

SWING : Balancing Coverage and Faithfulness for Dialogue Summarization

Kung-Hsiang Huang^{♦*} Siffi Singh[◇] Xiaofei Ma[◇] Wei Xiao[◇]
Feng Nan[◇] Nick Dingwall[◇] William Yang Wang[◇] Kathleen McKeown[◇]

[♦]University of Illinois Urbana-Champaign [◇]AWS AI Labs
khhuang3@illinois.edu

{siffis, xiaofeim, weixiaow, nanfen, nickding, wyw, mckeownk}@amazon.com

Abstract

Missing information is a common issue of dialogue summarization where some information in the reference summaries is not covered in the generated summaries. To address this issue, we propose to utilize natural language inference (NLI) models to improve coverage while avoiding introducing factual inconsistencies. Specifically, we use NLI to compute fine-grained training signals to encourage the model to generate content in the reference summaries that have not been covered, as well as to distinguish between factually consistent and inconsistent generated sentences. Experiments on the DIALOGSUM and SAMSUM datasets confirm the effectiveness of the proposed approach in balancing coverage and faithfulness, validated with automatic metrics and human evaluations. Additionally, we compute the correlation between commonly used automatic metrics with human judgments in terms of three different dimensions regarding coverage and factual consistency to provide insight into the most suitable metric for evaluating dialogue summaries.¹

1 Introduction

Dialogue summarization is a text generation task that aims to produce a compact summary given a piece of conversation. Conventional approaches to dialogue summarization rely on features of conversation data (Goo and Chen, 2018; Li et al., 2019; Oya et al., 2014). Recently, the rise of large pre-trained language models (LMs) has enabled coherent and fluent summaries to be generated without these features. However, low coverage and factual inconsistency remain two pressing issues as studies have shown that the summaries generated from these pre-trained LMs often do not fully cover the reference (Liu and Chen, 2021; Tang et al.,

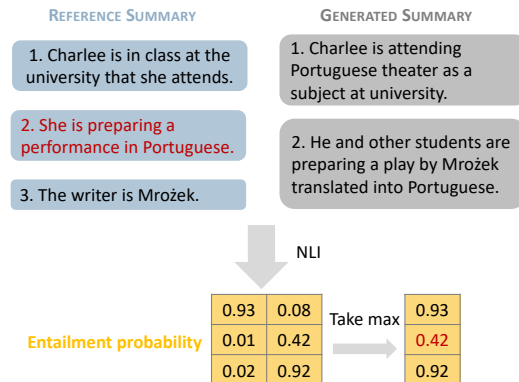


Figure 1: An illustration of how NLI can help determine whether a reference sentence is covered by the generated summary. We compute the entailment probability from each reference sentence (i.e. premise) to each generated sentence (i.e. hypothesis). By taking the max value along the row dimension, the resulting vector denotes the probability that each reference sentence entails a sentence in the generated summary. In this example, the entailment probability for the second reference sentence is low, indicating that this sentence is likely not covered by the generated summary.

2022) and that the generated summaries are often not factually consistent with the inputs (Zhang et al., 2020b; Maynez et al., 2020; Cao and Wang, 2021). If an unfaithful dialogue summarization model with low coverage is deployed for public use, it could spread misinformation and generate misleading content that only covers partial facts of a conversation. Hence, we are urgently in need of a solution to improve coverage without negatively impacting faithfulness for dialogue summarization.

Relatively little work addresses coverage and factual inconsistency for dialogue summarization. Some work addresses the issue of unfaithfulness with a controllable generation framework guided by person named entities (Liu and Chen, 2021) or summary sketches (Wu et al., 2021). Tang et al. (2022) categorize factual inconsistencies for dialogue summarization into different types of errors,

^{*}Work done while interning at Amazon.

¹We release our source code for research purposes: <https://github.com/amazon-science/AWS-SWING>.

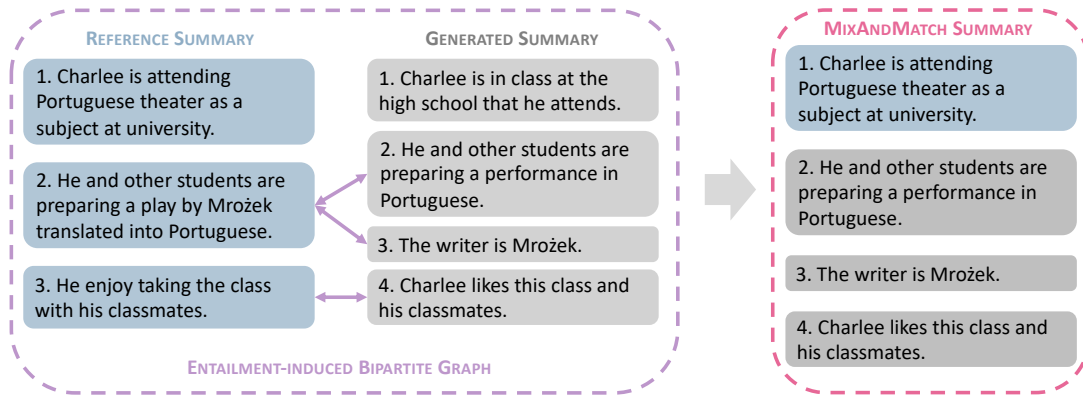



Figure 2: Illustration of how an entailment-induced bipartite graph is built and how a MIXANDMATCH summary is derived. With the NLI model, we determine which sentences from each summary contain equivalent information by computing the entailment probabilities between pairs of generated sentences and reference sentences, as indicated by the purple edges. Based on the graph, we determine that the generated summary does not cover the first reference sentence and that the first generated sentence is not faithful. Hence, the MIXANDMATCH summary is formed by combining the first reference sentence and the second to the fourth generated sentence.

such as missing information and wrong reference. Their framework integrates a contrastive loss and a self-supervised loss to reduce multiple types of errors. However, a great portion ($> 40\%$) of their outputs does not cover the full content of the reference summary. Thus, it is important to address coverage and factual consistency synergistically in dialogue summarization. The issue where the content in the reference does not occur in the generated summary is known as the missing information issue (Liu and Chen, 2021; Tang et al., 2022). In this work, we aim to mitigate missing information in the summary while being faithful to the dialogue.

We propose SWING , Summarizing Dialogue With NLI Guidance. Our approach samples a summary from the model and utilizes natural language inference (NLI) to determine (1) the faithfulness of each generated sentence and (2) whether each reference sentence has been covered by the generated summary. An example is shown in Figure 1. Based on the results computed by NLI, two losses are proposed to encourage the model to generate missing information and distinguish between factually consistent and inconsistent generated sentences.

Our contributions can be summarized as follows:

- We propose SWING, a dialogue summarization framework that effectively addresses missing information through two losses computed using NLI. The first loss encourages the model to recover content missing from the reference summaries. The second loss instructs the model to differentiate between factually consistent and inconsistent

generated sentences.

