

# Analyzing and Evaluating Faithfulness in Dialogue Summarization

Bin Wang<sup>†</sup>, Chen Zhang<sup>†</sup>, Yan Zhang<sup>†</sup>, Yiming Chen<sup>†</sup>, Haizhou Li<sup>‡, †, §</sup>

<sup>†</sup>National University of Singapore, Singapore

<sup>‡</sup>The Chinese University of Hong Kong, Shenzhen, China

<sup>§</sup>Kriston AI, China

bwang28c@gmail.com

## Abstract

Dialogue summarization is abstractive in nature, making it suffer from factual errors. The factual correctness of summaries has the highest priority before practical applications. Many efforts have been made to improve faithfulness in text summarization. However, there is a lack of systematic study on dialogue summarization systems. In this work, we first perform the fine-grained human analysis on the faithfulness of dialogue summaries and observe that over 35% of generated summaries are faithfully inconsistent respective the source dialogues. Furthermore, we present a new model-level faithfulness evaluation method. It examines generation models with multi-choice questions created by rule-based transformations. Experimental results show that our evaluation schema is a strong proxy for the factual correctness of summarization models. The human-annotated faithfulness samples and the evaluation toolkit are released to facilitate future research toward faithful dialogue summarization. Code available: <https://github.com/BinWang28/FacEval>.

## 1 Introduction

Text summarization aims to condense a document into a short paragraph or a single sentence while conveying the core information (El-Kassas et al., 2021). It can be either extractive or abstractive. Extractive summarization methods identify salient sequence spans from the source document and pasting them together (Dorr et al., 2003; Kobayashi et al., 2015; Zhong et al., 2020). Abstractive summarization methods generate completely new summary in a coherent manner (Paulus et al., 2018; Lewis et al., 2020; Zhang et al., 2020; Liu et al., 2021a; Wang et al., 2022). Previous work discovered that abstractive summarization suffers from unfaithful outputs, limiting its applicability in real-world scenarios (Kryscinski et al., 2020; Falke et al., 2019; Zhu et al., 2021a; Ladhak et al., 2022).

<b>Dialogue:</b> <i>Freddie:</i> Nanna, are you coming to visit us soon? <i>Winnie:</i> Oh darling, Nanna has broken her leg, you'll have to visit me instead. <i>Freddie:</i> I forgott. Well come soon. <i>Winnie:</i> Good, ask Mummy and Daddy and they will come when they can. <i>Freddie:</i> Yes love you. Leg better soon? <i>Winnie:</i> Yes, quite soon. Tell mummy to ring me. Bye darling xxxxx
<b>Summaries:</b> <i>Human:</i> Winnie has broken her leg and will not visit any time soon. Freddie will ask mummy to call Winnie up. ✓ <i>BART:</i> Nanna has broken her leg, so Freddie will have to visit her instead. Nanna will get better soon. ✓ <i>MV-BART:</i> Nanna has broken her leg and Freddie will have to visit Winnie instead. Mummy and Daddy will come to visit them soon. ✗ <i>Coref-BART:</i> Freddie wants to visit Winnie, but Nanna has broken her leg, so he will have to visit her instead. Mummy and Daddy will come when they can. ✗ <i>CondigSum-BART:</i> Winnie's Nanna has broken her leg and Freddie will have to visit her instead. ✗

Table 1: A real example from SAMSum dataset. Span of factual errors are marked with underline.

As an essential way of exchanging information, conversations usually involve multiple participants, informal language usage, repetition, and negotiations (Sacks et al., 1978; Chen and Yang, 2020). Therefore, dialogue summarization is vulnerable to factual issues due to its abstractive nature. Table 1 gives an example of factually incorrect dialogue summaries. The problem of factual correctness is broadly studied for text summarization in news and article domains (Nallapati et al., 2016; Narayan et al., 2018). The progress is primarily because of the availability of factually annotated data at both summary and token levels (Kryscinski et al., 2020; Wang et al., 2020; Pagnoni et al., 2021; Cao et al., 2022). Many studies are proposed to evaluate and reduce factual errors in the generated summaries. However, due to the interactive nature of dialogues, we cannot simply transfer these methods to dialogue summarization.

In this work, we first categorize the most frequently occurred factual errors for dialogue sum-

marization into 6 types. Then, we collect fine-grained factual annotations for human reference and the output of 4 recent dialogue summarization systems (§3). At least two annotators are involved, and a verification process is incorporated to ensure the annotation quality. As a result, our study on human-annotated data suggests that over 35% of the generated dialogue summaries contain at least one factual error. Similar observations have been made in the news summarization domain where 30%-80% of generated text are factually inconsistent (Cao et al., 2018; Pagnoni et al., 2021). More research attention should be made toward faithful dialogue summarization.

The unavailability of faithful evaluation methods hinders the development of effective dialogue summarization models. In this work, we present a model-level evaluation schema, FacEval, targeting dialogue summarisation models’ faithfulness (§4). First, we synthesize a set of positive and negative summaries for each dialogue with back-translation or rule-based transformations. Then, a summarization model is asked to distinguish positive and negative summaries based on conditional generation probabilities. More correct judgements indicate the model is more factually competent.

To compare the model-level performance of evaluation methods, we leverage two *ad-hoc* training schema to synthesize a series of models with different capability ranks. Then, the evaluation methods are used to predict the ranking of trained models. Seven non-factual and factual evaluation methods have been examined, followed by a detailed discussion of their properties. The effectiveness of FacEval is also proven by showing a strong correlation with the factual correctness of summarization models.

