

Mind the Gap! Injecting Commonsense Knowledge for Abstractive Dialogue Summarization

Seungone Kim^{1*}

Se June Joo^{1*}

Hyungjoo Chae^{1,2*}

Chaehyeong Kim^{1*}

Seung-won Hwang³

Jinyoung Yeo^{1,2†}

Yonsei University¹

Tutoring at Market Designers²

Seoul National University³

{louisdebroglie, sr7418, mapoout, cheris8, jinyeo}@yonsei.ac.kr
seungwonh@snu.ac.kr

Abstract

In this paper, we propose to leverage the unique characteristics of dialogues sharing commonsense knowledge across participants, to resolve the difficulties in summarizing them. We present **SICK**, a framework that uses commonsense inferences as additional context. Compared to previous work that solely relies on the input dialogue, **SICK** uses an external knowledge model to generate a rich set of commonsense inferences and selects the most probable one with a similarity-based selection method. Built upon **SICK**, **SICK++** utilizes commonsense as supervision, where the task of generating commonsense inferences is added upon summarizing the dialogue in a multi-task learning setting. Experimental results show that with injected commonsense knowledge, our framework generates more informative and consistent summaries than existing methods.

1 Introduction

Abstractive dialogue summarization is a task of generating a shorter summary while preserving the context of a conversation (Li et al., 2017; Gliwa et al., 2019). Unlike conventional document-to-document summarization (e.g., news articles and scientific publications) (Nallapati et al., 2016; Gehrmann et al., 2018), such *dialogue-to-document* summarization suffers from the discrepancy between input and output forms, which makes learning their mapping patterns more challenging.

There are two key challenges that make summarizing dialogues harder than documents. First, detecting unspoken intention is crucial for understanding an utterance (Mendelsohn, 1994; Ram et al., 2018). As shown in Figure 1, without understanding the intent “to make fun of someone”, it is hard to write a correct summary. Second, there exists information that can only be understood when

<p>Alyssa: What do you think about it? Derek: I can fart bright stripes and bright stars better than she sings. ⇒ xWant, to make fun of someone Alyssa: The best part is that she acts like she nailed it. But at least it's funny in a good way. Derek: It is 😊</p>
<p>Golden summary Derek and Alyssa make fun of Fergie's performance of the national anthem.</p>
<p>Melody: you're probably due for a new one anyway, no? Peggy: you're right. 5 years is a long time to own one. Melody: yes, that's ancient by laptop standards ⇒ HinderedBy, the laptop is too old Peggy: ok. i might just not bother getting it repaired after all. Melody: sounds like a good idea</p>
<p>Golden summary Melody's 5-year-old laptop is broken. Tomorrow she'll know what's wrong. She won't be repairing it, because her laptop is too old. Instead, she'll buy a new one.</p>

Figure 1: Example of dialogue-summary pairs. Capturing the intention and hidden meaning is important to generate a novel summary.

its hidden meaning is revealed (Talmy, 1988). For example, it is important to capture the hidden meaning “The laptop is too old” beyond the written text “yes, that's ancient by laptop standards” when writing the summary.

Commonsense knowledge models (Hwang et al., 2021; Gabriel et al., 2021; West et al., 2022) such as COMET can generate a set of event-centered (e.g., HINDEREDBY, XREASON, XNEED) and social-interaction (e.g., XINTENT, XWANT) commonsense inferences. We argue that the aforementioned issues can be mitigated using commonsense knowledge by *filling in the gap* in a dialogue.

Despite its effectiveness, it is non-trivial to use commonsense knowledge for improving abstractive dialogue summarization performance. While commonsense knowledge has been widely applied to commonsense reasoning (Bosselut and Choi, 2021; Liu et al., 2020; Chang et al., 2021; Wang et al., 2021; Kim et al., 2022) or question answering (Shwartz et al., 2020; Bosselut et al., 2021),

*Equal contribution

†Corresponding author

its usage for summarization is understudied (Feng et al., 2021).

In this paper, we present our framework **SICK** and its extension **SICK++** to properly inject commonsense knowledge into state-of-the-art language models (e.g., BART (Lewis et al., 2020)) for abstractive dialogue summarization. We argue a naïve adoption of commonsense only hurts performance in summarization, as (a) expanding source contents is counter-intuitive approach for the goal of condensation, and (b) simply adding additional inputs in pre-trained language models does not lead to robust inferences as reported in Zhou et al. (2021b,a). Our framework addresses this by (a) filtering and (b) robust training.

Based on analytical measurements, commonsense knowledge is selected and enumerated as an additional context of dialogue inputs. In **SICK++**, we also design a new auxiliary task named *commonsense supervision*. Using commonsense knowledge generated from gold summaries as additional supervision, the goal of the task is to generate the target commonsense. Then, the dialogue summarization and commonsense generation tasks are jointly learned in a multi-task learning setting to effectively inject commonsense knowledge into the shared encoder.

To validate our framework, we conduct a set of experiments on abstractive dialogue summarization benchmarks. Empirical results show that our framework can improve summarization performance with leveraged commonsense knowledge, outperforming other baselines. Human evaluation results prove that our method can generate informative and consistent summaries. In addition, we conduct experiments to analyze the effect of commonsense knowledge on abstractive dialogue summarization.

2 Related Work

2.1 Abstractive Dialogue Summarization

Compared to extractive summarization (Nallapati et al., 2017; Zhang et al., 2018; Zhong et al., 2020), abstractive summarization is considered more challenging and has received extensive attention (Rush et al., 2015; See et al., 2017). Benefiting from the advance of large-scale pre-trained language models, the performance of encoder-decoder models has achieved substantial improvements in document summarization (Nallapati et al., 2016; Gehrmann et al., 2018; Zhang et al., 2020a).