- Our approach achieves the best performance in mitigating missing information on two public dialogue summarization datasets, DIALOGSUM (Chen et al., 2021b) and SAMSUM (Gliwa et al., 2019), as validated by automatic metrics and human judges.
- We measure the correlation of human judgments with conventional and recently developed automatic metrics to provide intuition for future research on evaluating the faithfulness and coverage of dialogue summaries.

2 Method

Upon analyzing the dialogue summaries in SAMSUM, we observe that dialogues are often summarized linearly, consistent with the findings of Wu et al. (2021). Therefore, we segment the summaries into sentences and use a natural language inference (NLI) model to provide finer-grained training signals at the sentence level for two goals: (1) encourage generating sentences in the reference summaries that have not been covered by the generated sentences and (2) differentiate factually consistent generated sentences from inconsistent ones. To achieve these goals, we first determine the faithfulness of each sentence using an entailment-induced bipartite graph (§2.1). Then, we propose two new losses addressing each challenge in turn: an **Uncovered Loss** that encourages the model to recover missing information (§2.2) and a **Contrastive Loss** that brings closer the representations of the reference summary and the generated sentences that

Algorithm 1: Entailment-induced Bipartite Graph

Input: A reference summary $S^* = \{s_1^*, \dots, s_n^*\}$, a generated summary $S = \{s_1, \dots, s_m\}$;
Output: The bipartite mapping ϕ between sentences in S^* and S ;

- 1 Initialize ϕ as a zero matrix of size $n \times m$ where $n = |S^*|$ and $m = |S|$;
- 2 Let τ be the entailment threshold;
- 3 // Resolve 1-to-many mappings;
- 4 **for** $i \leftarrow 1$ **to** n **do**
- 5 $V \leftarrow \emptyset$;
- 6 **for** $j \leftarrow 1$ **to** m **do**
- 7 **if** $p_{\text{ent}}(s_i^*, s_j) > \tau$ **and** $\phi(i, j) = 0$ **then**
- 8 $V \leftarrow V \cup j$;
- 9 $s_V \leftarrow$ Concatenate sentences in $\{s_v, \forall v \in V\}$;
- 10 **if** V is consecutive **and** $p_{\text{ent}}(s_V, s_i^*) > \tau$ **then**
- 11 **for** $v \in V$ **do**
- 12 $\phi(i, v) \leftarrow 1$;
- 13 // Resolve many-to-1 mappings;
- 14 **for** $j \leftarrow 1$ **to** m **do**
- 15 $V \leftarrow \emptyset$;
- 16 **for** $i \leftarrow 1$ **to** n **do**
- 17 **if** $p_{\text{ent}}(s_j, s_i^*) > \tau$ **and** $\phi(i, j) = 0$ **then**
- 18 $V \leftarrow V \cup i$;
- 19 $s_V^* \leftarrow$ Concatenate sentences in $\{s_v^*, \forall v \in V\}$;
- 20 **if** V is consecutive **and** $p_{\text{ent}}(s_V^*, s_j) > \tau$ **then**
- 21 **for** $v \in V$ **do**
- 22 $\phi(v, j) \leftarrow 1$;
- 23 // Resolve 1-to-1 mappings;
- 24 **for** $i \leftarrow 1$ **to** n **do**
- 25 **for** $j \leftarrow 1$ **to** m **do**
- 26 **if** $\phi(i, j) = 0$ **and** $p_{\text{ent}}(s_j, s_i^*) > \tau$ **and**
- 27 $p_{\text{ent}}(s_i^*, s_j) > \tau$ **then**
- 28 $\phi(i, j) \leftarrow 1$;
- 28 Return ϕ ;

contain equivalent information to some sentences in the reference summary (§2.3). For the rest of this paper, we use *reference sentence* and *generated sentence* to refer to a sentence in the reference summary and the generated summary, respectively.

2.1 Entailment-induced Bipartite Graph

To determine which reference sentence has not been covered by the generated summary and which generated sentence is not faithful to the reference summary, we construct a bipartite graph that links sentences between a reference summary and a generated summary. An edge indicates the linked sentences contain equivalent information. If no edge connects to a reference sentence, we consider this sentence not covered by the generated summary. Similarly, if a generated sentence is not linked in the bipartite graph, this sentence is likely not

faithful to the reference summary. We use the entailment probabilities computed by an NLI model to determine whether a pair of sentences contain equivalent information. The procedure of constructing the bipartite graph is shown in Algorithm 1.

The NLI model takes in two sentences, a premise (P) and a hypothesis (H), and computes whether P entails, contradicts, or is neutral to H . Here, we only focus on the entailment probability from the i -th reference sentence to the j -th generated sentence $p_{\text{ent}}(s_i^*, s_j)$. We use the ROBERTA-LARGE model² trained on the MNLI dataset, achieving an accuracy of around 91%, which is on par with the performance of state-of-the-art models.

Let $\phi(i, j)$ denote the mapping between the i -th reference sentence and the j -th generated sentence. $\phi(i, j) = 1$ if a link exists between s_i^* and s_j ; otherwise, $\phi(i, j) = 0$. We first consider a simplified setting by assuming each reference sentence can be mapped to at most one generated sentence, and vice versa (i.e. $0 \leq \sum_j \phi(i, j) \leq 1$). In this setting, we can determine whether two sentences contain equivalent information by checking the entailment relation from both directions (lines 26-27).

$$\phi(i, j) = \begin{cases} 1, & p_{\text{ent}}(s_i^*, s_j) > \tau \wedge p_{\text{ent}}(s_j, s_i^*) > \tau \\ 0, & \text{otherwise} \end{cases} \quad (1)$$

Here, τ is a hyperparameter that indicates the entailment threshold.

However, one reference sentence may contain information equivalent to multiple generated sentences (one-to-many mappings) and vice versa (many-to-one mappings). In Figure 2, for example, the second reference sentence contains information equivalent to the second and the third generated sentences combined. This relation cannot be discovered if we only check the entailment relation between pairs of individual sentences.

Therefore, we must resolve one-to-many and many-to-one mappings before checking one-to-one mappings. To find one-to-many mappings, for every reference sentence s_i^* , we look for consecutive generated sentences $\{s_j, s_{j+1}, \dots, s_{j+k}\}$ s.t. $\max_i p_{\text{ent}}(s_i^*, s_m) > \tau \forall m \in \{j, \dots, j+k\}$ (lines 6-8). We only check for consecutive sentences based on our previous observation that dialogues are often summarized linearly. For every match, we concatenate the generated sentences $s_{j:j+k} = \{s_j, s_{j+1}, \dots, s_{j+k}\}$ and

²<https://huggingface.co/roberta-large-mnli>

check whether $s_{j:j+k}$ entails the reference sentence s_i^* (lines 8-9). If the entailment holds, we let $\phi(i, m) = 1 \quad \forall m \in \{j, \dots, j+k\}$ (lines 11-12). The same approach is used to address many-to-one mappings (lines 14-22). Following Algorithm 1, a bipartite graph is built between the generated summary and the reference summary. Henceforth, we denote the reference sentences that have not been covered as $\underline{S}^* = \{s_i^* | \forall j \phi(i, j) = 0\}$ and generated sentences that can be mapped to some of the reference sentences as $\underline{S} = \{s_j | \exists i \phi(i, j) = 1\}$.