## 2 Related Work

### 2.1 Summarization Methods

Text summarization is one of the most important tasks in natural language generation (NLG). With the development of pre-trained language models, a lot progress has been made to abstractive text summarization (See et al., 2017; Zhang et al., 2020; Liu et al., 2022), especially in news domain (Hermann et al., 2015; Narayan et al., 2018). With the availability of datasets (Carletta et al., 2005; Gliwa et al., 2019; Zhu et al., 2021b), dialogue summarization research has attracted a lot of attention. For dialogue summarization, fine-tuning pre-trained gen-

eration models including T5 (Raffel et al., 2020), PEGASUS (Zhang et al., 2020) and BART (Lewis et al., 2020) are served as a strong baseline, where BART achieves the SOTA performance on ROUGE scores. Some recent works consider the dialogue properties for more advanced summarization models. Chen and Yang (2020) and Liu et al. (2021a) incorporate the conversational structures into the semantic encoding process of dialogue. Conversations involve lots of co-references. Therefore, Liu et al. (2021b) proposes injecting co-reference information into the transformer layers by adapting attention maps or through graph convolutional networks (GCN). We include the outputs of recent dialogue summarization models in our analysis.

### 2.2 Faithfulness Analysis

Previous works spot that the factual consistency problem is one key aspect of improving text summarization (Kryscinski et al., 2020; Cao et al., 2020). The analysis of factual errors in summaries is mainly performed in the news domain. Kryscinski et al. (2019) and Falke et al. (2019) conducted the initial crowdsourcing of binary factual annotations and found that nearly 30% of the generated summaries are factually inconsistent. Recent extensions focus on more fine-grained analysis (Cao and Wang, 2021; Pagnoni et al., 2021) and also discovering factual evidences at entity level (Cao et al., 2022) or span level (Huang et al., 2020; Maynez et al., 2020; Goyal and Durrett, 2021).

Recently, CONFIT presented the first study on the faithfulness of dialogue summaries (Tang et al., 2022b). Similar to our work, they also define a taxonomy of factual errors and conduct fine-grained annotations. However, they focus on comparing reference summaries and generated summaries without referring to the whole dialogue. It is sub-optimal because the reference summary cannot fully represent the entire dialogue and also can be incorrect according to our analysis in Section 3. Besides, the missing and redundant information is categorized as factual errors, which we consider less proper. More recent advanced dialogue summarization models are also not included in their analysis.

### 2.3 Faithfulness Evaluation

The default evaluation metric for summarization, ROUGE, is based on n-gram overlaps between a generated summary and the corresponding references, rendering it less sensitive for capturing fac-

Speaker 1: Fiona	Speaker 2: Jonathan
<i>What should I prepare 4 my dad's birthday?</i>	<i>How old is he?</i>
<i>Turning 50.</i>	<i>Wow, a round birthday, it must be sth big.</i>
<i>I know, but I don't have any idea.</i>	<i>What does he like?</i>
<i>He watches a lot of military movies.</i>	<i>Well, a movie ticket is probably not what you thought of.</i>
<i>No, not even close.</i>	<i>U said he likes military... maybe paintball?</i>
<i>I don't know how my mum will react but I like it :D</i>	

<b>Ref. Summary:</b>	Fiona doesn't know what she should give to her dad as a birthday gift. He likes military. Jonathan suggests a paintball match.
<b>SubObjE:</b>	Jonathan doesn't know what she should give to her dad as a birthday gift. He likes military. Jonathan suggests a paintball match.
<b>ProE:</b>	Fiona doesn't know what he should give to her dad as a birthday gift. He likes military. Jonathan suggests a paintball match.
<b>NegE:</b>	Fiona doesn't know what she should give to her dad as a birthday gift. He hates military. Jonathan suggests a paintball match.
<b>ParE:</b>	Fiona doesn't know what she should give to her dad as a Christmas gift. He likes military. Jonathan suggests a paintball match.
<b>HalE:</b>	Fiona doesn't know what she should give to her dad as a birthday gift. He likes military. Jonathan invites Fiona to watch a military movie.

Table 2: An illustration of the taxonomy on factual error types.

tual errors. Therefore, several new metrics are proposed to evaluate the faithfulness in the news domain (Kryscinski et al., 2019; Fabbri et al., 2021; Tang et al., 2022a). There are two major groups, one is based on natural language inference, and the other is based on question-answering. Kryscinski et al. (2020) and Goyal and Durrett (2020) propose to leverage entailment relationship. Scialom et al. (2021) and Wang et al. (2020) involves question generation, answer generation and answer-overlap as the factual consistency measure. Zhao et al. (2021) proposes to evaluate the faithfulness of task-oriented dialogue summarization by calculating the amount of overlapped dialogue states, which requires additional human annotations.

### 3 Fine-grained Faithfulness Analysis

Previous studies of factuality analysis in summarization mainly focus on the news domain. The typology of factual errors for dialogues can be very different. Therefore, we first define a taxonomy of frequently occurred factual errors for dialogue summaries. A fine-grained analysis is then performed by measuring the factual consistency within dialogue summary pairs.

#### 3.1 Taxonomy of Factual Errors

We collect the generated summaries using four SOTA dialogue summarization models on the popular dialogue summarization dataset, SAMSum (Gliwa et al., 2019). The selected models are BART (Lewis et al., 2020), MV-BART (Chen and Yang, 2020), Coref-BART (Liu et al., 2021b) and CondigSum-BART (Liu et al., 2021a). We define five most frequently occurred error types in dialogue summaries as below. An example for each error type is shown in Table 2.

