madotto, and Juneyoung Park. 2022. Dialogue summaries as dialogue states (DS2), template-guided summarization for few-shot dialogue state tracking. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 3824–3846, Dublin, Ireland. Association for Computational Linguistics.
- Yusheng Su, Xiaozhi Wang, Yujia Qin, Chi-Min Chan, Yankai Lin, Huadong Wang, Kaiyue Wen, Zhiyuan Liu, Peng Li, Juanzi Li, Lei Hou, Maosong Sun, and Jie Zhou. 2022. On transferability of prompt tuning for natural language processing. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3949–3969, Seattle, United States. Association for Computational Linguistics.
- Alon Talmor, Yanai Elazar, Yoav Goldberg, and Jonathan Berant. 2020. oLMpics-on what language model pre-training captures. *Transactions of the Association for Computational Linguistics*, 8:743–758.
- Tu Vu, Brian Lester, Noah Constant, Rami Al-Rfou’, and Daniel Cer. 2022. SPoT: Better frozen model adaptation through soft prompt transfer. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5039–5059, Dublin, Ireland. Association for Computational Linguistics.
- Ke Wang and Xiaojun Wan. 2021. TransSum: Translating aspect and sentiment embeddings for self-supervised opinion summarization. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 729–742, Online. Association for Computational Linguistics.
- Jason D Williams and Steve Young. 2007. Partially observable markov decision processes for spoken dialog systems. *Computer Speech & Language*, 21(2):393–422.

- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. **Transformers: State-of-the-art natural language processing**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Kaige Xie, Sarah Wiegrefe, and Mark Riedl. 2022. **Calibrating trust of multi-hop question answering systems with decompositional probes**. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 2888–2902, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- 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.
- Lin Yuan and Zhou Yu. 2019. **Abstractive dialog summarization with semantic scaffolds**. *arXiv preprint arXiv:1910.00825*.
- Ruifeng Yuan, Zili Wang, Ziqiang Cao, and Wenjie Li. 2022. **Few-shot query-focused summarization with prefix-merging**. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 3704–3714, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Xiaoxue Zang, Abhinav Rastogi, Srinivas Sunkara, Raghav Gupta, Jianguo Zhang, and Jindong Chen. 2020. **MultiWOZ 2.2 : A dialogue dataset with additional annotation corrections and state tracking baselines**. In *Proceedings of the 2nd Workshop on Natural Language Processing for Conversational AI*, pages 109–117, Online. Association for Computational Linguistics.
- Jeffrey Zhao, Mahdis Mahdieh, Ye Zhang, Yuan Cao, and Yonghui Wu. 2021a. **Effective sequence-to-sequence dialogue state tracking**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7486–7493, Online and 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*.
- Lulu Zhao, Fujia Zheng, Weihao Zeng, Keqing He, Weiran Xu, Huixing Jiang, Wei Wu, and Yanan Wu. 2022. **Domain-oriented prefix-tuning: Towards efficient and generalizable fine-tuning for zero-shot dialogue summarization**. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4848–4862, Seattle, United States. Association for Computational Linguistics.
- Zexuan Zhong, Dan Friedman, and Danqi Chen. 2021. **Factual probing is [MASK]: Learning vs. learning to recall**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5017–5033, Online. Association for Computational Linguistics.
- Haojie Zhuang, Wei Emma Zhang, Jian Yang, Congbo Ma, Yutong Qu, and Quan Z. Sheng. 2022. **Learning from the source document: Unsupervised abstractive summarization**. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 4194–4205, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Yicheng Zou, Bolin Zhu, Xingwu Hu, Tao Gui, and Qi Zhang. 2021. **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 Data Annotation Study

We recruit 30 human participants on Prolific⁶, a crowdsourcing platform, to annotate 30 dialogues for their dialogue states and dialogue summaries. We split 30 participants into two batches and split 30 dialogues into two batches as well. We follow a Latin Square design, similarly to (Gonzalez and Sjøgaard, 2020), to make sure that each batch of participants only sees each batch of dialogues in one of the following two annotation settings: dialogue state and dialogue summary, yet each setting is tested on both all 30 annotators and all 30 dialogues. This ensures that no bias in the duration of annotation occurs due to annotators having previously seen the dialogues.