2.2 Uncovered Loss

The objective of the uncovered loss is to encourage the model to generate information from the reference summary that the generated summary has not covered. To this end, we train the model with MIXANDMATCH summaries, which are constructed by combining reference sentences that are not covered by the generated summary and generated sentences that contain information equivalent to some of the reference sentences. An example is shown in Figure 2.

The MIXANDMATCH summary \hat{S} is constructed by taking the union of \underline{S} and \underline{S}^* and sorting the sentences by their index,

$$\hat{S} = \text{SORT}(\underline{S} \cup \underline{S}^*). \quad (2)$$

The uncovered loss is effectively maximum likelihood estimation (MLE) with MIXANDMATCH summaries being the decoding targets:

$$\mathcal{L}_{\text{Uncovered}} = - \sum_t \log p(\hat{S}_t | \hat{S}_{<t}, \mathcal{D}), \quad (3)$$

where \mathcal{D} is the original dialogue and \hat{S}_t denotes the t -th token in the MIXANDMATCH summary.

The main advantages of constructing MIXANDMATCH summaries over other positive sample construction approaches, such as back translation and paraphrasing, are the two desired properties of this formulation. First, the model already has a high probability of generating sentences in \underline{S} . Therefore, the loss function (Equation (3)) does not penalize the model much for generating these sentences. Second, the penalty for generating sentences \underline{S}^* is larger since the model has a lower probability of generating those sentences.

2.3 Contrastive Loss

In the early stage of our experiment, the original goal was to discourage the model from generating

factually inconsistent sentences. We adopt unlikelihood training (Welleck et al., 2020) to decrease the probability of sampling these sentences from the model. However, we found that this objective causes the model to generate nonsense sequences. This phenomenon was also observed when we experimented with CONSEQ (Nan et al., 2021), which also incorporates such a loss function into its training process, as shown in §4.1. We hypothesize that it resulted from the fact that sentences in dialogue summaries share similar structures. Hence, using the unlikelihood training objective would confuse the model.

Instead, we pivoted our focus on differentiating factually consistent sentences from their inconsistent counterparts with the proposed contrastive loss. For each summary, we use the factually inconsistent sentences as negative samples (i.e. $s_j \notin \underline{S}$) and consistent sentences as positive samples (i.e. $s_j \in \underline{S}$). The contrastive learning objective takes a similar form as the InfoNCE loss (Oord et al., 2018):

$$\mathcal{L}_{\text{Contrastive}} = - \sum_{s_i \in \underline{S}} \frac{\exp(\cos(h_i, h_{S^*}))}{\sum_{s_j \in \underline{S}} \exp(\cos(h_j, h_{S^*}))} \quad (4)$$

, where h_i and h_j denote the representations of the generated sentences, h_{S^*} means the representations of the reference summary, and $\cos(\cdot, \cdot)$ denotes cosine similarity. The main difference between our contrastive objective and the other work (Cao and Wang, 2021; Tang et al., 2022) is granularity. Equation (4) operates at the sentence level rather than the summary level; therefore, it provides finer-grained training signals.

2.4 Training

The final loss function that our model is optimized with is a weighted sum of the two aforementioned loss functions and MLE,

$$\mathcal{L}_{\text{Final}} = \mathcal{L}_{\text{MLE}} + \alpha \mathcal{L}_{\text{Uncovered}} + \beta \mathcal{L}_{\text{Contrastive}}, \quad (5)$$

where \mathcal{L}_{MLE} is:

$$\mathcal{L}_{\text{MLE}} = - \sum_t \log p(S_t^* | S_{<t}^*, \mathcal{D}). \quad (6)$$

3 Experiments

3.1 Datasets

Experiments are conducted on two English-language dialogue summarization datasets: SAMSUM (Gliwa et al., 2019) and DIALOGSUM (Chen

Model	DIALOGSUM								SAMSUM							
	RL _F	RL _R	BS _F	BS _R	FC _F	FC _R	QS	QFE	RL _F	RL _R	BS _F	BS _R	FC _F	FC _R	QS	QFE
TextRank	27.74	29.16	-3.000	-3.039	60.55	59.54	-1.948	0.566	15.08	16.15	-4.374	-3.891	34.28	33.02	-2.172	0.237
BART-LARGE	50.82	56.78	-2.012	-1.960	82.90	85.86	-1.183	1.854	49.53	52.71	-2.248	-2.332	62.46	61.28	-0.912	2.335
CTRLDIASUMM	48.99	57.25	-2.145	-1.985	82.55	85.96	-1.214	1.817	47.79	51.17	-2.360	-2.414	61.50	61.76	-0.957	2.272
CODS	48.51	48.36	-2.379	-2.214	83.33	86.81	-1.246	1.860	48.39	47.68	-2.643	-2.593	61.21	62.01	-0.867	2.345
CONSEQ	22.82	19.50	-3.480	-3.588	84.24	73.14	-1.474	0.208	12.04	7.62	-5.908	-7.278	41.23	13.77	-2.058	0.035
CLIFF	51.87	56.22	-2.012	-1.973	85.38	86.30	-1.106	2.109	43.70	45.49	-2.485	-2.340	55.47	56.01	-1.063	1.891
CONFIT	50.44	55.65	-2.049	-2.016	83.34	86.37	-1.179	1.790	49.29	52.76	-2.188	-2.316	65.03	63.12	-0.819	2.343
SWING	51.96	59.04*	-1.999*	-1.904*	86.48	89.03	-1.082*	2.087	50.08	52.91	-2.228	-2.310*	64.19	63.52	-0.829	2.407*
- $\mathcal{L}_{\text{Uncovered}}$	50.94	60.06*	-2.044	-1.895*	83.26	87.45	-1.075*	2.339*	49.78	53.57	-2.231	-2.295*	63.81	63.11	-0.876	1.989
- $\mathcal{L}_{\text{Contrastive}}$	51.53	59.27*	-2.012	-1.901*	82.90	85.86	-1.130	2.399*	49.73	53.95	-2.185*	-2.143*	63.47	63.15	-0.886	2.027

Table 1: Performance comparison on DIALOGSUM and SAMSUM. - $\mathcal{L}_{\text{Uncovered}}$ and - $\mathcal{L}_{\text{Contrastive}}$ denote variants of SWING by ablating the corresponding loss. RL denotes ROUGE-L (%), BS denotes BARTSCORE, FC denotes FACTCC (%), QS denotes QUALS, and QFE denotes QAFACTEVAL. The subscripts F and R denote F1 score and recall, respectively. The proposed method outperforms previous systems on both DIALOGSUM and SAMSUM in most metrics, especially on the recall measures. Statistical significance over previous best systems computed with the permutation test (Fisher et al., 1937) is indicated with * ($p < .01$).

et al., 2021b). SAMSUM contains 16,369 online chitchat dialogues with an average of around 94 tokens per dialogue. DIALOGSUM is a spoken dialogue dataset that consists of 13,460 samples in total. With an average token count of about 131, the dialogues in DIALOGSUM are under real-life scenarios with clear communication patterns and intents. Details of the dataset statistics can be found in Appendix A.