Recently, abstractive dialogue summarization has become another emerging research area, where the goal is to generate concise summaries for conversations such as meetings (Zhu et al., 2021) and chit-chat (Chen et al., 2021). It is more difficult to capture the key points in dialogues than documents, because people do not state the obvious (Grice, 1975) and conversations have a more interactive flow of information between speakers (Li et al., 2021b). Based on the characteristic of the dialogues, many studies focused on organizing the information in the dialogues. Wu et al. (2021) propose to create a summary sketch for a given dialogue as weak supervision. Chen and Yang (2021) explicitly model structures in conversations by incorporating discourse relations and action triples in utterances through structured graphs. Instead of organizing the given dialogue for better understanding, our method adds additional knowledge to fill in the missing cues between dialogues.

2.2 Commonsense Knowledge Models

Recent research has focused on commonsense knowledge acquisition through different lines: commonsense knowledge graphs and commonsense knowledge models. Unlike static knowledge graphs such as ATOMIC (Sap et al., 2019) in which entities and relations between entities are represented in nodes and edges, commonsense knowledge models such as COMET (Bosselut et al., 2019) have been shown to generate implicit commonsense inferences along several dimensions depending on what knowledge graphs they are trained on. Commonsense knowledge models can be used to anticipate and reason unobserved causes and effects in relation to the observed event (Sap et al., 2019). Despite these functions, they are applied on defined domains (Shwartz et al., 2020; Bosselut et al., 2021). Especially, on dialogue summarization task, there has been limited usage of using commonsense directly as additional context. For example, Feng et al. (2021) and Zhou et al. (2022) utilized ConceptNet (Speer et al., 2017), a static knowledge graph with encyclopedic knowledge, to fill in the missing cues between dialogue.

In contrast to encyclopedic knowledge, our method uses event-centered and social-interaction knowledge as additional context. Also, instead of retrieving from a static knowledge graph, our method deploys on-the-fly commonsense knowledge models to acquire a rich set of commonsense

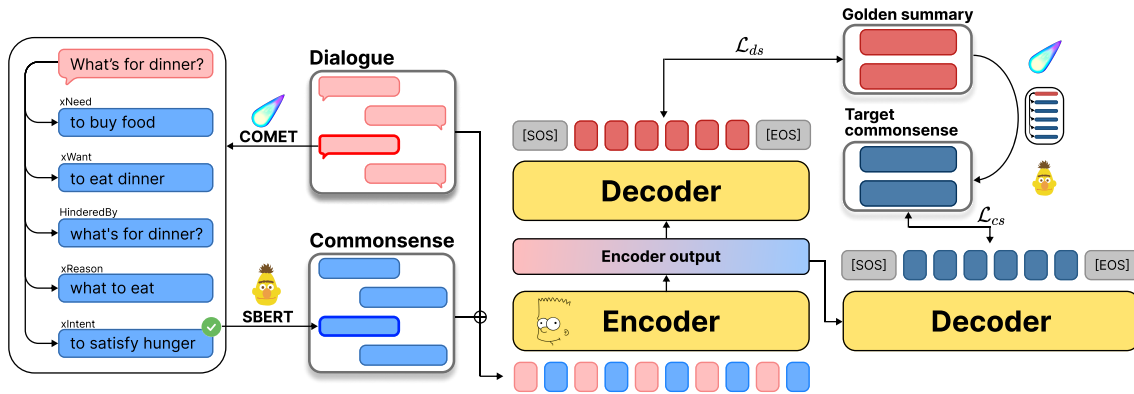


Figure 2: The overall framework of SICK and SICK++. The decoder generating target commonsense is used for SICK++.

inferences dynamically.

3 Proposed Framework

In this section, to inject commonsense knowledge for rich abstractive dialogue summarization, we introduce our new framework, SICK(Summarizing with Injected Commonsense Knowledge) and its extension SICK++, as shown in Figure 2.

3.1 Task Description

Our task definition follows a sequence-to-sequence learning problem setting. Based on pre-trained generative language models, our goal is to learn a mapping function $\mathbb{M} : \mathcal{D} \rightarrow \mathcal{Y}$ where $\mathcal{D} = \{u_1, u_2, \dots, u_n\}$ is a dialogue with n utterances, and $\mathcal{Y} = \{y_1, y_2, \dots, y_m\}$ is a corresponding summary of m sentences.

We further extend the task with two modifications. First, we generate and filter to acquire a set of commonsense knowledge $\mathcal{C} = \{c_1, c_2, \dots, c_n\}$ based on \mathcal{D} (Section 3.2, 3.3). Then, we adjust the mapping function as $\mathbb{M} : \mathcal{X} \rightarrow \mathcal{Y}$, where \mathcal{X} is a cross concatenation of \mathcal{D} and \mathcal{C} (Section 3.3). Second, we add an auxiliary task *commonsense supervision*, $\mathbb{M}^* : \mathcal{X} \rightarrow \mathcal{Z}$, where the target commonsense $\mathcal{Z} = \{z_1, z_2, \dots, z_m\}$ is acquired based on \mathcal{Y} (Section 3.4).

3.2 Commonsense Knowledge Generation

In SICK, commonsense knowledge is leveraged as a supplement to insufficient context of dialogues. As shown in Table 1, additional information can be derived from the given utterance in various aspects. There are some cases where the intention of the speaker is crucial in comprehending the dialogue (e.g., “to believe in something”, “to talk to someone

Utterance	Charlie : Do you really believe that dreams can mean something?
HINDEREDBY	Charlie doesn't believe in dreams.
XWANT	to talk to someone about dreams.
XINTENT	to believe in something.
XNEED	to have a dream.
XREASON	Charlie is a skeptic.