We measure the duration of the annotation processes for both dialogue state and dialogue sum-

⁶<https://www.prolific.co/>

mary. The average duration of annotating a dialogue for its dialogue states is 1.3 minutes; the average duration of annotating a dialogue for its dialogue summary is 3.8 minutes, which is much longer. These results are in line with our intuition: the annotation of a dialogue summary requires not only tracking the dialogue states, but also having an utterance-level detailed understanding of the dialogue, because only after understanding the whole dialogue progression can annotators write a fluent and faithful summary.

B Implementation Details

We use Hugging Face Transformers⁷ (Wolf et al., 2020) during implementation. We train the BART-large models using AdamW (Loshchilov and Hutter, 2019) with the default learning rate linearly decaying from $5E - 5$. All models with a prompt length of 200 are trained for 50 epochs on an NVIDIA TITAN Xp GPU (12 GB memory) with a batch size of 2 and they each take approximately 25 hours (for DST) / 0.3 hours (for 100-shot dialogue summarization) to train. During inference, we perform a beam search with a beam size of 6, and the decoding takes 1.5 seconds per batch.

All turns of the input dialogue are prepended with special tokens as speaker identifiers ([USER] or [SYSTEM]), and then concatenated into a single input sequence which is truncated to 1024 BPE tokens. We use the ROUGE-L F1 score as the textual similarity metric $\text{Sim}(\cdot, \cdot)$ in Algorithm 1. The dialogue skeletons are appended to the ground-truth dialogue states (or summaries), and there is a special token [SEP] between the dialogue states (or summaries) and skeletons.

C Details of Human Evaluation Metrics

Human participants are asked to read the summaries and give their ratings (0, 1, or 2) in terms of four evaluation metrics:

- **Informativeness** examines whether the critical information in the dialogue is missed in the summary:
 - ★ 0: lots of the critical information in the dialogue is missed;
 - ★ 1: a small amount of the critical information in the dialogue is missed;
 - ★ 2: no critical information in the dialogue is missed.

- **Faithfulness** examines whether the information presented in the summary is factually incorrect or unmentioned according to the dialogue:

- ★ 0: lots of the information presented in the summary is factually incorrect or unmentioned;
- ★ 1: a small amount of the information presented in the summary is factually incorrect or unmentioned;
- ★ 2: no information presented in the summary is factually incorrect or unmentioned.

- **Fluency** examines whether the sentences in the summary are ungrammatical or ill-formed:

- ★ 0: lots of the sentences in the summary are ungrammatical or ill-formed;
- ★ 1: a small amount of the sentences in the summary are ungrammatical or ill-formed;
- ★ 2: no sentence in the summary is ungrammatical or ill-formed.

- **Redundancy** examines whether the expressions of the summary can be simplified:

- ★ 0: lots of the expressions of the summary can be simplified;
- ★ 1: a small amount of the expressions of the summary can be simplified;
- ★ 2: no expression of the summary can be simplified.

D Case Study

We present a case study in Table 4 to illustrate the effectiveness of SAPT.

⁷<https://github.com/huggingface/transformers>

Dialogue Example	
I am looking for a taxi from yu garden arriving by 14:30 .	<i>I'll need a destination to give them before I can book that for you.</i>
I'd be going to holy trinity church . I'll need the car type and contact number as well.	<i>The car that is reserved for you is a grey Toyota, and the contact number is 07997985529. Can I do anything else for you?</i>
Nope that's all thank you.	<i>Alright, enjoy your trip!</i>
SPoT (Vu et al., 2022)	
The user asks the user to book a taxi from yu garden arriving by 14:30 , and the destination is holy trinity church .	
SAPT [DST]	
The user asks for the car type and contact number of the taxi. The taxi leaves at 14:30 on thursday, and the destination is holy trinity church .	
SAPT [SUMM]	
The user wants to know the car type and the contact number of the taxi. The taxi leaves at 14:30 , and the destination is holy trinity church .	
SAPT [DST+SUMM]	
The user asks the agent to check the car type and contact number of the taxi. The taxi leaves at 14:30 on thursday, from yu garden to holy trinity church .	
Ground Truth	
The user wonders if it is possible to know the car type and the phone number . The taxi arrives at 14:30 , from yu garden to holy trinity church .	

Table 4: A case study. We highlight all dialogue-state-related information. The summaries provided by all three SAPT variants provide more complete dialogue-state-related information coverage than the baseline method SPOT. Among those three variants, only SAPT [DST+SUMM] covers all dialogue-state-related information. However, compared to the ground truth, the summary provided by SAPT [DST+SUMM] contains information that is unmentioned in the dialogue (i.e. on Thursday), suggesting there is still room for SAPT to be improved.