3.2 Metrics

Our evaluation focuses on measuring the factual consistency, particularly the missing information challenge, of the summarization models. Therefore, we adopt recently developed metrics that have been shown to correlate well with human judgments in terms of faithfulness. BARTScore (Yuan et al., 2021) computes the semantic overlap between the generated summary and the reference summary by calculating the logarithmic probability of generating each summary conditioned on the other one. Since our goal is to assess how well the model reduce information missing from the reference summary, we consider the *Recall* (R) setting where we assess $p(S^* | S, \theta)$, the likelihood of generating the reference summary S given the generated summary S^* . FactCC (Kryscinski et al., 2020) is an entailment-based metric that predicts the faithfulness probability of a claim w.r.t. with the source texts. Similar to BARTScore, we use FactCC in the *Recall* setting where the claim is a reference sentence and the source text is the generated summary. We report the mean of the average CORRECT probability of each sentence within a generated summary.

In addition, we report the ROUGE-L metric (Lin, 2004), which has been also shown to better reflect

faithfulness compared to ROUGE-1 and ROUGE-2 (Pagnoni et al., 2021). For these metrics, we also consider the *F1* setting, where we compute each metric in the reverse direction ($S^* \rightarrow S$) and then take the average of both directions, to validate that the model is not generating too much redundant information. Finally, two recently introduced QA-based metrics that have demonstrated close approximation to human judgements in terms of factuality, QUALS (Nan et al., 2021) and QAFACTEVAL (Fabbri et al., 2022a), are also used for evaluation.

3.3 Implementation Details

We choose BART (Lewis et al., 2020) as the backbone seq2seq model as it has demonstrated better dialogue summarization performance than other pre-trained language models (Tang et al., 2022), such as PEGASUS (Zhang et al., 2020a) and T5 (Raffel et al., 2020). The proposed models are optimized using AdamW (Loshchilov and Hutter, 2019) with learning rate $3e-5$ and weight decay $1e-3$. The maximum input sequence length is set to 1024. For all baseline models, we use the best hyper-parameters reported in their papers. We fix τ to be 0.5 throughout all our experiments. α and β are both 1.0.

3.4 Baselines

We compare SWING with the following competitive baseline systems. **TextRank** (Mihalcea and Tarau, 2004) is a graph-based ranking algorithm that performs extractive summarization. **BART** (Lewis et al., 2020) is a seq2seq language model pre-trained on various denoising objectives. **CTRLDIASUMM** (Liu and Chen, 2021) and **CODS** (Wu et al., 2021) are controllable generation frameworks that generate summaries guided by named entity

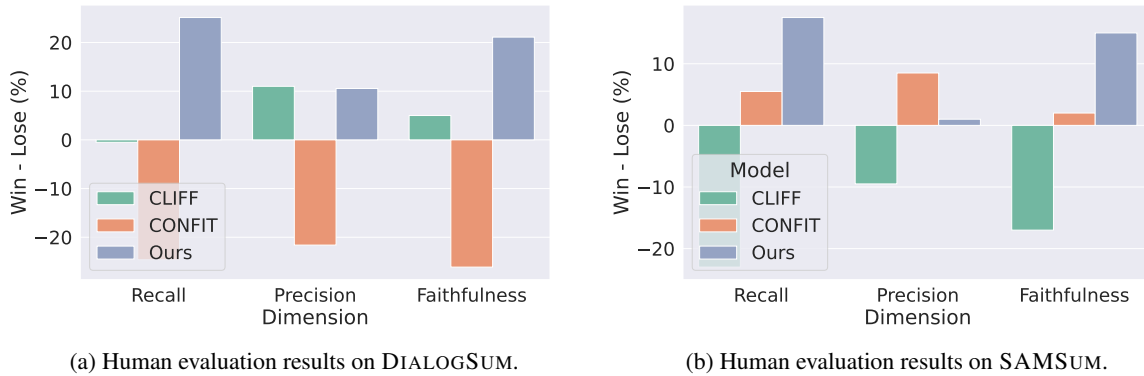


Figure 3: Human evaluation results. SWING achieves the highest RECALL and FAITHFULNESS scores on both datasets, suggesting the advantages of our approach in reducing missing information and improving the overall faithfulness of the generated dialogue summary.

planning and sketches, respectively. **CONSEQ** (Nan et al., 2021) learns a contrastive objective based on unlikelihood training, where positive and negative samples are selected by QUALS. **CLIFF** (Cao and Wang, 2021) and **CONFIT** (Tang et al., 2022) are trained with a similar contrastive learning loss that takes the form of the InfoNCE loss (Oord et al., 2018), except that CONFIT is optimized with an additional self-supervised loss that aims to reduce reference errors. BART-LARGE is used across all experiments that involve pre-trained language models for fair comparison.

4 Results

4.1 Main results

Table 1 summarizes the main results on DIALOGSUM and SAMSUM. SWING outperforms previous approaches in almost all metrics, especially recall measures. This result reflects that the proposed approach generates summaries that cover more content in the reference summaries lexically and semantically. One interesting observation was the deficient performance of CONSEQ on both datasets. We hypothesize that poor performance was the use of the unlikelihood training objective in their loss, as mentioned in §2.3. Since sentences of dialogue summaries often share similar structures, adopting such an objective could confuse the model. We verified this hypothesis by running a small experiment by training BART-LARGE with MLE and negative samples determined by QUALS, similar to CONSEQ. The resulting model also produces significantly lower performance than training with MLE alone. The finding confirms that the poor performance of CONSEQ is caused by the unlikelihood training and that such a loss function is unsuitable for dialogue summarization.