**Subject Object Error (SubObjE):** The subject(s)

or object(s) involved for an event is (partially) wrong. It includes substitution, addition and deletion of any related subject(s) or object(s).

**Pronoun Error (ProE):** Pronoun references are frequently occurred in dialogue summarization. This error includes wrong references and ambiguous ones that cannot be fully understandable relying on the summary.

**Negation Error (NegE):** Dialogues can contain confirmation utterances. This error means that the generated summary makes wrong conclusions when contradictory or unconfirmed events are presented in the dialogue.

**Particulars Error (ParE):** The summary presents related events, but some details are inaccurate or faulty. It can include incorrect information like date, time and location.

**Hallucination Error (HalE):** Generation models have the imaginary ability and can be triggered by certain prompt words in the dialogue. The hallucination error refers to the cases where the summary contains events not presented in the dialogue.

**Other Error (OthE):** It is used to classify factual errors that do not belong to any of the above types.

Note that the above-mentioned error types are not exclusive to each other. That is, one summary may contain multiple error types.

#### 3.2 Annotation Process

We random sample 150 dialogues from the test set of SAMSum. Five summaries are listed for each dialogue, including the human-written one and four model-generated summaries.

Falke et al. (2019) finds that it needs at least 12 annotations to reach an inter-annotator agreement of coefficient  $k = 0.75$ , which can lead to high annotation costs and unreliable results with fewer annotators (Kryscinski et al., 2020). There-

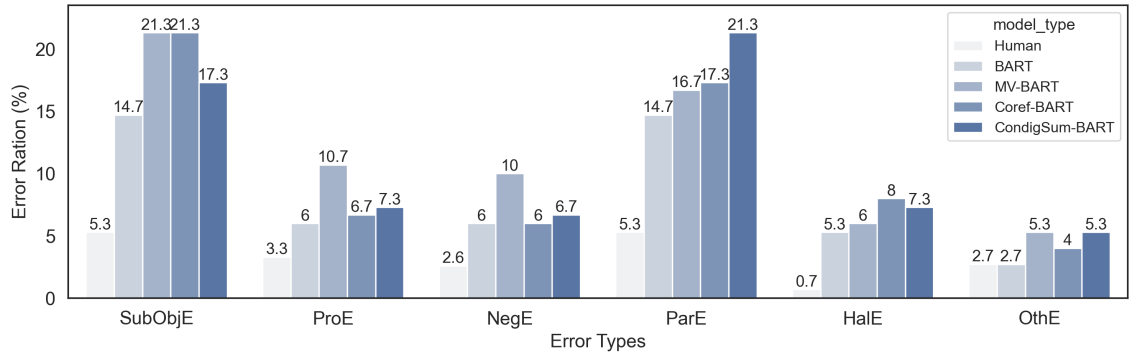


Figure 1: The proportion of summaries with different types of factual errors. Note that one summary can contain multiple error types.

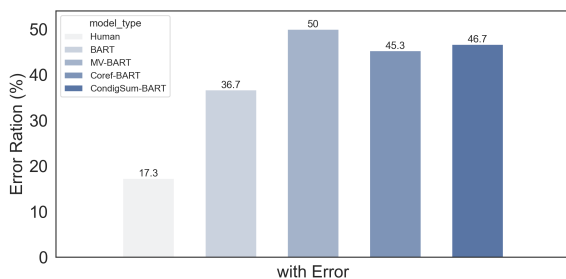


Figure 2: The proportion of summaries with at least one factual error.

fore, we perform a two-step verification process to ensure the annotation quality. First, each sample is annotated by two distinct annotators. If there is a disagreement about whether a summary contains factual errors, a third annotator is involved in making the final decision while considering inputs from the previous two annotators. As a result, we have collected 750 fine-grained faithfulness annotations from 30 participants.

### 3.3 Results and Analysis

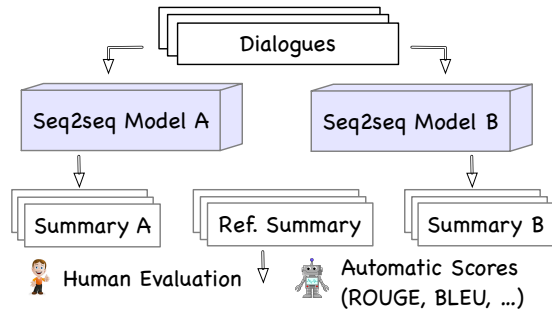
The detailed annotation results are shown in Figure 1. There are several exciting findings: 1) the human annotations contain non-negligible factual errors at around 17%; 2) 36% to 50% of generated summaries from dialogue summarization models contain at least one factual error; 3) three advanced dialogue summarization models perform worse than their baseline on factual consistency.

First, the popular SAMSum dataset (Gliwa et al., 2019) associates each dialogue with one human-written reference summary. However, we found that 17% of reference summaries have factual errors. Therefore, we encourage people to be aware of the issue, especially for evaluation. It is because


the dialogue annotation process for SAMSum only involved one annotator per sample, and no further verification process was executed. We notice that the source of factual errors for human summaries is also different from machine-generated ones. Some factual errors in human-written summaries are caused by typos, which rarely occur in machine-generated summaries.

For dialogue summarization models, we found that 35%-50% of generated summaries contain factual errors. The most frequent error types are SubObjE and ParE. Because dialogue often involves scattered information exchange with multiple speakers in multiple turns, it is very challenging to accurately locate *who* and *whom* in *who-did-what-to-whom*. That is the leading cause of SubObjE. ParE is the second most frequent error type, indicating that the generated summaries express the same topic but do not accurately capture the details. OthE occurs less frequently. It shows that our taxonomy of factual errors can cover the most frequent error types for dialogue summarization.

