

# Mutual Reinforcement of LLM Dialogue Synthesis and Summarization Capabilities for Few-Shot Dialogue Summarization

Yen-Ju Lu<sup>†\*</sup>, Ting-Yao Hu<sup>\*</sup>, Hema Swetha Koppula, Hadi Pouransari,  
Jen-Hao Rick Chang, Yin Xia, Xiang Kong, Qi Zhu,  
Simon Wang, Oncel Tuzel, Raviteja Vemulapalli

<sup>†</sup>Johns Hopkins University    Apple

## Abstract

In this work, we propose Mutual Reinforcing Data Synthesis (MRDS) within LLMs to improve few-shot dialogue summarization task. Unlike prior methods that require external knowledge, we mutually reinforce the LLM’s dialogue synthesis and summarization capabilities, allowing them to complement each other during training and enhance overall performances. The dialogue synthesis capability is enhanced by directed preference optimization with preference scoring from summarization capability. The summarization capability is enhanced by the additional high quality dialogue-summary paired data produced by the dialogue synthesis capability. By leveraging the proposed MRDS mechanism, we elicit the internal knowledge of LLM in the format of synthetic data, and use it to augment the few-shot real training dataset. Empirical results demonstrate that our method improves dialogue summarization, achieving a 1.5% increase in ROUGE scores and a 0.3% improvement in BERT scores in few-shot settings. Furthermore, our method attains the highest average scores in human evaluations, surpassing both the pre-trained models and the baselines fine-tuned solely for summarization tasks.

## 1 Introduction

Dialogue summarization focuses on producing concise and coherent summaries of conversations in various domains such as customer service (Feigenblat et al., 2021; Zhao et al., 2021), medical consultations (Chintagunta et al., 2021; Jain et al., 2022; Zeng et al., 2020), and casual interactions (Chen et al., 2021; Gliwa et al., 2019). While state-of-the-art large language models (LLMs) such as Llama3 (Dubey et al., 2024) have been shown to work well for a wide variety of natural language processing tasks, their dialogue summarization performance is

\*Research conducted during internship at Apple. Corresponding authors: ylu125@jhu.edu, tingyao\_hu@apple.com

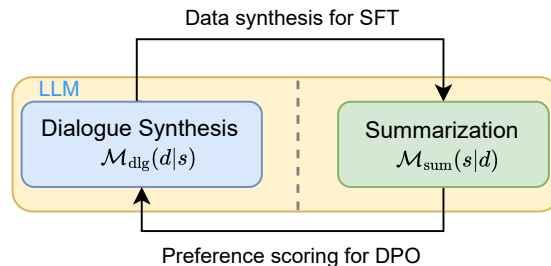


Figure 1: Mutual Reinforcement via Data Synthesis - We first leverage the summarization capability of the pretrained LLM for scoring dialogue preference pairs that are used to improve the dialogue synthesis capability with DPO, and then use the improved synthesis capability to generate SFT data for improving the summarization capability.

unsatisfactory in many target domains of interest. This could be because, these models, despite being trained on massive datasets, may have seen limited data related to dialogue summarization in specific target domains.

Collecting a large-scale, real-world dialogue dataset and manually annotating it with concise summaries is not only time-consuming and expensive, but it also raises privacy concerns in many target domains, as conversational data is often sensitive in nature. Consequently, there is a growing interest in developing few-shot learning approaches that improve the dialogue summarization capabilities of pretrained LLMs using limited real dialogue-summary pairs from the target domain.

Existing works that address data scarcity for dialogue summarization assume access to either additional dialogue-summary pairs from other domains (Li et al., 2023; Park et al., 2024; Yu et al., 2021; Zou et al., 2021; Zhong et al., 2022) or additional unlabeled dialogues from the target domain (Chen and Yang, 2021; He et al., 2024) or teacher models that are larger and more powerful than the target models (Ouyang et al., 2023; Pham et al., 2023). While these approaches have shown some perfor-

mance improvements, they need access to external resources that may not be available in all scenarios. Different from these works, we focus on improving the dialogue summarization capability of a pre-trained LLM using limited real dialogue-summary pairs from the target domain without relying on any additional data sources or external models.

LLMs possess a vast amount of implicit knowledge acquired during pre-training on large-scale text corpora, enabling them to generate contextually relevant text. Motivated by this, we introduce the framework of Mutual Reinforcement via Data Synthesis (MRDS) that harnesses the internal knowledge embedded within these models and their inherent data synthesis capability to address the data scarcity problem in dialogue summarization. Specifically, given a pretrained LLM, our method incorporates dialogue synthesis and summarization into a mutually reinforcing cycle to enhance both capabilities simultaneously (see Fig. 1).

To improve the dialogue synthesis capability, we create a dataset of synthetic dialogue preference pairs scored by leveraging the summarization capability of the pretrained LLM, and train a LoRA (Hu et al., 2022) adapter for dialogue synthesis using direct preference optimization (DPO) with the synthetic dialogue preference pairs in addition to supervised finetuning (SFT) with the limited real dialogue-summary pairs. This ensures that the generated dialogues are coherent and closely aligned with their corresponding summaries. To improve the summarization capability, we utilize the trained dialogue synthesis adapter to generate synthetic dialogue-summary pairs, and train a LoRA adapter for summarization by performing SFT with limited real data and generated synthetic data. Effectively, the initial summarization capability of the pretrained LLM helps improve the dialogue synthesis capability which in turn helps further improve the summarization capability.

### Major contributions:

- **Mutual reinforcement framework:** We introduce a framework that incorporates dialogue synthesis and summarization capabilities of an LLM into a mutually reinforcing cycle, and enhances both capabilities using limited target domain data.
- **DPO-Enhanced dialogue synthesis:** We propose to use direct preference optimization to improve the dialogue synthesis capability of a

pretrained LLM, ensuring that the generated dialogues are correctly formatted, coherent, and closely aligned with their summaries.

- **Experimental validation:** We demonstrate consistent improvements over various alternative approaches in terms of BERT and ROUGE scores on two widely-used benchmark datasets, namely SAMSum and DialogSum. The proposed approach also achieves the best average score in human evaluations. We also present several ablation results that demonstrate the effectiveness of individual components of the proposed approach.

## 2 Related Work

**Low-resource/Few-shot Dialogue Summarization.** Multiple lines of methods have been proposed to address the data sparsity problem of dialogue summarization. Yu et al. (2021) and Zou et al. (2021) employ datasets from multiple source domains to conduct pre-training. He et al. (2024) rely on semi-supervised learning techniques such as pseudo labeling to incorporate the additional dialogue data without summary annotation. Xie et al. (2024) and Zhao et al. (2022) design sophisticated prompt tuning strategies, enabling cross-task knowledge transfer. Recently, several methods have leveraged synthetic data generation using external knowledge or unlabeled datasets. For instance, GENDEX (Park et al., 2024) generates synthetic dialogues by utilizing external knowledge bases, thereby enriching the training data and improving model performance. Similarly, compositional data augmentation proposed by Ouyang et al. (2023) creates new training samples by recombining existing data, enhancing diversity without additional manual annotations. Additionally, Tian et al. (2024) employed a mixture-of-experts framework that integrates external knowledge to enhance summarization capabilities. While these approaches have demonstrated performance gains, they often depend on domain-specific resources, large unlabeled datasets, or external stronger models, which may not be available or practical in few-shot settings. This reliance on external data limits their applicability in scenarios where access to such resources is constrained.