4.2 Human Evaluation

To further validate the effectiveness of SWING, we use Amazon’s Mechanical Turk (AMT) to recruit workers to conduct human evaluations on three methods: CLIFF, CONFIT and SWING. We sampled 100 dialogues from the test set of DIALOGSUM and SAMSUM, respectively. For each dialogue, human judges are presented with a pair of summaries produced by two different approaches and asked to select the better one with respect to three dimensions. **RECALL** assesses the portion of information in the reference summary covered by the generated summary. **PRECISION** considers whether all the content in the generated summary occurs in the reference summary. **FAITHFULNESS** examines whether the generated summary is factually consistent with the dialogue. "Tie" is selected if the judges consider the two summaries to be of equal quality. The final score of each system is calculated as the percentage of times the system is selected as the better one minus the percentage of times the system is not. To evaluate the annotation quality, we compute the inter-annotator agreement. The average Cohan’s Kappa (Cohen, 1960) is 54.35%, indicating a moderate agreement. Details of the human evaluation setup can be found in Appendix B.

The human evaluation results are demonstrated in Figure 3. We have the following observations. First, SWING achieves the highest RECALL scores on both datasets, indicating that our approach is the best in addressing the missing information issue for dialogue summarization. Second, while SWING does not score the highest on PRECISION, we achieve the highest scores on FAITHFULNESS. This implies that even though our approach often generates summaries with extra information,

Reference Summary	CONFIT	SWING
Mike took his car into garage today. Ernest is relieved as someone had just crashed into a red Honda which looks like Mike’s.	Mike took his car to the garage today. Someone crashed into his car.	Mike took his car into the garage today. Someone just crashed into a red Honda looking like Mike’s.
Hilary has the keys to the apartment. Benjamin wants to get them and go take a nap. Hilary is having lunch with some French people at La Cantina. Hilary is meeting them at the entrance to the conference hall at 2 pm. Benjamin and Elliot might join them. They’re meeting for the drinks in the evening.	Benjamin, Elliot, Daniel and Hilary will meet at La Cantina at 2 pm to have lunch with some French people who work on the history of food in colonial Mexico. They will try to avoid talking about their subject of research.	Hilary has the keys to Benjamin, Elliot and Daniel’s apartment. They will meet at the entrance to the conference hall at 2 pm and go to La Cantina for lunch with some French people who work on the history of food in colonial Mexico.

Table 2: Qualitative analysis on the outputs of SWING and CONFIT. The two rows demonstrate the *missing details* and the *missing sentences* issue of the summaries generated by CONFIT, respectively. The extra information in the outputs of CONFIT that also occurs in the reference summaries is highlighted in blue. In both cases, SWING is able to cover more content presented in the reference summaries.

the additional content is likely still faithful to the input. To measure the amount of additional information produced, we compute the average number of tokens per summary for each model. As seen in Table 3, the summaries generated by SWING is only slightly longer than those produced by CLIFF and CONFIT. This suggests that SWING achieves significantly higher faithfulness and coverage than CLIFF and CONFIT while maintaining conciseness.

Model	DIALOGSUM	SAMSUM
CONFIT	29.46	22.45
CLIFF	27.34	22.30
BART-LARGE	28.03	23.19
SWING	31.32	24.23

Table 3: Average token count per summary generated by different models.

4.3 Qualitative Analysis

To provide better insight into the effectiveness of the proposed method, we conduct a qualitative analysis using the 100 dialogues randomly sampled from the SAMSUM dataset. Specifically, we further categorize missing information errors into two sub-types: (1) *missing details* where partial information of a sentence in the reference summary is missing in the generated summary and (2) *missing sentences* where the model fails to generate an entire sentence in the reference summary. An example of each sub-type is shown in Table 2. By comparing the test sets outputs of CONFIT and SWING, we see that there are 10 improved cases with less *missing details* and 6 cases where *missing sentences* is mitigated by SWING. Meanwhile, our

proposed approach only introduces *missing details* error and *missing sentences* error in 1 and 2 examples, respectively. This implies that our approach is effective in alleviating both sub-types of missing information error while particularly advantageous in reducing *missing details* errors.

4.4 Correlation with Human Judgements

Although recently proposed metrics have been shown to be highly correlated with human judgements on news summarization in terms of factuality (Kryscinski et al., 2020; Yuan et al., 2021), no previous work has studied the transferability of these metrics to dialogue summarization. We seek to answer this question by computing the correlation of the automatic metrics in Table 1 with the human annotations discussed in §4.2. Using Kendall’s Tau (Kendall, 1938) as the correlation measure, the results are summarized in Table 4. We observe that: (1) $BARTSCORE_R$ is the most consistent and reliable metric across the three dimensions. It performs the best in RECALL on both datasets, indicating that $BARTSCORE_R$ is most suitable for measuring how well a model resolves the missing information issue in dialogue summarization. (2) Although a large number of invalid questions and answers are generated, QUALS is the best metric for assessing PRECISION overall. (3) $FACTCC_F$ and $FACTCC_R$ are two of the worst metrics in general. This could be explained by the fact that FACTCC constructs negative samples with some semantically variant transformations. However, these transformations may not be comprehensive enough to cover all cases. Hence, the poor transferability of FACTCC on these two datasets.

Metric	DIALOGSUM			SAMSUM		
	RECALL	PRECISION	FAITHFULNESS	RECALL	PRECISION	FAITHFULNESS
ROUGE-L _F	23.50	24.21	10.29	6.07	10.24	-0.75
ROUGE-L _R	23.46	2.51	4.24	29.52	9.61	17.88
BARTSCORE _F	18.35	25.94	3.17	15.50	8.00	10.69
BARTSCORE _R	26.48	14.87	9.25	32.10	9.68	24.11
FACTCC _F	6.15	6.93	1.19	-3.43	5.12	-2.28
FACTCC _R	4.79	6.86	10.56	4.13	10.32	-1.43
QUALS	14.23	23.61	-0.83	1.55	15.35	4.50
QAFACTEVAL	14.06	16.20	16.80	5.03	2.83	6.26

Table 4: Correlation (%) of automatic metrics with human judgements. We first convert human evaluation results and automatic metric scores into a scale of $\{-1, 0, 1\}$, which corresponds to $\{\text{LOSE, TIE, WIN}\}$. Then, Kendall’s Tau (Kendall, 1938) is used to compute the correlation between two sequences.

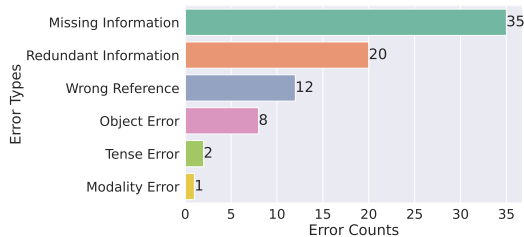


Figure 4: Remaining challenges.