Surprisingly, we found that MV-BART, Coref-BART and CondigSum-BART perform even worse than the baseline model, with an increase of around 10% overall factual error rate. They are accepted as more advanced summarization models and perform better on ROUGE scores. It indicates that enhancing topical information is not necessarily contributing much to factuality (Chen and Yang, 2020; Liu et al., 2021a). Coref-BART aims to improve BART with co-reference information (Liu et al., 2021b). However, our result shows it does not bring obvious benefits. In conclusion, we encourage the future development of summarization models to pay more attention to the factuality perspective, and a more diverse evaluation schema



(a) Sample-level evaluation schema.



(b) Model-level evaluation schema.

Figure 3: An illustration of two types of evaluation paradigms.

beyond ROUGE scores should be incorporated.

## 4 Model-level Faithfulness Evaluation

Some efforts have been made toward sample-level factual error evaluation. An example is shown in Figure 3. The sample-level evaluation methods are model-agnostic and examine a model solely based on its output sequences. Most existing evaluation methods, including ROUGE score, human evaluation and recent factual evaluation methods, belong to this type. One ultimate goal for factuality evaluation is to discriminate better summarization models. We propose directly probing models’ generation probability with a constrained search space. First, FacEval generates a set of positive and negative samples with variant factual errors by rule-based transformations. Then, the generation probabilities of positive and negative summaries are compared for each dialogue. A better summarization model should be more likely to generate positive summaries than negative ones.

### 4.1 Dialogue-summary Pair Generation

We design transformations to synthesize negative samples with factual errors. Given the source and target text, one or more modifications are performed to the target text while referring to the information of the source text. It is because the frequently occurred errors are conceptual confusions from the source. Our designed transformations are listed as follows:




Figure 4: An illustration of comparing the generation probability of positive and negative samples. Solid and dashed lines refer to probability comparison and sample construction, respectively.

- **Speaker Swap (SS)**: We first spot the name of speakers from the source text by colon symbol and then swap the names at the target text.
- **Entity / Pronoun / Date / Number Swap (ES / PS / DS / NS)**: An NER system is first applied to both source and target text. The entities from the target text are randomly swapped with entities from the source text if they share the same entity type.
- **Negation (NG)**: Negation is performed using a set of hand-crafted rules. Auxiliary verbs are first scanned. Then, positive verbs are negated by adding not or n’t. Similarly, negative sentences are inverted by negation removal.

First, we paraphrase the summary to create more positive samples through back-translation (BT). The Google Cloud API is leveraged for this task<sup>1</sup>. Then, we generate new summaries with factual errors by corrupting positive summaries, which means the summaries are treated as the target text, and the dialogue is the source text.

Ideally, the negative summaries should be prone to errors generated in real-world scenarios. Therefore, our designed transformations try to mimic that. In the context of the analysis presented in Section 3, we have the following list of correspondences: 1) SS-SubObjE; 2) PS-ProE; 3) NG-NegE; 4) ES/DS/NS-ParE.

### 4.2 Comparison of Generation Probabilities

An illustration of probability comparison is shown in Figure 4. Given a dialogue  $D$ , a summary  $S = [y_1, \dots, y_L]$  and a summarization model  $f_s(\cdot)$ , we can compute a generation score ( $GS$ ) for  $D$ - $S$  pair

<sup>1</sup><https://cloud.google.com/translate/>

from the generation probability:

$$\begin{aligned} GS(S|D) &= \frac{1}{L^\alpha} \log P(y_1, \dots, y_L|D) \\ &= \frac{1}{L^\alpha} \sum_{t=1}^L \log P(y_t|y_1, \dots, y_{t-1}, D) \end{aligned}$$

where the generation probability for each token is as follows:

$$\begin{aligned} P(y_t|y_1, \dots, y_{t-1}, D) \\ = P(f_s(y_1, \dots, y_{t-1}, D) = y_t) \end{aligned}$$

We leverage the above generation score from decision process of beam search algorithm (Graves, 2012), where the sequence length is taken into consideration. In default, we set the length penalty parameter  $\alpha$  as 1.0.

For dialogue  $D_i$ , there is positive summary set  $S = [S_1, \dots, S_M]$  and negative summary set  $\hat{S} = [\hat{S}_1, \dots, \hat{S}_N]$ . We evaluate the number of times the positive samples have higher scores than the negative samples concerning the same dialogue. The factuality score ( $FS$ ) of model  $f_s(\cdot)$  is then computed as follows:

$$\begin{aligned} FS(f_s) &= \frac{1}{|D|} \sum_{i=1}^{|D|} \frac{1}{MN} \\ &\quad \sum_{m=1}^M \sum_{n=1}^N \mathbb{1}[GS(S_m|D_i) > GS(\hat{S}_n|D_i)] \end{aligned}$$

where  $|D|$  is the number of dialogues.

### 4.3 Evaluation Preparation

A series of models need to be prepared with different faithfulness capabilities to evaluate the effectiveness of model-level evaluation methods. One option is to collect as many well-trained models as possible and refer to human annotations to rank models based on factuality. However, it is hard to reach a high agreement and may not be trustworthy with limited annotators, as indicated by Falke et al. (2019) and Kryscinski et al. (2020). Therefore, instead, we construct a series of models using the following two ad-hoc methods:

**Limited data training (LDT).** One joint agreement is that more training data lead to better model performance. Therefore, we train 20 models using different proportions of the training data from 5% to 100%.