**Synthetic Data from LLMs** Many previous works have shown that LLMs are capable of synthesizing high quality training data for machine

learning models. One line of methods primarily focus on zero-shot learning scenario (Ye et al., 2022a,b; Gao et al., 2023; Meng et al., 2022; Gupta et al., 2024), where they sample data from LLMs based on task related prompts, and use the synthetic data to train small, task-specific models from scratch. Other works also demonstrate the effectiveness of synthetic data from LLM in different domains such as speech recognition (Su et al., 2024), information extraction (Bonifacio et al., 2022; Josifoski et al., 2023), text-to-SQL (Yang et al., 2024), and dialogue state tracking (Kulkarni et al., 2024; Mehri et al., 2022). Recently, several works also investigate the idea of LLM self-improvement, suggesting that the synthetic data from LLMs can improve their own instruction following abilities. Self-instruct (Wang et al., 2023) samples from an LLM to create a synthetic prompt-response paired dataset, which can be used to finetune the original LLM. Li et al. (2024) introduce instruction backtranslation, which obtains synthetic instruction prompts from back-translating a web scale corpus with the same LLM. Gulcehre et al. (2023), Yuan et al. (2024), and Chen et al. (2024) pay attention to the generation of responses, but utilize them in different manners. Gulcehre et al. (2023) rely on an external scoring function to obtain the reward of synthetic responses. Yuan et al. (2024) propose a self-rewarding framework, using the LLM to score the response generated from itself. Chen et al. (2024) design a self-play mechanism, finetuning the LLM to distinguish the responses generated by the itself and human responses.

### 3 Methodology

**Notations** Let  $\mathcal{D}_p$  denote the limited real dataset of dialogue-summary pairs, and  $\mathcal{S}$  denote the set of summaries in  $\mathcal{D}_p$ .

Given a pretrained LLM, we characterize its dialogue synthesis and summarization capabilities using two separate LoRA (Hu et al., 2022) adapters. First, we leverage the summarization capability of the pretrained LLM to score synthetic dialogue preference pairs, and use this data to train the dialogue synthesis adapter using DPO in addition to SFT on real data  $\mathcal{D}_p$ . Then, we leverage the trained dialogue synthesis adapter to generate synthetic dialogue-summary pairs that are used (in addition to real data  $\mathcal{D}_p$ ) to train the summarization adapter. Effectively, the synthesis and summarization capabilities help improve each other without using any

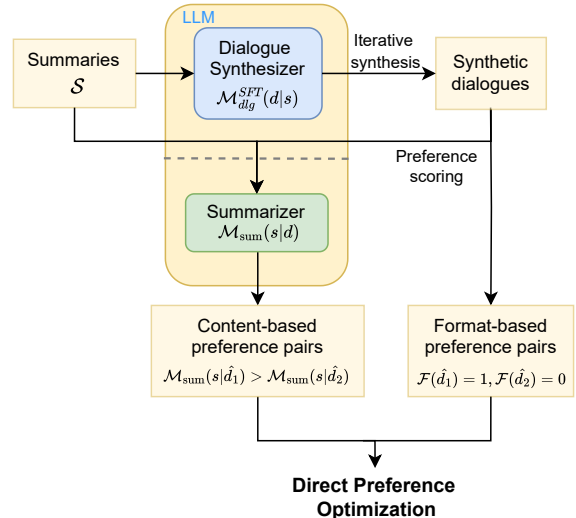


Figure 2: Preference pairs for DPO - Given a summary set  $\mathcal{S}$ , the LLM generates synthetic dialogues which are evaluated based on content alignment and format correctness. Content-based preferences leverage the LLM’s own summarization capability to assess how well the synthesized dialogues align with the input summaries. Format-based preferences ensure that the dialogues follow proper formatting.

external models or additional real data.

#### 3.1 Dialogue Synthesizer Training

The dialogue synthesizer (pretrained LLM with dialogue synthesis adapter) takes a summary  $s$  as input and generates a dialogue  $\hat{d}$ . A straightforward way to train this model is to perform SFT with the real dialogue-summary pairs  $\mathcal{D}_p$ . However, when the amount of such training data is limited, the resulting dialogue synthesizer  $\mathcal{M}_{dlg}^{SFT}$  often generates poor quality dialogues in terms of both dialogue format and content (see Table 8).

To encourage the dialogue synthesizer to generate higher quality dialogues, we construct a synthetic dataset of preference pairs  $\{s, \hat{d}_1, \hat{d}_2 \mid s \in \mathcal{S}, \mathcal{P}(\hat{d}_1) > \mathcal{P}(\hat{d}_2)\}$ , where  $\mathcal{P}$  denotes a dialogue quality scoring function, and train the synthesizer using DPO. Specifically, we generate two types of preference pairs using two different quality scoring functions, one focusing on the dialogue format and the other focusing on the dialogue content (see Fig. 2). Both the preferred and rejected dialogues in these pairs are synthetic dialogues generated using the SFT-trained dialogue synthesizer  $\mathcal{M}_{dlg}^{SFT}$ .

**Format-based Preference Pairs** Since the SFT-trained synthesizer  $\mathcal{M}_{dlg}^{SFT}$  often generates dialogues with several formatting errors, we develop

an iterative dialogue synthesis (IDS) method that generates correctly-formatted dialogues. First, we generate a dialogue conditioned on the input summary and check it for format errors. We discard the portion of the dialogue after the first detected error, concatenate the remaining correctly-formatted partial dialogue to the initial input prompt and give it as input to the dialogue synthesizer to complete the dialogue. This process is repeated until we get a dialogue without formatting errors. Table 11 shows an example run of this IDS process.

For a given summary, we generate multiple dialogues with format errors by directly sampling from the SFT-trained synthesizer  $\mathcal{M}_{dlg}^{SFT}$  and multiple clean dialogues by repeating the above IDS process several times. We use these samples to form preference pairs  $\{s, \hat{d}_1, \hat{d}_2 \mid s \in \mathcal{S}, \mathcal{F}(\hat{d}_1) = 1, \mathcal{F}(\hat{d}_2) = 0\}$ , where  $\mathcal{F}$  denotes the format check-based binary scoring function (1 for clean dialogues, 0 for those with errors). When trained with these preference pairs, the dialogue synthesizer learns to generate dialogues in correct format. Table 9 shows an example of format-based preference pair.

**Content-based Preference Pairs** A dialogue  $\hat{d}$  generated by a well-trained dialogue synthesizer should have high content alignment with the summary  $s$  used to generate  $\hat{d}$ . We leverage the summarization capability of the pretrained LLM to measure this alignment. The main motivation is that if we summarize the synthetic dialogue  $\hat{d}$ , the corresponding summary  $s$  should have high probability as the summarization output. Based on this, we use the likelihood  $\mathcal{M}_{sum}(s|\hat{d})$  of the summary  $s$  conditioned on the dialogue  $\hat{d}$  measured by the pretrained LLM as the alignment score.

To encourage the dialogue synthesizer to generate dialogues that have high content alignment with the input summaries, we construct a dataset of preference pairs using the content alignment score  $\mathcal{M}_{sum}(s|\hat{d})$ . Specifically, for each summary  $s$ , we first generate multiple clean dialogues from the SFT-trained dialogue synthesizer  $\mathcal{M}_{dlg}^{SFT}$  following the IDS process described above, and then pick the dialogues with the best and least content alignment scores to form preference pairs  $\{s, \hat{d}_1, \hat{d}_2 \mid s \in \mathcal{S}, \mathcal{M}_{sum}(s|\hat{d}_1) > \mathcal{M}_{sum}(s|\hat{d}_2)\}$ . Table 10 shows an example of content-based preference pair.

Instead of using the pretrained LLM, we also explored training the summarization adapter on the limited real data  $\mathcal{D}_p$  and using it for alignment scor-

ing to generate DPO training data. However, we did not observe significant improvements in the final summarization results (after training the dialogue synthesizer with DPO, generating synthetic dialogues, and using them to train the final summarization adapter).

**Training with DPO and SFT** In addition to DPO with synthetic preference pairs, we also use SFT with the limited real data  $\mathcal{D}_p$  to train the dialogue synthesizer. We accumulate gradients from both losses in each optimization step. While DPO with synthetic data encourages the model to generate contextually accurate dialogues in correct format, it does not explicitly encourage the model to generate dialogues that mimic the target distribution represented by real data  $\mathcal{D}_p$ . By combining DPO with SFT on real data, we encourage the model to generate dialogues closer to the target distribution while being contextually accurate and better-formatted.