4.5 Remaining Challenges

We analyzed the remaining errors by comparing 100 generated summaries with corresponding reference summaries on the SAMSUM datasets using the categories of factual errors defined in Tang et al. (2022). The results are shown in Figure 4. We observe that missing information still accounts for the largest portion of factual errors, even though our approach significantly exceeds prior methods in mitigating this issue. This reflects that this issue is challenging to tackle and that there is still a great opportunity to improve the reduction of missing information. As a comparison, we manually inspected outputs of BART-LARGE using the same 100 dialogues as input. We found 42 cases where information is missing from the dialogue summaries produced by BART-LARGE. This observation further confirms the effectiveness of SWING in addressing insufficient coverage. In addition, redundant information is another major source of errors. Although we have shown in §4.2 that the additional information generated by SWING is likely still faithful to the input dialogue, compactness is one of the important qualities of a summary. This can be improved by using NLI to guide the model to avoid generating extra information. Other common mistakes are wrong reference and object errors, both of which can be addressed with the self-supervised

loss discussed in Tang et al. (2022).³

5 Related Work

Dialogue Summarization Early work on dialogue summarization focus on the AMI meeting corpus (McCowan et al., 2005) due to the lack of dialogue summarization data. These studies enhance summarization performance by leveraging features of conversational data, such as dialogue act (Goo and Chen, 2018), visual features (Li et al., 2019), and the relationships between summary and dialogue (Oya et al., 2014). Later, Gliwa et al. (2019) released the SAMSUM dataset, the first large-scale dialogue summarization dataset, enabling abstractive summarization research on casual chat dialogue. With the rise of large language models (LMs), recent work focuses on improving the controllability of sequence-to-sequence models built upon large LMs. For instance, Wu et al. (2021) proposes to utilize a summary sketch to control the granularity of the summary generated. Liu and Chen (2021) conditions the generators with person name entities to control which people to include in the generating summary. Chan et al. (2021) improves controllability by formulating the summarization task as a constrained Markov Decision Process.

Factual Consistency Enhancement While factuality has been widely explored in the field of fact-checking and fake news detection (Thorne et al., 2018; Wadden et al., 2020; Huang et al., 2022b; Shu et al., 2018; Pan et al., 2021; Huang et al., 2022a), factual inconsistency remains a major challenge for abstractive summarization. One line of work attempts to improve the faithfulness of

³This analysis is not comparable to results reported in Tang et al. (2022) due to differences in the sampled examples.

the generated summary with a separate correction model that corrects the errors made by the summarization model (Dong et al., 2020; Cao et al., 2020; Fabbri et al., 2022b) or directly fix factual inconsistencies in the training data (Adams et al., 2022). Another line of work employs auxiliary loss functions to improve models’ representations or discourage the model from generating unfaithful outputs (Cao and Wang, 2021; Chen et al., 2021a; Nan et al., 2021; Tang et al., 2022). The main advantage of these approaches is the efficiency in inference time.

Some studies have attempted to use NLI to detect factual inconsistency in generated summaries. Early approaches rely on out-of-the-box NLI models, which did not yield satisfactory results (Falke et al., 2019). Barrantes et al. (2020) improved the detection accuracy by using an NLI model fine-tuned on the Adversarial NLI dataset (Nie et al., 2020). Laban et al. (2022) addresses the mismatch issue in input granularity between NLI datasets and inconsistency detection by passing sentence pairs as inputs instead of document-summary pairs. Kryscinski et al. (2020) and Yin et al. (2021) trains document-sentence entailment models to address the granularity mismatch issue. Utama et al. (2022) introduces a controllable generation framework that generates document-level NLI training data for identifying factual inconsistency. Our work leverages an NLI model to guide the dialogue summarization model to recover missing information.

6 Conclusion

We have proposed SWING, a dialogue summarization framework that generates summaries with mitigated missing information and improved faithfulness. To instruct the model to generate missing content from the reference summaries and to differentiate factually consistent generated sentences from their inconsistent counterparts, we propose two losses based on NLI. Experimental results on the DIALOGSUM and SAMSUM datasets showed that our approach achieves significantly higher faithfulness and coverage, while still maintaining conciseness, compared to prior methods. In addition, we measure the correlation between the reported automatic metrics and human judgments to provide insight into the most suitable metric for evaluating the coverage and factuality of dialogue summaries for future research.

7 Ethical Considerations

We acknowledge that the use of large language models pre-trained on the Web could lead to biased outputs. We did find out that our model may sometimes generate the incorrect pronouns for neutral names. For example, in Figure 1, Charlee is being referred to as a male in the generated summary, while Charlee is actually a female as shown in the reference summary. Such an issue is often caused by under-specified context (e.g. Charlee’s gender is not mentioned in the input dialogue). Fortunately, we found that such an error accounts for < 1% of the total outputs from our framework and the issue can be largely alleviated when enough context is provided.

8 Limitations

While our proposed approach is effective in mitigating missing information, this issue is still far from resolved, as shown in Figure 4. Significant effort is needed to ensure dialogue summarization models produce completely factual content. In addition, our method works as we found that most of the reference summaries in the two datasets we used are faithful to the corresponding dialogue. The proposed method may not work on other summarization datasets, such as XSum, which contains hallucinations in about 70% of the reference summaries (Maynez et al., 2020).

9 Acknowledgments

We would like to extend our gratitude to the reviewers for their valuable feedback and insights, which greatly contributed to the improvement of this paper. We would also like to thank the human evaluators for their time and effort in assessing the performance of our model. Their contributions have been essential in ensuring the quality of our research.

References

- Griffin Adams, Han-Chin Shing, Qing Sun, Christopher Winestock, Kathleen McKeown, and Noémie Elhadad. 2022. Learning to revise references for faithful summarization. *arXiv preprint arXiv:2204.10290*.
- Mario Barrantes, Benedikt Herudek, and Richard Wang. 2020. Adversarial nli for factual correctness in text summarisation models. *ArXiv*, abs/2005.11739.
- Meng Cao, Yue Dong, Jiapeng Wu, and Jackie Chi Kit Cheung. 2020. Factual error correction for abstractive summarization models. In *Proceedings of the*