	# Diag	# Spk	# Turn	Sum. Len.
Train	14,732	2.40	11.17	23.4
Val	818	2.39	10.83	23.4
Test	819	2.36	11.25	23.1

Table 3: The detailed statistics of the SAMSum dataset. The header refers to the number of dialogues, the average number of speakers, the average number of dialogue turns, and the average summary lengths.

**Mixed data training (MDT).** In this setting, we randomly replace the human-labelled training samples with noisy ones. The noisy samples are created by corrupting only the dialogue using transformations introduced in Section 4.1. Here, the source and target are both the dialogue. The trained model is more likely to be confused and generate more factual errors with noisy data. Here, we obtain 21 models with different replacement ratios from 100% to 0%.

LDT will cause a model to be less competent for generation in all aspects. In comparison, MDT will lead the model to generate summaries with more factuality errors while less affecting other properties like fluency. Therefore, we expect a better factuality evaluator to correlate more with MDT models. All correlations are computed on model-level instead of sample-level judgements.

## 5 Experiments

### 5.1 Experimental Settings

SAMSum dataset (Gliwa et al., 2019) is used for all experiments. It consists of 16,369 dialogue-summary pairs written by expert linguistics. One human-written reference summary is provided for each dialogue. The detailed dataset statistics are listed in Table 3. The samples from the test set are used for all evaluation methods.

For backbone models, we exam with  $BART_{Large}$ ,  $BART_{Base}$  (Lewis et al., 2020),  $T5_{Base}$  and  $T5_{Small}$  (Raffel et al., 2020), which are SOTA summarization models. Each model is trained with both LDT and MDT methods. As a result, we obtained 164 trained models, divided into eight groups. The models in each group are associated with increasing levels of capabilities. The Spearman’s rank correlation coefficient ( $\rho$ ) between these models and evaluation scores is reported. For sample-based evaluation methods, the scores on all test set samples are averaged as the model-level performance. We ensure

Train. Strategy	Model	NG	PS	SS	ES	DS	NS	All
LDT	<i>BART<sub>Large</sub></i>	79.79	31.28	87.82	-18.65	77.89	-2.41	84.70
	<i>BART<sub>Base</sub></i>	89.62	75.64	98.65	73.83	40.84	-2.26	99.40
	<i>T5<sub>Base</sub></i>	63.16	46.62	93.98	50.68	73.38	7.98	89.62
	<i>T5<sub>Small</sub></i>	96.54	90.53	97.29	91.13	87.52	43.98	98.05
MDT	<i>BART<sub>Large</sub></i>	64.03	53.12	99.87	67.14	-21.30	35.84	99.74
	<i>BART<sub>Base</sub></i>	40.91	95.97	100.0	77.53	0.78	22.18	100.0
	<i>T5<sub>Base</sub></i>	79.22	90.91	99.87	92.73	-4.29	43.34	99.87
	<i>T5<sub>Small</sub></i>	89.74	88.31	100.0	90.78	-12.27	-38.26	99.74
# Neg Samples		3,094	2,990	2,100	643	547	98	9,472

Table 4: Detailed correlation analysis between model series and negative sample types. For each column, one negative type is involved. ‘all’ indicates the usage of all negative types.

all models are appropriately trained and avoid training collapses by examining their ROUGE scores. The best hyper-parameters are used and kept the same for models from the same group.

## 5.2 Results and Analysis

Table 4 shows the fine-grained results of FacEval. First, we found that FacEval has a higher correlation with MDT models than LDT models. The LDT models are less competent in all aspects as fewer data are involved with training. The generated summaries are weaker in multiple elements, including factuality, fluency, coherence, and granularity. In contrast, the MDT models mainly deteriorate in factuality with factually corrupted training data. Therefore, it is desired that FacEval shows a higher correlation to MDT models.

Second, when considering each negative sample type, a relatively higher correlation is shown with negation (NG), pronoun swap (PS) and speaker swap (SS). It is because more comparison pairs are created with these methods. Also, for chit-chat dialogues, almost all summaries contain reasoning concerning speakers and personnel in the dialogue. And the confirmation of action is happening in multiple utterances. As a result, these several error types are more commonly witnessed in dialogue summarization, as illustrated in Figure 1. In contrast, the negative pairs generated by entity swap (ES), date swap (DS) and number swap (NS) show a lower correlation. It is because these samples are more related to particular errors which appear in various formats and are more challenging to simulate. Even though solely considering these samples shows a lower correlation, we still include them

in the overall comparison process to have a more comprehensive evaluation.

## 5.3 Comparison with Other Metrics

We include a list of popular evaluation methods for summarization to compare our evaluation schema with existing ones. It contains three generic evaluation methods and four dedicated faithfulness evaluation methods.

### 5.3.1 Baseline Metrics

Three generic evaluation methods are as follows:

**ROUGE** (Lin, 2004) score is the default evaluation metric for summarization. We experiment with the F-measure of ROUGE-1, ROUGE-2 and ROUGE-L<sup>2</sup>, which are derived from the uni-gram overlap, the bi-grams overlap and the longest common subsequence (LCS) between generated and reference summaries, respectively.

**BLEU** (Papineni et al., 2002) score is the primary evaluation metric for machine translation. It is mainly designed for corpus-level similarity computation derived from n-gram overlaps. In the following experiments, we report the most commonly used BLEU-4 score.

**BERTScore** (Zhang\* et al., 2020) leverages the pre-trained contextual embeddings from BERT and computes the similarity between text sequences by matching words in candidate and reference by cosine similarity.