Table 1: Example of commonsense knowledge generated by COMET given a dialogue.

about dreams”). Whereas in other cases, the hidden information is necessary (e.g., “Charlie doesn't believe in dreams”, “to have a dreams”, “Charlie is a skeptic”). We adopt an external commonsense knowledge model that generates a diverse and abundant set of commonsense inferences in natural language. Given a text x and a relation type r , the commonsense knowledge model gives an output c grounded to the relation type. i.e., $f : (x, r) \rightarrow c$.

Specifically, we use COMET (Hwang et al., 2021), a widely-used generative commonsense model as our external model. Among the 23 possible candidate relation types, we choose 5 unique relations that helps understand the speakers' intentions and find out the missing information. COMET generates 5 commonsense inferences per relation type, resulting in 25 per input.

Also, to attend to the previous utterances when generating commonsense inferences, we further explore a discourse-aware model, PARACOMET (Gabriel et al., 2021) that generates coherent inferences. More specifically, while COMET generates a set of commonsense inferences considering only one sentence at a time, PARACOMET adopts an internal memory module to consider previous dialogue history when generating an output.

Prev-Utterances	Jane : google maps says it is at least 3h Steven : I used to make it in 2, trust me Jane : but it's almost 300km Steven : the road is new , we will make it
Utterance	Jane : I don't want to stress out, let's meet at 4:30 instead of 5, ok?
XINTENT	to avoid stress.
XWANT	to not be late.
XREACT	annoyed
XEFFECT	PersonX sweats from nervousness.
XATTR	nervous.

Table 2: Example of commonsense knowledge generated by PARA-COMET given a dialogue.

In Table 2, when generating commonsense inferences of the current dialogue, PARA-COMET conditions on the previous utterance. Knowing what was previously stated, the intention of the speaker (e.g., “to not be bothered”, “to not be stressed”, “upset”) and the hidden knowledge (e.g., “annoyed”, “PersonX gets into trouble”) differs from COMET.

3.3 Summarizing with Injecting Commonsense Knowledge (SICK)

Filtering Compared to question answering and commonsense generation (Shwartz et al., 2020; Wang et al., 2021), summarizing dialogues has another difficulty. The amount of data provided as the input should be mapped into the output in a concise form. Therefore, simply providing extra input (i.e., commonsense knowledge) may confuse the model when generating a summary. Moreover, it is unable to add every possible commonsense knowledge to the dialogue due to the limited input sequence length of transformer-based models.

To address this issue, we propose to select the most favorable commonsense for each utterance. For 25 candidates, we measure the semantic relevance between the utterance and the commonsense inference concerning. One could imagine that filtering could choose only very similar “safe” examples that might not be as valuable/interesting in practice (i.e., *diversity vs. quality*). However, recent literature address that paradoxically, filtering increases diversity (West et al., 2022). We also discuss the impact of different filtering methods in Appendix E.

We employ SBERT (Reimers and Gurevych, 2019) to compute the similarity score between utterance and commonsense pairs. We select one commonsense inference c_i , with the highest score for each utterance u_i among the candidate relations

\mathcal{R} . As a result, we obtain the input commonsense $\mathcal{C} = \{c_i\}_{i=1}^n$ aligned with dialogue \mathcal{D} .

$$c_i = \operatorname{argmax}_{c_i^r} (\operatorname{score}(u_i, c_i^r)) \quad (r \in \mathcal{R}) \quad (1)$$

Cross Concatenation After obtaining the input commonsense for the dialogue, we concatenate the dialogue and its corresponding set of commonsense inferences. To encode the information that c_i is derived from u_i , we enforce to attend its neighbor token. Instead of concatenating \mathcal{D} and \mathcal{C} back and forth, we concatenate turn by turn considering *locality of reference* (Clark et al., 2019; Zaheer et al., 2020), where tokens tend to attend its neighboring tokens. In order to separate the modalities between dialogues and commonsense inferences, we add special tokens $\langle I \rangle$, $\langle /I \rangle$ in back and forth of each commonsense inference c_i . Thus the input sequence \mathcal{X} is formulated as:

$$\mathcal{X} = \mathcal{D} \oplus \mathcal{C} = \dots \parallel u_i \parallel \langle I \rangle c_i \langle /I \rangle \parallel \dots \quad (2)$$

Training SICK is built upon a transformer-based encoder-decoder architecture. The encoder fuses the information from two different modalities (i.e., dialogue and commonsense inference). By the stack of decoders, the encoder output is used for cross-attention with the summary. The training objective, a negative-log likelihood parameterized by θ_{ds} , can be formulated as:

$$\mathcal{L}_{ds} = - \sum_{i=1}^{|\mathcal{Y}|} \sum_{j=1}^{|y_i|} \log P(w_{i,j} | w_{i < j}, \mathcal{X}; \theta_{ds}) \quad (3)$$

where $w_{i,j}$ is j -th token of i -th sentence y_i in target summary \mathcal{Y} .

3.4 SICK++

Commonsense Supervision It is well known that models do not consider the actual input as a whole and only look at certain parts of the input therefore not performing the underlying task but some derivative (Branco et al., 2021). For example, in Figure 1, although it is critical to understand Derek’s intention (e.g., “to make fun of Fergie’s performance”), SICK may not utilize the commonsense to comprehend the dialogue.

To overcome this problem, we propose an auxiliary task named *commonsense supervision*. In addition to providing commonsense on the input side, we also leverage commonsense knowledge as additional target variable, which prevents the model

from disregarding commonsense and enforces actually to utilize commonsense. For instance, when the summary “*Derek and Alyssa make fun of Fergie’s performance of the national anthem.*” is given to COMET, we observe that a target commonsense “*to make fun of*” is generated. Generating both the summary and the target commonsense has an effect of emphasizing that the input commonsense inference “*to make fun of someone*” is important.