### 3.2 Summarizer Training

The summarization model (pretrained LLM with summarization adapter) is trained using SFT with limited real data  $\mathcal{D}_p$  and additional synthetic dialogue-summary pairs.

**Synthetic Summary Generation** To generate synthetic summaries that mimic the distribution represented by the real summary set  $\mathcal{S}$ , we first extract a (2-3 words) topic for each summary  $s \in \mathcal{S}$  using the pretrained LLM. Then, for each topic, we generate multiple new synthetic summaries by using a topic-based summary synthesizer. This synthesizer is obtained by supervised LoRA finetuning of the pretrained LLM using the real summaries in  $\mathcal{S}$  and the corresponding extracted topics.

**Synthetic Dialogue Generation** For each synthetic summary, we generate a dialogue by directly sampling from the dialogue synthesizer that has been trained with both DPO and SFT, as explained in the previous section. Since this improved synthesizer already generates dialogues with high quality, we do not use the costly iterative synthesis approach in this step.

**Training with Synthetic and Real Data** Models trained only on limited real data tend to overfit quickly. While combining real and synthetic data samples in each minibatch could address this issue to some extent, finding the perfect ratio between the

two data sources is challenging and highly dependent on the quality of the synthetic data. Moreover, using a fixed ratio of real and synthetic data samples throughout the training may not be optimal. If the minibatch is dominated by synthetic data, then the model may inherit the artifacts present in the synthetic data, and if the minibatch is dominated by real data, the model may start to overfit before taking full advantage of the synthetic data.

To address these issues, we follow a two-stage training strategy, where we train only on synthetic data in the first stage and only on real data in the second stage. The synthetic-only first stage allows the model to learn general dialogue summarization skills without the risk of quickly overfitting on limited real data. The real-only second stage allows the model to adapt to the distribution of the real data mitigating the artifacts learned from synthetic data. Following this two-stage approach, we make effective use of both synthetic and real data, resulting in a more accurate summarization model.

## 4 Experiments

**Datasets** We experiment with two widely-used benchmark datasets, namely SAMSum (Gliwa et al., 2019) and DialogSum (Chen et al., 2021). The SAMSum dataset contains over 16,000 casual conversations mimicking everyday chats among friends and family. The DialogSum dataset includes about 13,460 face-to-face spoken dialogues between friends, colleagues, and between service providers and customers, covering various daily-life topics. These datasets also provide a human-written summary for each dialogue. Since this work focuses on few-shot settings, for each dataset, we experiment with either 100 or 300 dialogue-summary pairs as the few-shot training dataset  $\mathcal{D}_p$ .

### 4.1 Alternative Methods

We compare our MRDS method with several dialogue summarization alternative approaches, divided into two categories: methods using the pre-trained Llama3 model without fine-tuning and methods fine-tuned on real or synthetic data.

#### Pre-trained Methods

- **Zero-shot:** Zero-shot summarization performance of Llama3.
- **ICL:** Summarization performance of Llama3 using in-context learning with  $k = 7$  examples.<sup>1</sup>

<sup>1</sup>We experimented with different number of in-context

#### Fine-tuned Methods

- **Real only:** Fine-tuning with real data only.
- **SFT:** Two-stage training using synthetic dialogues generated by the SFT dialogue synthesizer.
- **SFT + Post-processing:** Two-stage training using synthetic dialogues from the SFT dialogue synthesizer, enhanced with Iterative Dialogue Synthesis (IDS) and content alignment filtering.

### 4.2 Implementation Details

We use Llama3-8B-Instruct (Dubey et al., 2024) as the pretrained base LLM. We use a rank of 16 and an alpha of 32 for all LoRA adapters, and keep the base model parameters frozen while training the LoRA adapters. Table 6 shows all the prompts used for topic extraction, topic-based summary synthesis, summary-based dialogue synthesis, and dialogue summarization tasks. All presented results are averaged over three runs.

**Dialogue Summarization** For the baseline model trained exclusively on real data, we optimized the hyperparameters and applied the same settings to all subsequent experiments for consistency. Our training strategy includes a batch size of 10 and a maximum learning rate of  $2.0 \times 10^{-4}$  with a warmup over the first 50 batches. We use the ReduceLRonPlateau scheduler with a patience of 5 and a reduction factor of 0.7. Training is stopped if the loss does not improve for 100 steps. We select the best checkpoint based on the validation loss obtained during the real data training phase.

In synthetic data experiments, we employ a two-stage training approach using the same hyperparameters. In the first phase, we train exclusively on synthetic data until the learning rate reduces to  $2.0 \times 10^{-5}$ , effectively serving as a pre-training phase. In the second phase, we apply the same training strategy as in the real-only experiments to ensure a fair comparison.

**Dialogue Synthesis** For the dialogue synthesizer trained with SFT only, we use a learning rate of  $2.0 \times 10^{-4}$  along with the ReduceLRonPlateau scheduler. The batch size and other hyperparameters are the same as those used for dialogue summarization. When training the synthesizer using both SFT and DPO, we start from the SFT checkpoint. In this combined training, we use a batch size of examples and  $k = 7$  worked best.

four for DPO and one for SFT, jointly updating the dialogue synthesizer by combining the losses from both objectives. A fixed learning rate of  $1 \times 10^{-5}$  is used during this phase. We validate the synthesizer checkpoints on the official validation set of the dataset, evaluating both format correctness and summarization cross-entropy loss. We select the checkpoint with the lowest summarization CE loss, ensuring at least 85% format correctness. Detailed training hyperparameters are provided in Table 13.

## 5 Results

### 5.1 Dialogue Summarization

We conducted experiments on the SAMSum and DialogSum datasets to evaluate the effectiveness of our proposed mutual reinforcing data synthesis (MRDS) method. The results are presented in Table 1, comparing various summarization approaches under 100-shot and 300-shot settings using metrics such as ROUGE-1 (R-1), ROUGE-2 (R-2), ROUGE-L (R-L) (Lin, 2004), and BERTScore (Zhang et al., 2020).

In the zero-shot setting, the pre-trained LLM (Zero shot) achieves R-1 scores of 31.3 on SAMSum and 28.2 on DialogSum. In-context learning approach (ICL) improves the R-1 score on SAMSum to 39.5 and DialogSum to 31.4, demonstrating efficiency of ICL in low-resource scenarios.

When fine-tuning with 100 real shots, the real-only method significantly improves performance over zero-shot methods, achieving R-1 scores of 50.9 on SAMSum and 44.0 on DialogSum. Incorporating dialogues from the SFT model maintains similar performance, while post-processing techniques (SFT + Post-processing) further enhance results, increasing R-1 scores to 51.8 on SAMSum and 44.7 on DialogSum, indicating the effectiveness of IDS and content alignment filtering. In the 300-shot setting, all methods benefit from additional training data. The real-only method reaches R-1 scores of 51.1 on SAMSum and 45.2 on DialogSum, with SFT showing incremental gains, and SFT + Post-processing achieving R-1 scores of 52.7 on SAMSum and 46.1 on DialogSum.

Our proposed MRDS method outperforms all approaches in both 100-shot and 300-shot settings. In the 100-shot setting, MRDS achieves the highest R-1 scores of 52.1 on SAMSum and 45.5 on DialogSum, along with improvements in R-2, R-L, and BERTScore metrics, highlighting its ability to leverage synthesized data effectively. In the 300-

shot setting, MRDS continues to deliver the best performance, matching the top R-1 score of 52.7 on SAMSum and setting a new high of 47.0 on DialogSum. The method consistently delivers superior ROUGE and BERTScore values, highlighting its robustness and scalability with increased data. This shows that MRDS not only improves efficiency by eliminating post-processing steps but also significantly boosts summarization performance.