Four faithful evaluation methods are as follows:

**FactCC<sub>v1</sub>** (Kryscinski et al., 2020) first augment summaries by applying rule-based transformations

<sup>2</sup>The py-rouge package is used. Different implementations can lead to discrepant results, as discussed in Feng et al. (2021). [pypi.org/project/py-rouge/](https://pypi.org/project/py-rouge/)

Train. Strategy Model	LDT				MDT				Avg.
	<i>BART<sub>Large</sub></i>	<i>BART<sub>Base</sub></i>	<i>T5<sub>Base</sub></i>	<i>T5<sub>Small</sub></i>	<i>BART<sub>Large</sub></i>	<i>BART<sub>Base</sub></i>	<i>T5<sub>Base</sub></i>	<i>T5<sub>Small</sub></i>	
<i>Non-Factual Evaluation Schema</i>									
ROUGE-1	81.35	95.79	94.44	95.94	84.16	91.04	95.58	85.84	90.52
ROUGE-2	86.77	96.84	96.39	96.09	90.13	97.01	95.45	93.90	94.07
ROUGE-L	75.64	96.24	96.39	92.63	86.23	98.31	97.14	94.16	92.09
BLEU	<b>91.88</b>	90.08	92.33	86.02	89.87	94.16	93.38	84.03	90.22
BERTScore	88.87	97.14	<b>97.29</b>	95.79	91.69	94.42	95.97	92.47	94.20
<i>Factual Evaluation Schema</i>									
FactCC <sub>v1</sub>	—	—	—	—	—	—	—	—	—
FactCC <sub>v2</sub>	82.39	84.57	42.45	97.07	96.01	99.22	98.27	<b>100.0</b>	87.50
FEQA	6.02	30.08	-60.15	33.23	57.92	54.29	75.06	85.58	35.25
NLI	39.40	31.28	93.08	90.53	91.17	82.99	93.77	92.99	76.90
<i>FacEval</i> (ours)	83.70	<b>99.40</b>	89.62	<b>98.05</b>	<b>99.74</b>	<b>100.0</b>	<b>99.87</b>	99.74	<b>96.27</b>

Table 5: Comparison of a series of automatic evaluation metrics. The result shown is Spearman’s rank correlation between model ranks and predicted scores.

from the document sentences and fine-tune a pre-trained language model BERT to classify whether the summary is consistent or inconsistent with the documents. It is initially trained in the news summarization domain.

**FactCC<sub>v2</sub>** is an adapted FactCC<sub>v1</sub> to the dialogue domain by us. The negative summaries are generated using our transformations discussed in Section 4.1. We train a *T5<sub>Small</sub>* model as the classifier and take dialogue and summary as input to predict their consistency.

**FEQA** (Durmus et al., 2020) is a question generation and answering method for faithfulness evaluation. It first extracts question-answer pairs from summaries with pre-trained models. Then, a QA model pulls answers from the document with the same questions. A matching method is deployed to measure the similarity between both answer responses from the summary and document as the factuality score. Note that the model is designed for documents.

**NLI** (Falke et al., 2019) is an entailment-based method which takes the maximal entailment probability between summary and document sentence as the factual consistency score. As no dialogue-based entailment model is available, we compute the entailment probability between reference and generated summaries with a BERT-based entailment model trained on SNLI and MultiNLI datasets.

### 5.3.2 Results and Analysis

The experimental results are shown in Table 5.

**Non-factual Evaluator:** The non-factual evalua-

tion methods measure the similarity between reference and generated summaries. ROUGE and BLEU are derived from n-gram overlaps, which indicate the overall generation quality. It is expected that evaluators have a reasonable correlation with LDT models as training with fewer data will resulting quality degradation of the summary in all aspects. For MDT models, they also show a good correlation. We observe that R-2 and R-L are better indicators than R-1 for factuality evaluation. It is because simply replacing isolated tokens can easily change the factual correctness of a summary without much influence on the R-1 score.

**Factual Evaluator:** As FactCC<sub>v1</sub> is trained for the news summarization, we found that the released model is incapable of making predictions for dialogues. Similarly, FEQA is not a good indicator of model performance because the question and answer generation models are optimized for documents, which limits its transferability to the dialogue domain. In comparison, FactCC<sub>v2</sub> and NLI are better evaluation methods for factuality and can make good predictions on MDT models.

**FacEval Properties:** FacEval is the only model-level evaluation schema. The examined model requires reasonable predictions on single sentences and differentiation between positive and negative pairs. Therefore, FacEval shows a strong correlation with LDT and MDT models. The exceptional performance on MDT models indicates that FacEval can effectively reflect model’s capability on factuality.



Model	# Params	All	NG	PS	SS	ES	DS	NS
<i>BART<sub>Large</sub></i>	400M	86.87	89.03	89.08	84.19	89.46	84.22	68.37
<i>BART<sub>Base</sub></i>	140M	85.73	87.73	89.43	79.46	89.96	82.50	68.63
<i>T5<sub>Base</sub></i>	220M	85.32	86.81	88.27	79.92	90.71	84.90	67.75
<i>T5<sub>Small</sub></i>	60M	79.87	82.09	85.94	69.68	87.17	79.88	62.96

Table 6: Benchmarking results on 4 different models

### 5.3.3 Benchmarking and Future Directions

It is beneficial to provide benchmarking performance and analysis on popular dialogue summarization models. As discussed in Sec. 3, dedicated dialogue summarization models do not outperform their baseline models in terms of faithfulness. Therefore, we evaluate on *T5* and *BART* models instead.