We generate a set of target commonsense inferences \mathcal{Z} with the summary \mathcal{Y} using an external knowledge model f . Then we filter and select the most plausible target commonsense.

$$z_i = \operatorname{argmax}_{z_i^r} (\operatorname{score}(y_i, z_i^r)) \quad (r \in \mathcal{R}) \quad (4)$$

To adopt commonsense knowledge as additional supervision, we further include commonsense summarization decoder \mathbb{D}_{cs} , which learns to generate target commonsense \mathcal{Z} .

Training With the target commonsense \mathcal{Z} , we train the commonsense summarization decoder \mathbb{D}_{cs} to minimize a negative log-likelihood loss function such as:

$$\mathcal{L}_{cs} = - \sum_{i=1}^{|\mathcal{Z}|} \sum_{j=1}^{|z_i|} \log P(w_{i,j} | w_{i < j}, \mathcal{X}; \theta_{cs}) \quad (5)$$

where w_j^i is a j -th word token of sentence c_i^y from the target commonsense \mathcal{Z} .

We linearly combine the two loss functions, Equation 3 and Equation 5, in a multi-task learning setting as follows:

$$\mathcal{L}_{total} = \lambda \cdot \mathcal{L}_{ds} + (1 - \lambda) \cdot \mathcal{L}_{cs} \quad (6)$$

where \mathcal{L}_{ds} and \mathcal{L}_{cs} denote the loss function for dialogue summarization decoder \mathbb{D}_{ds} and commonsense summarization decoder \mathbb{D}_{cs} , respectively. λ is a predefined hyperparameter to adjust the scale of each loss. In our setting, we set $\lambda = 0.66$.

Inference During inference, given an input dialogue \mathcal{D}_{test} , we first obtain input commonsense \mathcal{C}_{test} for the dialogue, and specify input sequence as $\mathcal{X}_{test} = \mathcal{D}_{test} \oplus \mathcal{C}_{test}$ by concatenating turn by turn. Then, the model predicts summary $\hat{\mathcal{Y}}_{test} = \mathbb{M}(\mathcal{X}_{test})$ for the dialogue \mathcal{D}_{test} . Note that while we train the model in a dual-decoder setting, we only use the dialogue summarization decoder \mathbb{D}_{ds} and discard the commonsense prediction decoder \mathbb{D}_{cs} at inference time.

	SAMSum	DialogSum
Train	14,732	12,460
Dev	818	500
Test	819	500
#Tokens/dialogue	82.57	121.56
#Tokens/summary	20.30	22.64
#Turns	11.2	9.5
#Speaker	2.4	2.0
#Compression rate	0.3538	0.2001

Table 3: Statistics of dialogue summarization datasets. # stands for the average number. The compression rate is a ratio of the length of summary divided by the length of dialogue.

4 Experimental Setup

4.1 Datasets and Baselines

We perform experiments on SAMSum (Gliwa et al., 2019) and DialogSum (Chen et al., 2021) datasets. SAMSum is the most widely used resource for abstractive dialogue summarization task. It consists of natural messenger-like conversations in English created by linguists with manually annotated summaries. DialogSum (Chen et al., 2021) is a recently released dataset for a more challenging task with a lower compression ratio. It contains multi-turn dialogues of real-life scenarios collected from three dialogue corpora. The data statistics are in Table 3.

We adopt three different types of baselines: (i) generative language models (See et al., 2017; Wu et al., 2019; Vaswani et al., 2017); (ii) pre-trained language models (Zhang et al., 2020c; Dong et al., 2019; Zhang et al., 2020a; Lewis et al., 2020); (iii) dialogue summarization Models (Feng et al., 2021; Chen and Yang, 2021; Wu et al., 2021). We provide more details in Appendix A.

4.2 Implementation Details

We employ two automatic evaluation metrics as: (i) ROUGE (Lin, 2004) scores, including ROUGE-1, ROUGE-2, and ROUGE-L, which compares word-level uni-gram and bi-gram, and the longest common sequence overlap with the gold summary respectively; (ii) BERTScore (Zhang et al., 2020b)¹, the recent popular metric for text generation, which computes the contextual similarity score between generated and reference summaries. We report F1 scores for both metrics. For simplicity, we use R-1, R-2, R-L, and B-S to denote ROGUE-1, ROUGE-2,

¹We follow https://github.com/Tiiiger/bert_score to calculate BERTScore. Note that different tools may result in different BERTScore.

Model	SAMSum				DialogSum			
	R-1	R-2	R-L	B-S	R-1	R-2	R-L	B-S
PointerGenerator (See et al., 2017)*	32.27	14.42	34.36	/	/	/	/	/
DynamicConv (Wu et al., 2019)*	41.07	17.11	37.27	/	/	/	/	/
Transformer (Vaswani et al., 2017)*	42.37	18.44	39.27	/	/	/	/	/
DialoGPT (Zhang et al., 2020c)†	39.77	16.58	38.42	/	/	/	/	/
BART-xsum (Lewis et al., 2020)†	51.74	26.46	48.72	53.87	/	/	/	/
UniLM (Dong et al., 2019)†	47.85	24.23	46.67	/	42.38	16.88	34.36	69.40
PEGASUS (Zhang et al., 2020a)†	50.50	27.23	49.32	53.35	38.40	13.84	33.41	68.20
BART-xsum (Lewis et al., 2020)‡	52.50	27.67	48.75	68.16	45.15	19.78	36.57	71.09
D-HGN (Feng et al., 2021)	42.03	18.07	39.57	64.20	/	/	/	/
S-BART (Chen and Yang, 2021)	50.70	25.50	48.08	70.07	/	/	/	/
CODS (Wu et al., 2021)	52.65	27.84	50.79	66.55	44.27	17.90	36.98	70.49
SICK w/ COMET (Ours)	53.04	27.60	48.49	71.61	45.70	20.08	40.26	71.08
SICK++ w/ COMET (Ours)	53.24	28.10	48.90	71.71	46.26	20.95	41.05	71.30
SICK w/ PARA-COMET (Ours)	53.39	28.42	49.12	71.83	46.01	20.30	40.75	71.57
SICK++ w/ PARA-COMET (Ours)	53.73	28.81	49.50	71.92	46.20	20.39	40.83	71.32