### 5.2 Human Evaluation

Table 2 presents the results of human evaluations, comparing four groups of summaries: two from the pretrained instructed Llama3 model—zero-shot results (Zero) and in-context learning (ICL)—and two from the fine-tuned Llama3-based summarization model: trained on real data only (Real) and trained with the proposed MRDS approach, which combines real and synthetic summaries and dialogues. Five human evaluators assessed the summaries based on informativeness, faithfulness, fluency, and redundancy, using a scale from zero to two, following the evaluation protocol from (Xie et al., 2024). The average scores across the four metrics are also reported.

For the pre-trained models, while the zero-shot model (Zero) achieves the highest score in informativeness (1.95), it scores poorly in redundancy (0.83), indicating a tendency to produce overly long summaries by including too much information – undesirable in summarization tasks. The in-context learning model (ICL) shows slight improvements in faithfulness (1.69 vs. 1.50) and redundancy (1.23 vs. 0.83) compared to zero-shot, indicating that it generates more concise and faithful summaries but still inherits some limitations of the pre-trained model. For the fine-tuned models, the Real approach outperforms the pre-trained models in overall average score (1.72 vs. 1.56 for Zero and 1.64 for ICL), demonstrating the benefit of fine-tuning with real data in improving summary quality. However, our proposed MRDS method achieves the highest average score (1.84), outperforming both the real-only and pre-trained models, particularly in faithfulness and redundancy. This suggests that incorporating synthesized data helps the model produce more precise, concise summaries. The results highlight that our MRDS approach significantly enhances summarization quality from the perspective of human evaluators. Example summaries from different approaches are shown in Table 7.

Table 1: Comparison of Summarization Methods on 100 and 300 shots.

Approach	SAMSum				DialogSum			
	R-1	R-2	R-L	BERTScore	R-1	R-2	R-L	BERTScore
Zero shot	31.3	12.3	23.9	81.2	28.2	10.0	21.4	81.6
ICL (k=7)	39.5	17.6	30.9	83.2	31.4	11.9	24.5	83.1
<b>100 Real shots</b>								
Real only	50.9	26.5	42.6	86.6	44.0	18.2	36.0	86.8
SFT	50.9	26.5	42.6	86.5	45.1	19.0	36.9	86.9
SFT + Post-processing	51.8	27.3	<b>43.5</b>	86.7	44.7	18.8	36.5	87.1
MRDS (ours)	<b>52.1</b>	<b>27.5</b>	43.4	<b>86.8</b>	<b>45.5</b>	<b>19.3</b>	<b>37.2</b>	<b>87.2</b>
<b>300 Real shots</b>								
Real only	51.1	26.9	42.8	86.5	45.2	19.5	37.3	87.2
SFT	52.1	27.6	43.7	86.8	45.9	19.8	37.7	87.2
SFT + Post-processing	<b>52.7</b>	28.1	44.1	<b>87.0</b>	46.1	20.0	37.6	87.3
MRDS (ours)	<b>52.7</b>	<b>28.3</b>	<b>44.4</b>	<b>87.0</b>	<b>47.0</b>	<b>20.4</b>	<b>38.6</b>	<b>87.5</b>

Table 2: Human evaluation results on SAMSum with 300-shot data, covering informativeness (inf.), faithfulness (fai.), fluency (flu.), redundancy (red.), and the average (ave.) of the four scores.

	Inf.	Fai.	Flu.	Red.	Ave.
<b>Zero</b>	<b>1.95</b>	1.50	1.94	0.83	1.56
<b>ICL</b>	<u>1.65</u>	1.69	<b>2.00</b>	1.23	1.64
<b>Real</b>	1.32	1.61	<b>2.00</b>	1.94	1.72
<b>MRDS</b>	1.55	<b>1.88</b>	<u>1.94</u>	<b>2.00</b>	<b>1.84</b>

### 5.3 Dialogue Synthesis Efficiency

We compared the efficiency of the post-processing approach—which includes iterative synthesis and content alignment filtering—with the DPO-based MRDS approach (Table 3). In the 100-shot scenarios, although MRDS still required iterative synthesis due to less consistent format correctness with limited data, it achieved 550 dialogues per hour on SAMSum compared to 63 dialogues per hour with post-processing—an 8.7-fold improvement. In the 300-shot scenarios, MRDS significantly increased throughput, generating 2,900 dialogues per hour on SAMSum versus 37.5 dialogues per hour with post-processing—a 77-fold improvement—by eliminating the need for IDS or content alignment filtering due to enhanced format correctness and summarization alignment.

Table 3: Comparison of efficiency for dialogue synthesis with (Dialogues/hour).

Approach	100 shots	300 shots
<b>SAMSum</b>		
Post-Processing	63	37.5
MRDS	<b>550</b>	<b>2900</b>
<b>DialogSum</b>		
Post-Processing	85	66.7
MRDS	<b>315</b>	<b>2500</b>

### 5.4 Model Analysis

**Summaries Effectiveness.** To evaluate the impact of different summary types on dialogue synthesis, we experimented with two training strategies: (1) a fixed 1:1 ratio of real to synthesized data and (2) a two-stage training approach, training first on synthesized data and then switching to real data. We tested these strategies using three types of summaries: the same summaries (identical to those in the real data), synthetic summaries generated from our synthesis process in Sec. 3.2, and unseen summaries drawn from additional training data.

Table 4 presents the results. Under the fixed ratio strategy, using the same summaries did not yield any improvement over the real-only approach, likely due to overfitting to the limited real summaries. Synthetic summaries provided some improvement in the 100-shot setting but introduced

Table 4: Comparison of different summaries with synthesis data on DialogSum 100 and 300 real shots. Synthesis dialogues are generated by SFT-trained dialogue synthesizer with IDS and content alignment filtering.

Approach	100 Real shots				300 Real shots			
	R-1	R-2	R-L	BERTScore	R-1	R-2	R-L	BERTScore
Real only	44.0	18.2	36.0	86.8	45.2	19.5	37.3	87.2
<b>Fixed Ratio Training</b>								
Same Summ	42.8	17.4	35.3	86.4	44.6	19.0	36.7	86.7
Synthetic Summ	44.1	18.2	36.1	86.7	44.0	18.2	35.5	86.6
Unseen Summ	44.5	18.8	36.3	87.0	46.6	20.5	38.1	87.5
<b>Two-Stage Training</b>								
Same Summ	42.5	16.8	34.6	86.3	45.8	19.9	37.6	87.2
Synthetic Summ	44.7	18.8	36.5	87.1	46.1	20.0	37.6	87.3
Unseen Summ	45.3	19.7	37.4	87.0	46.8	20.8	38.8	87.6

artifacts that degraded performance in the 300-shot setting. This suggests that mixing synthetic and real data in a fixed ratio can lead to the model learning undesirable patterns from the synthetic data as the amount of real data increases.

In contrast, the two-stage training approach produced better results. Training with synthetic summaries first allowed the model to learn general patterns from the synthesized data before fine-tuning on real data, which improved performance in both the 100-shot and 300-shot settings. Specifically, using synthetic summaries in two-stage training significantly outperformed using the same summaries, mitigating overfitting and enhancing generalization. This highlights the value of synthetic summaries in the more effective two-stage training framework. Additionally, incorporating unseen summaries improved both training strategies, but the two-stage training still provided superior results.

**Dialogue Generation and Filtering.** To evaluate the benefits of generating and filtering dialogues, we experimented with unseen summaries using the two-stage training strategy. In the first stage, we tested three types of synthetic data: summaries only, summaries with unfiltered dialogues, and approaches involving content alignment filtering or DPO training (MRDS), as shown in Table 5. We found that training with summaries alone improved upon the real-only approach; however, adding dialogues to the training data further enhanced the results. Finally, the models utilizing filtering or DPO outperformed all other methods, demonstrating that including dialogues with filtering and DPO

Table 5: Ablation study for dialogue generation, filtering, and DPO on unseen summaries. The training is conducted in two stages training.