The benchmarking results are shown in Tab 6. There are several interesting findings. First, *BART<sub>Large</sub>* has the largest model size as well as the overall best performance. We can also conclude that larger pre-trained models are more faithful based on our evaluation. Second, *BART* model is generally better than *T5* in factuality with model size taken into consideration. This may be because that *BART* is designed for the generation with various denoising objectives, while *T5* is a sequence-to-sequence model for different tasks including but not limited to generation. Third, from fine-grained analysis, we can see that speaker information (from SS) is a major challenge for dialogue summarization. This is because dialogue involves multiple speakers and their roles are tightly involved in the ideal summarization. Therefore, how to improve the model’s understanding capability on speaker roles is an interesting direction to explore (Liu et al., 2021b). Meantime, because some faithful errors are coming from lack of commonsense for existing models (Wang et al., 2021). How to effectively combine hidden semantics (Wang and Kuo, 2020) and well-structured knowledge (Ge et al., 2022) are also worth exploration.

## 6 Conclusion

We believe our faithfulness analysis and evaluation method can facilitate the development of dialogue summarization systems. Instead of measuring faithfulness on generated summaries, we directly assess the model’s capability by multi-choice questions. We expect FacEval to be effectively extended to other generation scenarios.

## 7 Limitations

The testing samples used in our method are obtained by rule-based transformations of the reference and back-translated summaries. It is still limited to the types of transformations designed. More transformation methods need to be proposed to have a comprehensive evaluation. To obtain more natural summaries, we can gather generated summaries and perform annotation by humans. The model can be evaluated in more aspects and closer to real-world scenarios with more available samples.

Verifying the effectiveness of the model-level evaluation schema requires various models and their corresponding rankings. However, such model rankings are currently unavailable because 1) there are not enough varieties of dialogue summarization models as it is still a developing field; 2) the annotations on the faithfulness of dialogue summaries are not adequate. Therefore, in this work, we refer to heuristic methods to manually create a series of models with desired capability levels. When new evaluators are proposed, the best practice is to leverage model-level human rankings for performance benchmarking.

## Acknowledgement

This research is supported by the Agency for Science, Technology and Research (A\*STAR) under its AME Programmatic Funding Scheme (Project No. A18A2b0046) and Science and Engineering Research Council, Agency of Science, Technology and Research (A\*STAR), Singapore, through the National Robotics Program under Human-Robot Interaction Phase 1 (Grant No. 192 25 00054). This work is also supported by the Internal Project Fund from Shenzhen Research Institute of Big Data under Grant T00120220002.

## References

- Meng Cao, Yue Dong, and Jackie Cheung. 2022. **Hal-lucinated but factual! inspecting the factuality of hallucinations in abstractive summarization.** In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3340–3354, Dublin, Ireland. Association for Computational Linguistics.
- 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.
- Ziqiang Cao, Furu Wei, Wenjie Li, and Sujian Li. 2018. **Faithful to the original: Fact aware neural abstractive summarization.** In *thirty-second AAAI conference on artificial intelligence*.
- Jean Carletta, Simone Ashby, Sebastien Bourban, Mike Flynn, Mael Guillemot, Thomas Hain, Jaroslav Kadlec, Vasilis Karaiskos, Wessel Kraaij, Melissa Kronenthal, et al. 2005. **The ami meeting corpus: A pre-announcement.** In *International workshop on machine learning for multimodal interaction*, pages 28–39. Springer.
- Jiaao Chen and Diyi Yang. 2020. **Multi-view sequence-to-sequence models with conversational structure for abstractive dialogue summarization.** In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4106–4118, Online. Association for Computational Linguistics.
- Bonnie Dorr, David Zajic, and Richard Schwartz. 2003. **Hedge trimmer: A parse-and-trim approach to headline generation.** In *Proceedings of the HLT-NAACL 03 Text Summarization Workshop*, pages 1–8.
- Esin Durmus, He He, and Mona Diab. 2020. **FEQA: A question answering evaluation framework for faithfulness assessment in abstractive summarization.** In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5055–5070, Online. Association for Computational Linguistics.
- Wafaa S El-Kassas, Cherif R Salama, Ahmed A Rafea, and Hoda K Mohamed. 2021. **Automatic text summarization: A comprehensive survey.** *Expert Systems with Applications*, 165:113679.
- Alexander R. Fabbri, Wojciech Kryściński, Bryan McCann, Caiming Xiong, Richard Socher, and Dragomir Radev. 2021. **SummEval: Re-evaluating summarization evaluation.** *Transactions of the Association for Computational Linguistics*, 9:391–409.
- 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.
- Xiachong Feng, Xiaocheng Feng, and Bing Qin. 2021. **A survey on dialogue summarization: Recent advances and new frontiers.** *arXiv preprint arXiv:2107.03175*.
- Xiou Ge, Yun-Cheng Wang, Bin Wang, and C-C Jay Kuo. 2022. **CompoundE: Knowledge graph embedding with translation, rotation and scaling compound operations.** *arXiv preprint arXiv:2207.05324*.
- 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.
- Tanya Goyal and Greg Durrett. 2020. **Evaluating factuality in generation with dependency-level entailment.** In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3592–3603, Online. Association for Computational Linguistics.
- Tanya Goyal and Greg Durrett. 2021. **Annotating and modeling fine-grained factuality in summarization.** In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1449–1462, Online. Association for Computational Linguistics.
- Alex Graves. 2012. **Sequence transduction with recurrent neural networks.** *arXiv preprint arXiv:1211.3711*.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. **Teaching machines to read and comprehend.** *Advances in neural information processing systems*, 28.
- Dandan Huang, Leyang Cui, Sen Yang, Guangsheng Bao, Kun Wang, Jun Xie, and Yue Zhang. 2020. **What have we achieved on text summarization?** In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 446–469, Online. Association for Computational Linguistics.
- Hayato Kobayashi, Masaki Noguchi, and Taichi Yasuka. 2015. **Summarization based on embedding distributions.** In *Proceedings of the 2015 Conference*