Table 4: Automatic evaluation on abstractive dialogue summarization benchmarks, *i.e.*, SAMSum and DialogSum. Results on SAMSum with * are obtained from (Gliwa et al., 2019), † are obtained from (Wu et al., 2021) and ‡ is a re-implemented version trained under the same conditions with ours for fair comparison. Results on DialogSum for all models are all reimplemented under the same conditions with ours.

ROUGE-L, and BERTScore (see Appendix C).

Our implementation is based on the Huggingface implementation (Wolf et al., 2020) of BART language model. Specifically, we use the weight checkpoint of BART-xsum². We use a maximum input length of 1024 tokens and output length of 100 tokens. Note that the input is either padded or truncated after each utterance and its corresponding commonsense is concatenated during pre-processing. We use a learning rate of 3e-6 and a batch size of 32 when fine-tuning our model on both benchmarks. We use linear warm-up over the first 600 steps, apply linear decay and use the Adam optimizer (Kingma and Ba, 2015). In our experiments, we use beam search with beam size of 20. We fine-tune our model on SAMSum for 20 epochs and DialogSum for 25 epochs. All experiments are run on one A100 NVIDIA GPU. More implementation details about commonsense knowledge generation is included in Appendix B.

5 Experimental Results

5.1 Automatic Evaluation

Performance Table 4 presents the performance on SAMSum and DialogSum test sets. SICK++ outperforms all baselines on ROUGE-1, ROUGE-2 and BERTScore by a consistent margin in both datasets.

²<https://huggingface.co/facebook/bart-large-xsum>

Comparison with State-of-the-Art We find that pre-trained language models (*e.g.*, DialoGPT, UniLM, PEGASUS, BART-xsum), outperform models that are not pre-trained (*e.g.*, PointerGenerator, DynamicConv, Transformer), confirming the impact of pre-training on abstractive dialogue summarization. Among the pre-trained generative language models examined, PEGASUS and BART-xsum are the most competitive models with ROUGE-1 higher than 50. SICK++ shows improvement on all metrics compared to BART-xsum (*e.g.*, without additional input, commonsense supervision) in both benchmarks showing that our method can be applied in different settings.

Among methods that alter the input to seek additional useful information in a dialogue setting, (*e.g.*, D-HGN, SBART, and CODS), CODS achieves better performance over other baselines in SAMSum. However, on DialogSum, a more challenging setting due to higher abstractiveness, CODS is not able to get as much gain of performance compared to other baselines. Meanwhile, SICK++ outperforms all baselines and shows competitive results implying the robustness of our framework.

Commonsense Models While SICK++ shows better performance regardless of which commonsense generation model is used, the excelling choice differs depending on the dataset. In SAMSum, SICK++ shows better performance with PARA-COMET than with COMET, however it shows op-

Model	SAMSum		DialogSum	
	Info.	Cons.	Info.	Cons.
BART-xsum	3.71	3.48	3.71	3.68
SICK++	3.85	3.81	3.79	3.97
Gold	4.00	3.96	4.03	4.21

Table 5: Human evaluation on SAMSum and DialogSum datasets. Info. and Cons. denotes informativeness and factual consistency respectively.

posite result in DialogSum. We conjecture this due to the characteristic of datasets and commonsense models hold. PARA-COMET has an advantage of using parametric memory to consider previous sentences, which may be sensitive in terms of length. Since SAMSum has shorter length of dialogues than DialogSum, the recurrent memory component of PARA-COMET is less likely to forget the previous sentences. We expect to get better performance with the help of commonsense-models that maintains longer memories of sentences/dialogues and leave this as future research.

5.2 Human Evaluation

We conduct human evaluation to verify the quality of the generated summaries. We randomly sample 50 dialogues from test sets of SAMSum and DialogSum, respectively. Annotators were asked to score the quality of a set of summaries from BART-xsum, SICK++, and ground-truth using a Likert scale from 1 (worst) to 5 (best) in terms of **informativeness** (*i.e.*, covers adequate information) and **factual consistency** (*i.e.*, consistent with the original input). Each summary was evaluated by three different annotators. Also, the win-loss ratio, which is not biased by subjectivity, is 51.33 (informativeness) and 54.16 (factual consistency), which is consistent to the observations made from the absolute scores.

In Table 5, human annotated summaries receive the best scores on all dimensions. SICK++ gets better scores than BART-xsum for informativeness, which matches the results of ROUGE scores in Section 5.1. Neural abstractive models often suffer from hallucinations that affect their reliability (Zhao et al., 2020). SICK++ also produces more consistent summaries even though factual consistency is not explicitly modeled. We assume that incorporating commonsense knowledge helps the model recognize the hidden meanings and better understand the dialogue, resulting in fewer factual

errors without improper reasoning over conversational flow.