Approach	DialogSum 300 shots			
	R-1	R-2	R-L	B-S
Real only	45.2	19.5	37.3	87.2
<b>Unseen Summaries</b>				
w/o dialogue	45.8	19.9	37.3	87.2
w/o filtering	46.3	20.1	37.9	87.3
w/ filtering	<u>46.8</u>	<b>20.8</b>	<b>38.8</b>	<b>87.6</b>
MRDS	<b>47.0</b>	<u>20.7</u>	<u>38.7</u>	<u>87.5</u>

effectively improves the final outcomes.

## 6 Conclusion

We introduce a novel approach that mutually reinforces dialogue synthesis and summarization capabilities of a large language model (LLM) to improve few-shot dialogue summarization without relying on external data. By leveraging the dialogue synthesis capability enhanced by DPO, we synthesize well-formatted, coherent dialogues to augment the few-shot real dataset. Furthermore, the two-stage training strategy effectively incorporated synthesized dialogues without introducing artifacts, improving summarization accuracy. Empirical results demonstrated significant improvements: a 1.5% increase in ROUGE scores, a 0.3% improvement in BERT scores. Human evaluations confirmed that our method outperforms the real-only



baseline and, in certain aspects, surpasses human-annotated ground truth summaries. Our approach offers a practical solution for real-world applications with limited data, utilizing the model’s inherent capabilities without external resources. Future work could extend this self-alignment framework to other NLP tasks affected by data scarcity and explore its integration with larger or more diverse LLM architectures.

## 7 Limitations

Comparing to the baseline method, training summarization adapters on few shot real data only, our method requires additional computation cost for data synthesis adapter training and sampling. Also, our method makes an assumption that part of the internal knowledge of LLM is useful for the target domain, which might be incorrect for a highly specialized target domain.

## 8 Ethics Statement

Our research introduces Mutual Reinforcing Data Synthesis (MRDS) within LLMs to enhance few-shot dialogue summarization tasks. While advancing natural language processing, we acknowledge ethical considerations associated with our methodology. By leveraging synthetic data generated by LLMs, we reduce reliance on large-scale real-world datasets that may contain sensitive or personally identifiable information (PII). We strive to protect user privacy and adhere to data protection regulations by using publicly available datasets (CC BY-NC-ND 4.0 and CC BY-NC-SA 4.0) and implementing data anonymization techniques.

Our method utilizes LLMs pre-trained on vast corpora that might contain biases and stereotypes. Although we train our synthesis and summarization models on public datasets with daily conversations, we recognize that biases may persist. We encourage future work to identify and reduce biases in synthetic data generation and in models trained on such data. We are committed to transparency. All experimental details—including data preprocessing, model configurations, and evaluation metrics—are thoroughly documented to ensure reproducibility and allow critical assessment by the research community. By openly sharing our methods and findings, we aim to foster collaboration and uphold ethical standards in AI development.

## References

- Luiz Bonifacio, Hugo Abonizio, Marzieh Fadaee, and Rodrigo Nogueira. 2022. [Inpars: Unsupervised dataset generation for information retrieval](#). In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR ’22*, page 2387–2392, New York, NY, USA. Association for Computing Machinery.
- 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.
- Yulong Chen, Yang Liu, Liang Chen, and Yue Zhang. 2021. Dialogsum: A real-life scenario dialogue summarization dataset. *arXiv preprint arXiv:2105.06762*.
- Zixiang Chen, Yihe Deng, Huizhuo Yuan, Kaixuan Ji, and Quanquan Gu. 2024. Self-play fine-tuning converts weak language models to strong language models. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 6621–6642. PMLR.
- Bharath Chintagunta, Namit Katariya, Xavier Amatriain, and Anitha Kannan. 2021. [Medically aware GPT-3 as a data generator for medical dialogue summarization](#). In *Proceedings of the Second Workshop on Natural Language Processing for Medical Conversations*, pages 66–76, Online. Association for Computational Linguistics.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Guy Feigenblat, Chulaka Gunasekara, Benjamin Sznajder, Sachindra Joshi, David Konopnicki, and Ranit Aharonov. 2021. [TWEETSUMM - a dialog summarization dataset for customer service](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 245–260, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Jiahui Gao, Renjie Pi, LIN Yong, Hang Xu, Jiacheng Ye, Zhiyong Wu, WEIZHONG ZHANG, Xiaodan Liang, Zhenguo Li, and Lingpeng Kong. 2023. [Self-guided noise-free data generation for efficient zero-shot learning](#). In *The Eleventh International Conference on Learning Representations*.
- 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.

- Caglar Gulcehre, Tom Le Paine, Srivatsan Srinivasan, Ksenia Konyushkova, Lotte Weerts, Abhishek Sharma, Aditya Siddhant, Alex Ahern, Miaosen Wang, Chenjie Gu, et al. 2023. Reinforced self-training (rest) for language modeling. *arXiv preprint arXiv:2308.08998*.
- Himanshu Gupta, Kevin Scaria, Ujjwala Ananthaswaran, Shreyas Verma, Mihir Parmar, Saurabh Arjun Sawant, Chitta Baral, and Swaroop Mishra. 2024. TarGEN: Targeted data generation with large language models. In *First Conference on Language Modeling*.
- Jianfeng He, Hang Su, Jason Cai, Igor Shalyminov, Hwanjun Song, and Saab Mansour. 2024. Semi-supervised dialogue abstractive summarization via high-quality pseudolabel selection. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 5976–5996, Mexico City, Mexico. Association for Computational Linguistics.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*.
- Raghav Jain, Anubhav Jangra, Sriparna Saha, and Adam Jatowt. 2022. A survey on medical document summarization. *arXiv preprint arXiv:2212.01669*.
- Martin Josifoski, Marija Sakota, Maxime Peyrard, and Robert West. 2023. Exploiting asymmetry for synthetic training data generation: SynthIE and the case of information extraction. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 1555–1574, Singapore. Association for Computational Linguistics.
- Atharva Kulkarni, Bo-Hsiang Tseng, Joel Ruben Antony Moniz, Dhivya Piraviperumal, Hong Yu, and Shruti Bhargava. 2024. SynthDST: Synthetic data is all you need for few-shot dialog state tracking. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1988–2001, St. Julian’s, Malta. Association for Computational Linguistics.
- Xian Li, Ping Yu, Chunting Zhou, Timo Schick, Omer Levy, Luke Zettlemoyer, Jason E Weston, and Mike Lewis. 2024. Self-alignment with instruction back-translation. In *The Twelfth International Conference on Learning Representations*.
- Yu Li, Baolin Peng, Pengcheng He, Michel Galley, Zhou Yu, and Jianfeng Gao. 2023. DIONYSUS: A pre-trained model for low-resource dialogue summarization. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1368–1386, Toronto, Canada. 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.
- Shikib Mehri, Yasemin Altun, and Maxine Eskenazi. 2022. LAD: Language models as data for zero-shot dialog. In *Proceedings of the 23rd Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 595–604, Edinburgh, UK. Association for Computational Linguistics.
- Yu Meng, Jiaxin Huang, Yu Zhang, and Jiawei Han. 2022. Generating training data with language models: Towards zero-shot language understanding. In *NeurIPS*.
- Siru Ouyang, Jiaao Chen, Jiawei Han, and Diyi Yang. 2023. Compositional data augmentation for abstractive conversation summarization. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1471–1488.
- Sangwon Park, Hongseok Choi, Dongha Choi, and Hyunju Lee. 2024. Gendex: Generative data augmentation strategy leveraging external data for abstractive dialogue summarization. In *Findings of the Association for Computational Linguistics ACL 2024*, pages 3171–3185.
- Minh-Quang Pham, Sathish Indurthi, Shamil Chollampatt, and Marco Turchi. 2023. Select, prompt, filter: Distilling large language models for summarizing conversations. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12257–12265, Singapore. Association for Computational Linguistics.
- Hsuan Su, Ting-Yao Hu, Hema Swetha Koppula, Raviteja Vemulapalli, Jen-Hao Rick Chang, Karren Yang, Gautam Varma Mantena, and Oncel Tuzel. 2024. Corpus synthesis for zero-shot asr domain adaptation using large language models. In *ICASSP*.
- Yuanhe Tian, Fei Xia, and Yan Song. 2024. Dialogue summarization with mixture of experts based on large language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7143–7155.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A. Smith, Daniel Khashabi, and Hannaneh Hajishirzi. 2023. Self-instruct: Aligning language models with self-generated instructions. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13484–13508, Toronto, Canada. Association for Computational Linguistics.
- Kaige Xie, Tong Yu, Haoliang Wang, Junda Wu, Handong Zhao, Ruiyi Zhang, Kanak Mahadik, Ani Nenkova, and Mark Riedl. 2024. Few-shot dialogue summarization via skeleton-assisted prompt transfer in prompt tuning. In *Proceedings of the 18th Conference of the European Chapter of the Association for*