- on *Empirical Methods in Natural Language Processing*, pages 1984–1989, Lisbon, Portugal. Association for Computational Linguistics.
- Wojciech Kryscinski, Nitish Shirish Keskar, Bryan McCann, Caiming Xiong, and Richard Socher. 2019. [Neural text summarization: A critical evaluation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 540–551, Hong Kong, China. Association for Computational Linguistics.
- 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.
- Faisal Ladhak, Esin Durmus, He He, Claire Cardie, and Kathleen McKeown. 2022. [Faithful or extractive? on mitigating the faithfulness-abstractiveness trade-off in abstractive summarization](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1410–1421, Dublin, Ireland. Association for Computational Linguistics.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. ROUGE: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Junpeng Liu, Yanyan Zou, Hainan Zhang, Hongshen Chen, Zhuoye Ding, Caixia Yuan, and Xiaojie Wang. 2021a. [Topic-aware contrastive learning for abstractive dialogue summarization](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 1229–1243, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Yixin Liu, Pengfei Liu, Dragomir Radev, and Graham Neubig. 2022. [BRIO: Bringing order to abstractive summarization](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2890–2903, Dublin, Ireland. Association for Computational Linguistics.
- Zhengyuan Liu, Ke Shi, and Nancy Chen. 2021b. [Coreference-aware dialogue summarization](#). In *Proceedings of the 22nd Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 509–519, Singapore and Online. Association for Computational Linguistics.
- 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.
- Ramesh Nallapati, Bowen Zhou, Cicero dos Santos, Çağlar Gulçehre, and Bing Xiang. 2016. [Abstractive text summarization using sequence-to-sequence RNNs and beyond](#). In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, pages 280–290, Berlin, Germany. Association for Computational Linguistics.
- Shashi Narayan, Shay B. Cohen, and Mirella Lapata. 2018. [Don’t give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1797–1807, Brussels, Belgium. 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.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Romain Paulus, Caiming Xiong, and Richard Socher. 2018. [A deep reinforced model for abstractive summarization](#). In *International Conference on Learning Representations*.
- 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.
- Harvey Sacks, Emanuel A Schegloff, and Gail Jefferson. 1978. A simplest systematics for the organization of turn taking for conversation. In *Studies in the organization of conversational interaction*, pages 7–55. Elsevier.
- Thomas Scialom, Paul-Alexis Dray, Sylvain Lamprier, Benjamin Piwowarski, Jacopo Staiano, Alex Wang, and Patrick Gallinari. 2021. [QuestEval: Summarization asks for fact-based evaluation](#). In *Proceedings of*

- the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6594–6604, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. [Get to the point: Summarization with pointer-generator networks](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1073–1083, Vancouver, Canada. Association for Computational Linguistics.
- Liyan Tang, Tanya Goyal, Alexander R Fabbri, Philippe Laban, Jiacheng Xu, Semih Yahvuz, Wojciech Kryściński, Justin F Rousseau, and Greg Durrett. 2022a. Understanding factual errors in summarization: Errors, summarizers, datasets, error detectors. *arXiv preprint arXiv:2205.12854*.
- Xiangru Tang, Arjun Nair, Borui Wang, Bingyao Wang, Jai Desai, Aaron Wade, Haoran Li, Asli Celikyilmaz, Yashar Mehdad, and Dragomir Radev. 2022b. [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.
- Alex Wang, Kyunghyun Cho, and Mike Lewis. 2020. [Asking and answering questions to evaluate the factual consistency of summaries](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5008–5020, Online. Association for Computational Linguistics.
- Bin Wang and C-C Jay Kuo. 2020. SBERT-WK: A sentence embedding method by dissecting bert-based word models. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 28:2146–2157.
- Bin Wang, Guangtao Wang, Jing Huang, Jiaxuan You, Jure Leskovec, and C-C Jay Kuo. 2021. Inductive learning on commonsense knowledge graph completion. In *2021 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE.
- Bin Wang, Chen Zhang, Chengwei Wei, and Haizhou Li. 2022. A focused study on sequence length for dialogue summarization. *arXiv preprint arXiv:2209.11910*.
- Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter J. Liu. 2020. Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. In *Proceedings of the 37th International Conference on Machine Learning, ICML'20*. JMLR.org.
- 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*.
- Ming Zhong, Pengfei Liu, Yiran Chen, Danqing Wang, Xipeng Qiu, and Xuanjing Huang. 2020. [Extractive summarization as text matching](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6197–6208, Online. Association for Computational Linguistics.
- Chenguang Zhu, William Hinthorn, Ruochen Xu, Qingkai Zeng, Michael Zeng, Xuedong Huang, and Meng Jiang. 2021a. [Enhancing factual consistency of abstractive summarization](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 718–733, Online. Association for Computational Linguistics.
- Chenguang Zhu, Yang Liu, Jie Mei, and Michael Zeng. 2021b. [MediaSum: A large-scale media interview dataset for dialogue summarization](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5927–5934, Online. Association for Computational Linguistics.