6 Analysis

To evaluate the effectiveness of our method, we address the following research questions to guide our experiments:

- **RQ1:** Does commonsense help summarizing dialogues?
- **RQ2:** Is our method worth using in terms of efficiency despite the extra effort?
- **RQ3:** Does *commonsense supervision* lead SICK++ to inject commonsense knowledge?

6.1 RQ1: Commonsense Applicability

We experiment in a zero-shot setting to examine how commonsense knowledge solely affects dialogue summarization. While there exist many factors that could affect performance besides commonsense during training (*e.g.*, hyperparameter configurations), in a zero-shot setting, we can directly compare when commonsense is given and not. We evaluate BART-xsum and SICK on the SAMSum and DialogSum test sets. Note that we use SICK (*i.e.*, only provided input commonsense) instead of SICK++ for zero-shot evaluation, since we cannot access ground-truth summary to generate target commonsense inferences \mathcal{Z} .

Table 6 presents zero-shot evaluation results on SAMSum and DialogSum respectively. We find that SICK outperforms BART-xsum, where the performance gain comes from additional commonsense. Since the only difference between BART-xsum and SICK is the input commonsense, providing extra commonsense for each utterance as Equation 2 helps the model generate more accurate and semantically informative summaries. This also supports the idea that commonsense is essential in resolving the discrepancy between dialogues and documents.

6.2 RQ2: Data Efficiency

Generating commonsense inferences requires irresistible effort, further described in Appendix B. Our approach has limitations in terms of time efficiency. However, we find that our method is helpful in situations where data is insufficient, meaning there is a trade-off (*time vs data efficiency*).

We hypothesize that due to providing additional knowledge and *commonsense supervision*, SICK++

Model	SAMSum				DialogSum			
	R-1	R-2	R-L	B-S	R-1	R-2	R-L	B-S
BART-xsum	20.83	4.28	15.28	46.59	17.40	4.16	13.80	42.97
SICK	23.12	5.09	17.45	47.69	18.32	3.80	14.98	43.97

Table 6: Zero-shot evaluation on SAMSum and DialogSum test set.




Figure 3: Performance of BART-xsum and SICK++ on SAMSum by varying the size of training data. We use BART_{BASE} for both of them. Details are shown in Appendix.

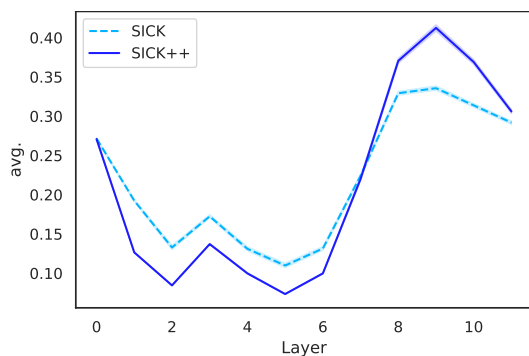


Figure 4: Attention visualization of SICK/SICK++. Each point of the line corresponds to the average attention a particular SICK encoder attention head puts towards commonsense inferences.

can show comparable performance even if only a small amount of training data is available (*i.e.*, data efficiency). As shown in Figure 3, with only 30% of training data, SICK++ shows better performance than BART-xsum trained with 70% of training data. Furthermore, SICK++ consistently outperforms BART-xsum regardless of training data size, proving the robustness of SICK++. The performance gap between SICK++ and BART-xsum can be viewed as a consequence of the leveraged commonsense, based on the fact that BART-xsum is the base architecture of SICK++.

6.3 RQ3: Effect of commonsense supervision on Injecting Commonsense Knowledge

We observe that SICK++ shows better performance than SICK, as we show in Table 4, but the reason for the performance improvement is somewhat unclear. To analyze the role of *commonsense supervision*, we now take a look at how the dual decoder setting impacts commonsense utilization of the en-

coder, the difference between SICK and SICK++.

Attention weights can be viewed as governing how “important” every other token is when producing the next representation for the current token (Clark et al., 2019). We conduct an experiment of measuring the averaged attention value of the commonsense inferences compared to utterances using validation sets of DialogSum, which is more abstractive (*i.e.*, more challenging to comprehend) compared to SAMSum.

The results are illustrated in Figure 4. Rogers et al. (2020) mentioned that final layers of language models are most task-specific, and we observe that SICK++ has marginally higher attention values. We conjecture this is due to the supervision provided by generating \mathcal{Z} instead of relying on distant supervision, meaning that our goal of enforcing the model to use commonsense inferences is successful. SICK++ enforces the encoder to fuse the two different modalities (*e.g.*, utterances, commonsense inferences). Meanwhile in lower and middle layers, SICK++’s attention values tend to be lower than SICK. One possible reason is that lower layers tend to look at syntactic and word-level information (Rogers et al., 2020), whereas the commonsense inferences generated by COMET or PARA-COMET is only meaningful when understood conceptually.

7 Conclusion

In this work, we propose SICK and SICK++ framework in order to resolve the two key challenges: i) *filling in the gap* in dialogues; ii) injecting commonsense knowledge into a model. We show that the difficulties in dialogues are resolved with commonsense knowledge and demonstrated that our

framework can successfully inject commonsense knowledge. As a result of injected commonsense knowledge, we obtain competitive results on SAM-Sum and DialogSum benchmarks.

Acknowledgements

The authors would like to thank the anonymous reviewers for their helpful comments and suggestions. This work was partly supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No. 2020-0-01361, Artificial Intelligence Graduate School Program (Yonsei University)), (No.2021-0-02068, Artificial Intelligence Innovation Hub), and (No. 2022-0-00077, AI Technology Development for Commonsense Extraction, Reasoning, and Inference from Heterogeneous Data). Jinyoung Yeo is a corresponding author.