- 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), pages 6251–6258, Online. Association for Computational Linguistics.
- Shuyang Cao and Lu Wang. 2021. **CLIFF: Contrastive learning for improving faithfulness and factuality in abstractive summarization**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6633–6649, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Hou Pong Chan, Lu Wang, and Irwin King. 2021. **Controllable summarization with constrained Markov decision process**. *Transactions of the Association for Computational Linguistics*, 9:1213–1232.
- Wang Chen, Piji Li, Hou Pong Chan, and Irwin King. 2021a. **Dialogue summarization with supporting utterance flow modelling and fact regularization**. *Knowl. Based Syst.*, 229:107328.
- Yulong Chen, Yang Liu, Liang Chen, and Yue Zhang. 2021b. **DialogSum: A real-life scenario dialogue summarization dataset**. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 5062–5074, Online. Association for Computational Linguistics.
- Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and psychological measurement*, 20(1):37–46.
- Yue Dong, Shuohang Wang, Zhe Gan, Yu Cheng, Jackie Chi Kit Cheung, and Jingjing Liu. 2020. **Multi-fact correction in abstractive text summarization**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9320–9331, Online. Association for Computational Linguistics.
- Alexander Fabbri, Chien-Sheng Wu, Wenhao Liu, and Caiming Xiong. 2022a. **QAFactEval: Improved QA-based factual consistency evaluation for summarization**. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2587–2601, Seattle, United States. Association for Computational Linguistics.
- Alexander R Fabbri, Prafulla Kumar Choubey, Jesse Vig, Chien-Sheng Wu, and Caiming Xiong. 2022b. **Improving factual consistency in summarization with compression-based post-editing**. *arXiv preprint arXiv:2211.06196*.
- Tobias Falke, Leonardo F. R. Ribeiro, Prasetya Ajie Utama, Ido Dagan, and Iryna Gurevych. 2019. **Ranking generated summaries by correctness: An interesting but challenging application for natural language inference**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2214–2220, Florence, Italy. Association for Computational Linguistics.
- Ronald Aylmer Fisher et al. 1937. The design of experiments. *The design of experiments.*, (2nd Ed).
- 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.
- 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 2018, Athens, Greece, December 18-21, 2018*, pages 735–742. IEEE.
- Kung-Hsiang Huang, Kathleen McKeown, Preslav Nakov, Yejin Choi, and Heng Ji. 2022a. **Faking fake news for real fake news detection: Propaganda-loaded training data generation**. *arXiv preprint arXiv:2203.05386*.
- Kung-Hsiang Huang, ChengXiang Zhai, and Heng Ji. 2022b. **CONCRETE: Improving cross-lingual fact-checking with cross-lingual retrieval**. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 1024–1035, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Maurice G Kendall. 1938. A new measure of rank correlation. *Biometrika*, 30(1/2):81–93.
- Wojciech Kryscinski, Bryan McCann, Caiming Xiong, and Richard Socher. 2020. **Evaluating the factual consistency of abstractive text summarization**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9332–9346, Online. Association for Computational Linguistics.
- Philippe Laban, Tobias Schnabel, Paul N. Bennett, and Marti A. Hearst. 2022. **SummaC: Re-visiting NLI-based models for inconsistency detection in summarization**. *Transactions of the Association for Computational Linguistics*, 10:163–177.
- 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.
- Manling Li, Lingyu Zhang, Heng Ji, and Richard J. Radke. 2019. **Keep meeting summaries on topic: Abstractive multi-modal meeting summarization**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2190–2196, Florence, Italy. 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.
- Zhengyuan Liu and Nancy Chen. 2021. [Controllable neural dialogue summarization with personal named entity planning](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 92–106, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *International Conference on Learning Representations*.
- Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. 2020. [On faithfulness and factuality in abstractive summarization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1906–1919, Online. Association for Computational Linguistics.
- I. McCowan, J. Carletta, W. Kraaij, S. Ashby, S. Bourban, M. Flynn, M. Guillemot, T. Hain, J. Kadlec, V. Karaiskos, M. Kronenthal, G. Lathoud, M. Lincoln, A. Lisowska, W. Post, Dennis Reidsma, and P. Wellner. 2005. The ami meeting corpus. In *Proceedings of Measuring Behavior 2005, 5th International Conference on Methods and Techniques in Behavioral Research*, pages 137–140. Noldus Information Technology.
- Rada Mihalcea and Paul Tarau. 2004. [TextRank: Bringing order into text](#). In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, pages 404–411, Barcelona, Spain. Association for Computational Linguistics.
- Feng Nan, Cicero Nogueira dos Santos, Henghui Zhu, Patrick Ng, Kathleen McKeown, Ramesh Nallapati, Dejian Zhang, Zhiguo Wang, Andrew O. Arnold, and Bing Xiang. 2021. [Improving factual consistency of abstractive summarization via question answering](#). 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 6881–6894, Online. Association for Computational Linguistics.
- Yixin Nie, Adina Williams, Emily Dinan, Mohit Bansal, Jason Weston, and Douwe Kiela. 2020. [Adversarial NLI: A new benchmark for natural language understanding](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4885–4901, Online. Association for Computational Linguistics.
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Tatsuro Oya, Yashar Mehdad, Giuseppe Carenini, and Raymond Ng. 2014. [A template-based abstractive meeting summarization: Leveraging summary and source text relationships](#). In *Proceedings of the 8th International Natural Language Generation Conference (INLG)*, pages 45–53, Philadelphia, Pennsylvania, U.S.A. Association for Computational Linguistics.
- Artidoro Pagnoni, Vidhisha Balachandran, and Yulia Tsvetkov. 2021. [Understanding factuality in abstractive summarization with FRANK: A benchmark for factuality metrics](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4812–4829, Online. Association for Computational Linguistics.
- Liangming Pan, Wenhui Chen, Wenhan Xiong, Min-Yen Kan, and William Yang Wang. 2021. [Zero-shot fact verification by claim 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 2: Short Papers)*, pages 476–483, Online. 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.
- Kai Shu, Deepak Mahudeswaran, Suhang Wang, Dongwon Lee, and Huan Liu. 2018. Fakenewsnet: A data repository with news content, social context and dynamic information for studying fake news on social media. *arXiv preprint arXiv:1809.01286*.
- Xiangru Tang, Arjun Nair, Borui Wang, Bingyao Wang, Jai Desai, Aaron Wade, Haoran Li, Asli Celikyilmaz, Yashar Mehdad, and Dragomir Radev. 2022. [CONFIT: Toward faithful dialogue summarization with linguistically-informed contrastive fine-tuning](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5657–5668, Seattle, United States. Association for Computational Linguistics.
- Yuqing Tang, Chau Tran, Xian Li, Peng-Jen Chen, Naman Goyal, Vishrav Chaudhary, Jiatao Gu, and Angela Fan. 2020. [Multilingual translation with extensible multilingual pretraining and finetuning](#).
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. [FEVER: a large-scale dataset for fact extraction and VERification](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 809–819, New Orleans, Louisiana. Association for Computational Linguistics.