- Computational Linguistics (Volume 1: Long Papers)*, pages 2408–2421, St. Julian’s, Malta. Association for Computational Linguistics.
- Jiayi Yang, Binyuan Hui, Min Yang, Jian Yang, Junyang Lin, and Chang Zhou. 2024. [Synthesizing text-to-SQL data from weak and strong LLMs](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7864–7875, Bangkok, Thailand. Association for Computational Linguistics.
- Jiacheng Ye, Jiahui Gao, Qintong Li, Hang Xu, Jiangtao Feng, Zhiyong Wu, Tao Yu, and Lingpeng Kong. 2022a. [ZeroGen: Efficient zero-shot learning via dataset generation](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11653–11669, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Jiacheng Ye, Jiahui Gao, Zhiyong Wu, Jiangtao Feng, Tao Yu, and Lingpeng Kong. 2022b. [ProGen: Progressive zero-shot dataset generation via in-context feedback](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 3671–3683, Abu Dhabi, United Arab Emirates. 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.
- Weizhe Yuan, Richard Yuanzhe Pang, Kyunghyun Cho, Xian Li, Sainbayar Sukhbaatar, Jing Xu, and Jason E Weston. 2024. Self-rewarding language models. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 57905–57923. PMLR.
- Guangtao Zeng, Wenmian Yang, Zeqian Ju, Yue Yang, Sicheng Wang, Ruisi Zhang, Meng Zhou, Jiaqi Zeng, Xiangyu Dong, Ruoyu Zhang, Hongchao Fang, Penghui Zhu, Shu Chen, and Pengtao Xie. 2020. [MedDialog: Large-scale medical dialogue datasets](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9241–9250, Online. Association for Computational Linguistics.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with bert](#). In *International Conference on Learning Representations*.
- Lulu Zhao, Fujia Zheng, Keqing He, Weihao Zeng, Yuejie Lei, Huixing Jiang, Wei Wu, Weiran Xu, Jun Guo, and Fanyu Meng. 2021. 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.
- Ming Zhong, Yang Liu, Yichong Xu, Chenguang Zhu, and Michael Zeng. 2022. Dialoglm: Pre-trained model for long dialogue understanding and summarization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 11765–11773.
- 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 Appendix

### A.1 Prompt Templates

Table 6 shows the prompt templates we use in this work for different purposes.

### A.2 Examples of Summarization Result

Table 7 presents an example dialogue along with summaries produced by various methods, including the human-labeled ground truth summary, the output from the zero-shot pre-trained model, summaries generated using in-context learning, and those fine-tuned with real data only as well as with the MRDS approach.

### A.3 Synthetic Dialogue from SFT Trained Synthesizer

To show the limited capability of dialogue synthesizer trained on few shot real data, we provide an example of low quality synthetic dialogue in Table 8.

### A.4 Preference Pair Examples

Tables 10 and 9 provide examples of content-based and format-based preference pairs, respectively. In Table 10, the summarization model scores the content alignment of the dialogues using  $\mathcal{M}_{sum}(s|\hat{d})$ , selecting the best and worst ones as the preferred and rejected dialogues. In Table 9, the preferred dialogues are synthesized using IDS, where  $\mathcal{F}(\hat{d}) = 1$ , while the rejected dialogues are generated through

<p><b>Topic extraction</b></p> <p>Please determine the main topic of the provided summary. The topic should be a brief phrase that captures the essence of the summary.</p> <p>Instructions:</p> <ul style="list-style-type: none"> <li>- Ensure the topic is around 2 words in length.</li> <li>- It should clearly reflect the core idea of the summary.</li> <li>- Avoid specific details and names; the topic should be general enough to apply to various summaries.</li> <li>- Use precise and specific language.</li> </ul> <p>Summary: [summary] Topic:</p>
<p><b>Summary synthesis given topic</b></p> <p>Please write a summary for a document based on the provided topic.</p> <p>Instructions:</p> <ul style="list-style-type: none"> <li>- The summary should be approximately [word count] words in length.</li> <li>- Ensure the summary captures the main idea related to the topic.</li> </ul> <p>Topic: [topic] Summary:</p>
<p><b>Dialogue synthesis given summary</b></p> <p>Please write a dialogue based on the given summary. Each line should start with the anonymous speakers' index followed by a colon.</p> <p>Instruction:</p> <ul style="list-style-type: none"> <li>- Use the [num of speaker] speakers in the dialogue: #1, #2 ...</li> <li>- Ensure each line starts with the speaker's index followed by a colon.</li> <li>- Write the dialogue clearly and naturally.</li> <li>- The dialogue should be around [num of turns] turns and [num of words in dialogue] words in length.</li> </ul> <p>Summary: [summary]</p>
<p><b>Dialogue summarization</b></p> <p>Document: [Dialogue] Summarize the provided document</p>
<p><b>Dialogue Summarization (zero shot)</b></p> <p>Dialogue: [Dialogue] Summarize the provided dialogue.</p>
<p><b>Dialogue Summarization (in-context learning)</b></p> <p>Your task is to summarize the provided dialogue. The following are some examples: Example 1: Dialogue: [Example dialogue 1] Summary: [Example summary 1]</p> <p>Example 2: Dialogue: [Example dialogue 2] Summary: [Example summary 2] ...</p> <p>Now, here is the target dialogue: [Dialogue] Summarize the target dialogue.</p>
<p><b>Dialogue Summarization (zero shot + length control)</b></p> <p>Dialogue: [Dialogue] Summarize the provided dialogue. The summary should be around [average word count of summaries in training set] words in length.</p>

Table 6: Prompt templates.

---

**Input dialogue**

---

Avery: You went to Ethan's house?  
David: yeah I had to babysit  
Avery: Aww, how do you babysit, just curious  
David: I had to go through a lot :/  
Avery: Was his sister naughty  
David: Tooo much  
Avery: Lol  
David: I will just refuse net time :/  
Avery: As you wish  
David: :/  
Avery: I just got his text  
David: What is he saying  
Avery: He is asking me to say thanks to you  
David: yeah whatever<3  
Avery: He was saying that your phone was switched off  
David: Yeah i have just turned it on  
Avery: I have told him about that  
David: k  
Avery: Gotta go now

---

**Ground Truth Summary**

---

David was looking after Ethan's sister. Ethan is grateful. David won't do it again.

---

**Generated Summary (Zero Shot)**

---

The dialogue is between Avery and David, discussing David's experience babysitting Ethan's sister. David mentions that the sister was "too much" and that he will refuse to babysit again. Avery shares that Ethan sent a text asking Avery to thank David, and also mentions that David's phone was switched off. David confirms that he just turned it on and the conversation ends with Avery saying they have to go.