References

- Antoine Bosselut and Yejin Choi. 2021. Dynamic knowledge graph construction for zero-shot commonsense question answering. In *Proceedings of AAAI*.
- Antoine Bosselut, Ronan Le Bras, and Yejin Choi. 2021. Dynamic neuro-symbolic knowledge graph construction for zero-shot commonsense question answering. In *Proceedings of AAAI*.
- Antoine Bosselut, Hannah Rashkin, Maarten Sap, Chaitanya Malaviya, Asli Celikyilmaz, and Yejin Choi. 2019. Comet: Commonsense transformers for automatic knowledge graph construction. In *Proceedings of ACL*.
- Ruben Branco, António Horta Branco, João António Rodrigues, and J. Silva. 2021. Shortcutted commonsense: Data spuriousness in deep learning of commonsense reasoning. In *Proceedings of EMNLP*.
- Tuhin Chakrabarty, Yejin Choi, and Vered Shwartz. 2022. It’s not rocket science: Interpreting figurative language in narratives. *Transactions of the Association for Computational Linguistics*, 10:589–606.
- Ting-Yun Chang, Yang Liu, Karthik Gopalakrishnan, Behnam Hedayatnia, Pei Zhou, and Dilek Hakkani-Tur. 2021. Incorporating commonsense knowledge graph in pretrained models for social commonsense tasks. In *EMNLP Workshop*.
- Jiaao Chen and Diyi Yang. 2021. Structure-aware abstractive conversation summarization via discourse and action graphs. In *Proceedings of NAACL*.
- Yulong Chen, Yang Liu, Liang Chen, and Yue Zhang. 2021. Dialogsum: A real-life scenario dialogue summarization dataset. In *Proceedings of ACL Findings*.
- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D Manning. 2019. What does bert look at? an analysis of bert’s attention. In *ACL Workshop*.
- Li Dong, Nan Yang, Wenhui Wang, Furu Wei, Xiaodong Liu, Yu Wang, Jianfeng Gao, Ming Zhou, and Hsiao-Wuen Hon. 2019. Unified language model pre-training for natural language understanding and generation. In *Proceedings of NeurIPS*.
- Xiachong Feng, Xiaocheng Feng, and Bing Qin. 2021. Incorporating commonsense knowledge into abstractive dialogue summarization via heterogeneous graph networks. In *Proceedings of China National Conference on Chinese Computational Linguistics*.
- Saadia Gabriel, Chandra Bhagavatula, Vered Shwartz, Ronan Le Bras, Maxwell Forbes, and Yejin Choi. 2021. Paragraph-level commonsense transformers with recurrent memory. In *Proceedings of AAAI*.
- Sebastian Gehrmann, Yuntian Deng, and Alexander M Rush. 2018. Bottom-up abstractive summarization. In *Proceedings of EMNLP*.
- Bogdan Gliwa, Iwona Mochol, Maciej Biesek, and Aleksander Wawer. 2019. Samsun corpus: A human-annotated dialogue dataset for abstractive summarization. In *ACL Workshop*.
- Herbert P Grice. 1975. Logic and conversation. In *Proceedings of Speech acts*, pages 41–58. Brill.
- Jena D Hwang, Chandra Bhagavatula, Ronan Le Bras, Jeff Da, Keisuke Sakaguchi, Antoine Bosselut, and Yejin Choi. 2021. (comet-) atomic 2020: On symbolic and neural commonsense knowledge graphs. In *Proceedings of AAAI*.
- Yu Jin Kim, Beong woo Kwak, Youngwook Kim, Reinald Kim Amplayo, Seung won Hwang, and Jinyoung Yeo. 2022. Modularized transfer learning with multiple knowledge graphs for zero-shot commonsense reasoning. In *Proceedings of NAACL*.
- Diederik P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *Proceedings of ICLR*.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2020. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of ACL*.
- Jiangnan Li, Zheng Lin, Peng Fu, and Weiping Wang. 2021a. Past, present, and future: Conversational emotion recognition through structural modeling of psychological knowledge. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 1204–1214.
- Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017. Dailydialog: A manually labelled multi-turn dialogue dataset. In *Proceedings of IJCNLP*.