- Prasetya Utama, Joshua Bambrick, Nafise Moosavi, and Iryna Gurevych. 2022. [Falsesum: Generating document-level NLI examples for recognizing factual inconsistency in summarization](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2763–2776, Seattle, United States. Association for Computational Linguistics.
- David Wadden, Shanchuan Lin, Kyle Lo, Lucy Lu Wang, Madeleine van Zuylen, Arman Cohan, and Hannaneh Hajishirzi. 2020. [Fact or fiction: Verifying scientific claims](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7534–7550, Online. Association for Computational Linguistics.
- Sean Welleck, Ilya Kulikov, Stephen Roller, Emily Dinan, Kyunghyun Cho, and Jason Weston. 2020. [Neural text generation with unlikelihood training](#). In *International Conference on Learning Representations*.
- Chien-Sheng Wu, Linqing Liu, Wenhao Liu, Pontus Stenetorp, and Caiming Xiong. 2021. [Controllable abstractive dialogue summarization with sketch supervision](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 5108–5122, Online. Association for Computational Linguistics.
- Wenpeng Yin, Dragomir Radev, and Caiming Xiong. 2021. [DocNLI: A large-scale dataset for document-level natural language inference](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 4913–4922, Online. Association for Computational Linguistics.
- Weizhe Yuan, Graham Neubig, and Pengfei Liu. 2021. [BARTScore: Evaluating generated text as text generation](#). In *Advances in Neural Information Processing Systems*.
- Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter J. Liu. 2020a. Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. In *Proceedings of the 37th International Conference on Machine Learning, ICML'20*. JMLR.org.
- Yuhao Zhang, Derek Merck, Emily Tsai, Christopher D. Manning, and Curtis Langlotz. 2020b. [Optimizing the factual correctness of a summary: A study of summarizing radiology reports](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5108–5120, Online. Association for Computational Linguistics.

A Dataset Statistics

We present the detailed statistics of DIALOGSUM and SAMSUM in Table 5.

Dataset	# Dialogues	Avg. Dialogue Words	Avg. Summ. Words
DIALOGSUM	13,460	187.5	31.0
SAMSUM	16,369	124.1	23.4

Table 5: Statistics of DIALOGSUM and SAMSUM. We use the NLTK tokenizer to compute word counts for both datasets.

B Human Evaluation Details

In this section, we describe the details of our human evaluation. We recruit AMT workers from the United States for ensuring language fluency. Qualification requirements are set such that only workers who have an acceptance rate greater than 99% and have more than 10,000 accepted HITs in the past are allowed to work on our annotation task. To further ensure annotation quality, we conducted two rounds of annotations. In the first round, we launched 100 HITs to select high-quality annotators in the first round. 8 qualified annotators are selected to enter the second round to conduct the remaining evaluation. We set the reward to \$0.8 per HIT to encourage experienced annotators to participate. Our annotation interface is displayed in Figure 5.

For each HIT, annotators are provided with a piece of dialogue and a corresponding reference summary as well as two summaries generated from different systems, demonstrated on the left segment of the interface. Based on the summaries and the dialogue, annotators are tasked to answer three questions shown on the right segment of the interface, each of which corresponds to RECALL, PRECISION, and FAITHFULNESS. They need to determine which summary is better with regard to each prompt.

C Comparison with Other Data Augmentation Methods

We compared our MIXANDMATCH summary construction technique with other data augmentation methods, including back translation (BACKTRANSLATE) and paraphrasing (PARAPHRASING). For back translation, we use mBART-50 (Tang et al., 2020) to translate a summary from English to German and then back to English. For paraphrase generation, we use this open

source package⁴. The experimental results are summarized in Table 6. Training with MIXANDMATCH summaries achieves the highest scores on most metrics, indicating that our proposed method is the most effective in improving the factuality of the generated summaries.

D Hardware and Software configurations

All experiments are conducted on a Linux machine with NVIDIA V100. We use PyTorch 1.11.0 with CUDA 10.1 as the Deep Learning framework and utilize Transformers 4.19.2 to load all pre-trained language models.

E Validation Set Performance

We report the validation set performance of our proposed model in Table 7.

F Number of Parameters

We do not introduce additional parameters to the backbone language model, BART-LARGE. During training time, the number of parameters equals to the sum of the number of parameters in BART-LARGE and ROBERTA-LARGE. In inference time, since we do not need the NLI component, the number of parameters is the same as that of BART-LARGE.

G Scientific Artifacts

The licenses for all the models and software used in this paper are listed below in parentheses: BART (MIT License), FACTCC (BSD-3-Clause License), QAFACTEVAL (BSD-3-Clause License), BARTSCORE (Apache License 2.0), QUALS (MIT License), py-ROUGE (Apache License 2.0), NLTK (Apache License 2.0).

⁴<https://github.com/Vamsi995/Paraphrase-Generator>

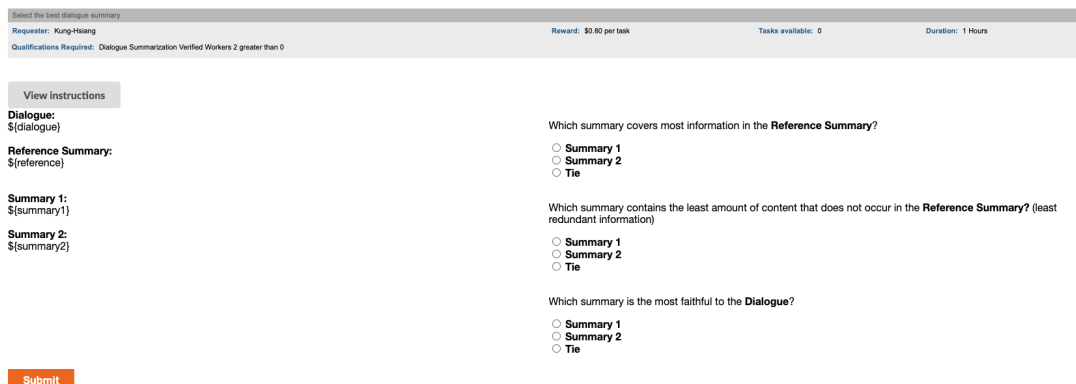


Figure 5: MTurk UI for our human evaluation.

Model	DIALOGSUM								SAMSUM							
	RL _F	RL _R	BS _F	BS _R	FC _F	FC _R	QS	QFE	RL _F	RL _R	BS _F	BS _R	FC _F	FC _R	QS	QFE
MIXANDMATCH	51.53	59.27	-2.012	-1.901	82.90	85.86	-1.130	2.399	49.73	53.95	-2.185	-2.143	63.47	63.15	-0.886	2.027
BACKTRANSLATE	50.41	58.22	-2.012	-2.032	83.20	84.23	-1.230	2.245	49.02	52.93	-2.234	-2.159	64.69	62.10	-1.230	1.984
PARAPHRASING	50.32	59.22	-2.133	-1.936	82.20	87.62	-1.198	2.333	49.23	53.94	-2.320	-2.178	64.78	63.98	-1.130	2.015

Table 6: Performance comparison on DIALOGSUM and SAMSUM with other positive data augmentation methods.

Model	DIALOGSUM						SAMSUM					
	RL _F	RL _R	BS _F	BS _R	FC _F	FC _R	RL _F	RL _R	BS _F	BS _R	FC _F	FC _R
SWING	48.45	51.27	-2.149	-2.169	71.36	70.65	50.61	53.74	-2.212	-2.134	64.27	64.56

Table 7: Validation set performance.