---

**Generated Summary (In-context Learning)**

---

David babysat at Ethan's house and had a difficult time. Avery asks how he managed it and David says it was a lot of work. Avery asks if Ethan's sister was naughty and David confirms that she was. David jokes that he will refuse to babysit again. Avery mentions Ethan sent a text asking Avery to thank David, and David agrees to let Avery do so. Avery also mentions that Ethan's phone was switched off, which David had just turned back on.

---

**Generated Summary (Real only)**

---

David had to babysit Ethan's sister.

---

**Generated Summary (MRDS, Ours)**

---

David had to babysit Ethan's sister. Avery got a text from Ethan saying thanks to David.

---

Table 7: Summarization results from our method and baselines.

<b>Summary <math>s</math>:</b> #2 can't join #1 and Mike for a game because he is studying.
<b>Generated dialogue <math>\hat{d}</math>:</b> #1: So will you come at 2? #2: <b>sure bro</b> , what time you gonna be with Mike? #1: about 20 min.00, Mike will call you 2: <b>he is going</b> 3 <b>no way, he's still sleeping hehehe</b> 2 <b>he's studying</b> 1: <b>and?</b> <b>not coming</b>

Table 8: Low quality dialogue generated by SFT-trained dialogue synthesizer. #1 and #2 are anonymized speaker names. The text highlighted in red shows content that is not aligned with the summary (#2 can't join and #2 is studying instead of Mike). The lines highlighted in blue are format errors, since they do not start with a proper speaker name (#1, #2, ...).

the SFT synthesizer's one-shot synthesis, where  $\mathcal{F}(\hat{d}) = 0$ .

## A.5 Data Formatting and Anonymization

To simplify the training of the data synthesizer, we perform anonymization preprocessing on both summaries and dialogues. Specifically, we extract the number of speakers and their names from the dialogues as metadata and replace the names in both dialogues and summaries with uniform identifiers (e.g., "#1", "#2").

### A.5.1 Data Formatting

To ensure that the synthetic data adheres to the correct format, we have established a set of formatting rules:

- **Speaker Identity:** each sentence in the dialogues must begin with a speaker identifier followed by a colon (e.g., "#1:", "#2:").
- **Consistency of Names:** Synthetic dialogues and summaries should not contain incorrect or extra names—for example, "#4" in dialogue-summary pairs with fewer than four speakers, or a "#" symbol without a number.
- **Inclusion of Speaker Names:** Synthetic summaries must contain at least one anonymized speaker name.

For summary synthesis, we discard any summaries that do not comply with these formatting rules. For dialogue synthesis, we apply Iterative Dialogue Synthesis (IDS) to ensure that the synthetic dialogues conform to the correct format.

### A.5.2 Iterative Dialogue Synthesis

Table 11 illustrates an example of iterative dialogue synthesis, showing partial inputs given to the synthesizer and the corresponding synthesized outputs. During the synthesis process, we check for formatting issues as described in Sec. A.5.1. If an error is found, the flawed sentence, along with all subsequent sentences, is discarded, and the correctly formatted segment (e.g., "#1:", "#2:") is used to re-synthesize the dialogue.

### A.5.3 Synthetic Dialogue Recovery

After data synthesis, we restore the anonymized dialogues and summaries before training the summarization model by replacing placeholders with real names using metadata extracted during the anonymization process. For synthetic summaries that lack original metadata, we seed the synthesis process with metadata from the real data in the few-shot training set, then replace the placeholders accordingly. Alternatively, we can generate random names and speaker numbers to populate the placeholders.

## A.6 Human Evaluation Details

We hired five machine learning researchers, and received their consent to report the results from their annotation work. We adopt the human evaluation method proposed in the previous work (Xie et al., 2024). For completeness, we describe the 4 metrics (informativeness, faithfulness, fluency and redundancy) and the corresponding instructions in Table 12.

---

**Summary:**  
On Sunday #1 and #2 are meeting with Dominica and James who got married recently.

---

**Preferred dialogue:**  
#1: Hahaha we must meet Dominica tomorrow with James.  
#2: yep  
#1: I dont think she will be very available after the marriage.  
#2: ye ye, she's very busy  
#1: I asked :) :) <file\_gifile>  
#2: so 8 o'clock?  
#1: sounds super

---

**Rejected dialogue:**  
#1: are u free next Sunday?  
#1 file\_gifiles : yeah  
#2: ok  
#1: wanna meet Domin?  
#2: ok  
# fileifile  
# gileile  
# : ok: <file :>  
file :) <file :) : :file :) <file < :>  
# Ok, ok :) <file :>  
:) it sounds cool

---

Table 9: Example of a format-based preference pair. #1 and #2 are anonymized speaker names. The lines highlighted in blue are format errors, since they do not start with a proper speaker name (#1, #2).

---

**Summary:**  
#1 is outside waiting for #2, who is late.

---

**Preferred dialogue:**  
#1: Hey, it's #1. I'm sitting outside where are you?.  
#2: <file\_gifile>  
#1: You're late! And I was waiting for you there.

---

**Rejected dialogue:**  
#1: I can't come to the meeting  
#2: what did?  
#1: I have not enough money for cab. I must wait for #1  
#2: #1 is waiting for us!  
#1: how can I get there?  
#2: use bus, it's the fastest, come back me

---

Table 10: Example of a content-based preference pair. #1 and #2 are anonymized speaker names.

<b>Input summary</b>
Both #1 and #2 forgot about Monica’s birthday today. As they do not know what present to buy for her, #2 will phone Monica’s boyfriend and ask his advice.
<b>First partial response</b>
<b>#1:</b>
<b>First synthesis result</b>
#1: I forgot about Monicas birthday today haha #2: Me to, me too #1: what now? #2: we should make something? #1: what do you propose to buy? #1: I’m asking my brother’s opinion, he knows her a bit more #2: same here, I will phone her boyfriend #1: you must have more info <i>2: ya, I will get some advice</i> <i>1:</i> <i>I asked and we decided to buy her something</i> <i>2.: I can not see.</i> <i>5. This 2 1 3’s got 1 6.75.</i>
<b>Second partial response</b>
#1: I forgot about Monicas birthday today haha #2: Me to, me too #1: what now? #2: we should make something? #1: what do you propose to buy? #1: I’m asking my brother’s opinion, he knows her a bit more #2: same here, I will phone her boyfriend #1: you must have more info <b>#2:</b>
<b>Second synthesis result</b>
#1: I forgot about Monicas birthday today haha #2: Me to, me too #1: what now? #2: we should make something? #1: what do you propose to buy? #1: I’m asking my brother’s opinion, he knows her a bit more #2: same here, I will phone her boyfriend #1: you must have more info <b>#2: sure</b> #1: I am already looking for some ideas #2: it looks hopeless #2: I will call and he phone #1: yes

Table 11: An example of iterative dialogue synthesis (IDS), condensed for brevity. In each step, the correctly-formatted partial response is concatenated to the initial dialogue synthesis prompt and given as input for generation. Newly generated text appears in blue, while discarded text is italicized. In each step, we also add an anonymized speaker name (shown in yellow) at the end of the partial response before using it for generation. We always use #1 as the first partial response, and randomly choose an anonymized speaker in the subsequent steps.



<b>Informativeness</b>
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.
<b>Faithfulness</b>
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.
<b>Fluency</b>
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.
<b>Redundancy</b>
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.

Table 12: Human evaluation metrics and their corresponding instructions

## A.7 Training Hyperparameters

Table 13 presents the hyperparameters used for the summarization and dialogue synthesis models. We conducted an extensive search for hyperparameters on the summarization model trained exclusively with real data. All other SFT training follows the same set of hyperparameters. The LR threshold for the two-stage summarization approach refers to the minimum learning rate reached during the first stage of training with synthetic data. Once this threshold is met, the second stage begins using real data, applying the same hyperparameters as those used for the real-only summarization model.

The DPO synthesis model is initialized from the checkpoints of the SFT synthesis model and is trained using both DPO and SFT loss. The batch size for DPO training is four, randomly sampled from two preference sets: format-based and content-based preference pairs. Additionally, SFT training data with a batch size of one is included in the process, and the combined losses are used to update the synthesis model.

## A.8 Dialogue Synthesis Quality

We evaluated the performance of different dialogue synthesis methods, including direct synthesis (One-shot), iterative dialogue synthesis (Iterative), and

various configurations of DPO, as presented in Table 14. In our DPO training, we constructed different sets of preference pairs based on the definitions provided earlier. For the single joint preference set, we used dialogues from the post-processing step as the preferred data ( $\hat{d}_1$ ) and the raw dialogues without iterative synthesis as the rejected data ( $\hat{d}_2$ ), forming preference pairs:  $\{s, \hat{d}_1, \hat{d}_2 \mid s \in \mathcal{S}, \mathcal{F}(\hat{d}_1) = 1, \mathcal{F}(\hat{d}_2) = 0, \mathcal{M}_{\text{sum}}(s \mid \hat{d}_1) > \mathcal{M}_{\text{sum}}(s \mid \hat{d}_2)\}$ . For the separated preference sets, we included both formatting preference pairs:  $\{s, \hat{d}_1, \hat{d}_2 \mid s \in \mathcal{S}, \mathcal{F}(\hat{d}_1) = 1, \mathcal{F}(\hat{d}_2) = 0\}$  and summarization preference pairs:  $\{s, \hat{d}_1, \hat{d}_2 \mid s \in \mathcal{S}, \mathcal{M}_{\text{sum}}(s \mid \hat{d}_1) > \mathcal{M}_{\text{sum}}(s \mid \hat{d}_2)\}$ . We also incorporated supervised fine-tuning data into the DPO training, indicated as "+SFT."

From Table 14, the SFT model’s one-shot synthesis had low format correctness (24%) and a high summarization cross-entropy (CE) loss of 4.74. While Iterative synthesis ensured 100% format correctness, the CE loss remained similar at 4.69, showing little improvement in content alignment. Applying DPO with the joint preference set improves format correctness to 78% and reduces the summarization CE loss to 4.62. However, training with DPO loss alone was unstable, so adding SFT loss (Joint Preference Set + SFT), stabilized

Table 13: Hyper-parameters used for training dialogue summarization, two-stage dialogue summarization, dialogue/summary synthesis, and DPO-based dialogue synthesis models.

	Summarization	Two-stages Summ.	Synthesis	DPO Syn.
Loss	CE	CE	CE	DPO and CE
Batch Size	10	10	10	4 DPO & 1 SFT
Learning Rate	2.0e-4	2.0e-4	2.0e-4	1.0e-5
Optimizer	ReduceLROnPlateau	ReduceLROnPlateau	ReduceLROnPlateau	Fixed LR
Validation Steps	2	2	2	20
Patience	5	5	5	-
Factor	0.7	0.7	0.7	-
Warmup Steps	50	50	50	-
Early Stopping	50	50	50	-
LR Threshold	-	2.0e-5	-	-
LoRA Rank		16		
LoRA Alpha		32		
Dropout Rate		0.4		

Synthesis Model	Synthesis Method	Format Corr.	Summ. CE
SFT	One-shot	24%	4.74
SFT	Iterative	100%*	4.69
Joint Preference Set	One-shot	78%	4.62
Joint Preference Set + SFT	One-shot	<b>96%</b>	<u>4.45</u>
Separate Preference Sets + SFT (MRDS)	One-shot	<u>92%</u>	<b>4.13</b>

Table 14: Dialogue Synthesis with different training strategy on SAMSum 300 shots experiments.

training, achieving 96% format correctness and a CE loss of 4.45. Yet, this approach prioritized formatting over content quality, limiting further improvement in CE loss.

By separating the preference sets and including both formatting and summarization preferences in DPO training (Separate Preference Sets + SFT), our MRDS approach effectively balanced format correctness (92%) and significantly lowered the CE loss to 4.13, indicating better alignment with summaries. This demonstrates that separating preference sets allows concurrent optimization of formatting and content, leading to the best overall performance.

### **A.9 Statistical Analysis of Dialogue Summarization Experiments**

To further validate the improvements introduced by the MRDS method, we conduct each experiment three times under the 100-shot setting on both the SAMSum and DialogSum datasets. For each run, we compute the mean and sample standard deviation (using  $n-1$  in the denominator) for ROUGE-1, ROUGE-2, ROUGE-L, and BERTScore. In addition, independent two-sample t-tests (with degrees of freedom  $df = 4$ ) were performed to assess the statistical significance of the differences between the baseline (summarization adapter fine-tuned with real data only) and MRDS results.

Tables 15 and 16 summarize the results for SAMSum and DialogSum, respectively. For instance, on SAMSum the ROUGE-1 score improved from an average of 50.90% (std = 0.18%) to 52.10% (std = 0.25%), a gain of 1.20 percentage points ( $t = 6.80$ ,  $p \approx 0.003$ ). Similar statistically significant improvements are observed across all metrics and on both datasets.

The t-test results confirm that the improvements achieved by MRDS over the baseline are statistically significant ( $p$ -values  $< 0.05$ ) across all metrics on both datasets.

### **A.10 Comparison with State-of-the-art Methods**

In Table 17, we compare the proposed MRDS with three existing state-of-the-art methods that use data augmentation and pseudo labeling to address the data scarcity problem for dialogue summarization task. The results show that our MRDS method achieves competitive or better performance in Rouge-1, Rouge-2 and Rouge-L. In the table, we also compare the different backbone models,

the amount of few-shot paired data and the need of unlabelled data among existing methods and MRDS.

Table 15: Statistical analysis on SAMSum 100-shot experiments (values in percentages). The baseline method adopts summarization adapter finetuned with real data only.

<b>Metric</b>	<b>Baseline</b> (mean $\pm$ std)	<b>MRDS (Ours)</b> (mean $\pm$ std)	<b>Difference</b>	<b>t-value</b>	<b>p-value</b>
ROUGE-1	50.90 $\pm$ 0.18	52.10 $\pm$ 0.25	+1.20	6.80	0.003
ROUGE-2	26.54 $\pm$ 0.04	27.54 $\pm$ 0.31	+1.00	5.52	0.005
ROUGE-L	42.62 $\pm$ 0.09	43.42 $\pm$ 0.35	+0.80	3.87	0.018
BERTScore	86.59 $\pm$ 0.04	86.81 $\pm$ 0.07	+0.22	4.84	0.011

Table 16: Statistical analysis on DialogSum 100-shot experiments (values in percentages). The baseline method adopts summarization adapter finetuned with real data only.

<b>Metric</b>	<b>Baseline</b> (mean $\pm$ std)	<b>MRDS (Ours)</b> (mean $\pm$ std)	<b>Difference</b>	<b>t-value</b>	<b>p-value</b>
ROUGE-1	44.04 $\pm$ 0.55	45.47 $\pm$ 0.15	+1.43	4.35	0.015
ROUGE-2	18.23 $\pm$ 0.35	19.26 $\pm$ 0.27	+1.03	4.06	0.017
ROUGE-L	35.97 $\pm$ 0.49	37.25 $\pm$ 0.13	+1.28	4.32	0.016
BERTScore	86.81 $\pm$ 0.11	87.19 $\pm$ 0.02	+0.38	5.86	0.008

Table 17: Comparison among the proposed MRDS and three existing state-of-the-art methods (Semi-CODA (Chen and Yang, 2021), SiCF (He et al., 2024) and COMPO (Ouyang et al., 2023)) on SAMSum dataset. \*Using unlabeled data.

Methods	backbone	num. of paired data	Rouge-1	Rouge-2	Rouge-L
Semi-CODA*	BART-large	147	42.16	17.82	38.89
SiCF*	DialogLM	147	45.85	19.90	35.96
COMPO	BART-large	147	49.78	24.65	45.41
MRDS (ours)	Llama3-8B-Instruct	100	52.1	27.5	43.4