- Zekang Li, Jinchao Zhang, Zhengcong Fei, Yang Feng, and Jie Zhou. 2021b. Conversations are not flat: Modeling the dynamic information flow across dialogue utterances. In *Proceedings of ACL*.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Proceedings of ACL*.
- Ye Liu, Tao Yang, Zeyu You, Wei Fan, and Philip S Yu. 2020. Commonsense evidence generation and injection in reading comprehension. In *Proceedings of SIGDIAL*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint*.
- David J Mendelsohn. 1994. *Learning to listen: A strategy-based approach for the second language learner*. Dominic Press.
- Ramesh Nallapati, Feifei Zhai, and Bowen Zhou. 2017. Summarunner: A recurrent neural network based sequence model for extractive summarization of documents. In *Proceedings of AAAI*.
- Ramesh Nallapati, Bowen Zhou, Cícero Nogueira dos Santos, Çağlar Gülçehre, and Bing Xiang. 2016. Abstractive text summarization using sequence-to-sequence rnns and beyond. In *Proceedings of CoNLL*.
- 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 EMNLP*.
- Ashwin Ram, Rohit Prasad, Chandra Khatri, Anu Venkatesh, Raefer Gabriel, Qing Liu, Jeff Nunn, Behnam Hedayatnia, Ming Cheng, Ashish Nagar, et al. 2018. Conversational ai: The science behind the alexa prize. *arXiv preprint*.
- Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of EMNLP*.
- Anna Rogers, Olga Kovaleva, and Anna Rumshisky. 2020. A primer in bertology: What we know about how bert works. *Transactions of the Association for Computational Linguistics*, 8:842–866.
- Alexander M Rush, Sumit Chopra, and Jason Weston. 2015. A neural attention model for abstractive sentence summarization. In *Proceedings of EMNLP*.
- Maarten Sap, Ronan Le Bras, Emily Allaway, Chandra Bhagavatula, Nicholas Lourie, Hannah Rashkin, Brendan Roof, Noah A Smith, and Yejin Choi. 2019. Atomic: An atlas of machine commonsense for if-then reasoning. In *Proceedings of AAAI*.
- Abigail See, Peter J Liu, and Christopher D Manning. 2017. Get to the point: Summarization with pointer-generator networks. *arXiv preprint*.
- Vered Shwartz, Peter West, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2020. Unsupervised commonsense question answering with self-talk. In *Proceedings of EMNLP*.
- Robyn Speer, Joshua Chin, and Catherine Havasi. 2017. Conceptnet 5.5: An open multilingual graph of general knowledge. In *Proceedings of AAAI*.
- Leonard Talmy. 1988. Force dynamics in language and cognition. *Cognitive science*, 12(1):49–100.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proceedings of NeurIPS*.
- Han Wang, Yang Liu, Chenguang Zhu, Linjun Shou, Ming Gong, Yichong Xu, and Michael Zeng. 2021. Retrieval enhanced model for commonsense generation. In *Proceedings of ACL Findings*.
- Peter West, Chandra Bhagavatula, Jack Hessel, Jena D Hwang, Liwei Jiang, Ronan Le Bras, Ximing Lu, Sean Welleck, and Yejin Choi. 2022. Symbolic knowledge distillation: from general language models to commonsense models. In *Proceedings of NAACL*.
- Adina Williams, Nikita Nangia, and Samuel R. Bowman. 2018. A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of NAACL*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of EMNLP*.
- Chien-Sheng Wu, Linqing Liu, Wenhao Liu, Pontus Stenetorp, and Caiming Xiong. 2021. Controllable abstractive dialogue summarization with sketch supervision. In *Proceedings of ACL Findings*.
- Felix Wu, Angela Fan, Alexei Baevski, Yann Dauphin, and Michael Auli. 2019. Pay less attention with lightweight and dynamic convolutions. In *Proceedings of ICLR*.
- Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, et al. 2020. Big bird: Transformers for longer sequences. In *Proceedings of NeurIPS*.
- Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter Liu. 2020a. Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. In *Proceedings of ICML*.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020b. Bertscore: Evaluating text generation with bert. In *Proceedings of ICLR*.

Xingxing Zhang, Mirella Lapata, Furu Wei, and Ming Zhou. 2018. Neural latent extractive document summarization. In *Proceedings of EMNLP*.

Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2020c. Dialogpt: Large-scale generative pre-training for conversational response generation. In *Proceedings of ACL*.

Zheng Zhao, Shay B Cohen, and Bonnie Webber. 2020. Reducing quantity hallucinations in abstractive summarization. In *Proceedings of EMNLP Findings*.

Ming Zhong, Pengfei Liu, Yiran Chen, Danqing Wang, Xipeng Qiu, and Xuan-Jing Huang. 2020. Extractive summarization as text matching. In *Proceedings of ACL*.

Pei Zhou, Karthik Gopalakrishnan, Behnam Hedayatnia, Seokhwan Kim, Jay Pujara, Xiang Ren, Yang Liu, and Dilek Hakkani-Tur. 2022. Think before you speak: Explicitly generating implicit commonsense knowledge for response generation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1237–1252.

Pei Zhou, Pegah Jandaghi, Bill Yuchen Lin, Justin Cho, Jay Pujara, and Xiang Ren. 2021a. Probing commonsense explanation in dialogue response generation. In *Proceedings of EMNLP Findings*.

Pei Zhou, Rahul Khanna, Seyeon Lee, Bill Yuchen Lin, Daniel Ho, Jay Pujara, and Xiang Ren. 2021b. Rica: Evaluating robust inference capabilities based on commonsense axioms. In *Proceedings of EMNLP*.

Chenguang Zhu, Yang Liu, Jie Mei, and Michael Zeng. 2021. Mediasum: A large-scale media interview dataset for dialogue summarization. In *Proceedings of NAACL*.

A Baselines

Generative Language Models

- **PointerGenerator** (See et al., 2017) is a RNN-based method designed for text summarization that deploys copy-attention mechanism.
- **DynamicConv** (Wu et al., 2019) is a lightweight convolutional model that can perform competitively to self-attention.
- **Transformer** (Vaswani et al., 2017) is a random-initialized (*i.e.*, not pre-trained) encoder-decoder architecture with self attention and multi-head attention.

Pre-trained Generative Language Models

- **DialoGPT** (Zhang et al., 2020c) is a GPT-2 model pre-trained on open-domain Reddit data designed for response generation.

- **UniLM** (Dong et al., 2019) is a unified language model which can be used for both natural language understanding and generation tasks by pre-trained using three types of language modeling tasks: unidirectional, bidirectional, and sequence-to-sequence prediction on English Wikipedia and BookCorpus.

- **PEGASUS** (Zhang et al., 2020a) is a model specifically designed for summarization tasks where it is pre-trained with an gap-sentence objective. Important sentences are masked from input and is trained to generate the missing parts, similar to an extractive summary approach.

- **BART** (Lewis et al., 2020) is trained by corrupting text with an arbitrary noising function and learning to reconstruct the original text.

- **BART-xsum**³ denotes a BART (Lewis et al., 2020) model fine-tuned on XSUM (Narayan et al., 2018) dataset.

Dialogue Summarization Models

- **CODS** (Wu et al., 2021) finds key phrases, and generates length-controllable summary from the key phrases.

- **D-HGN** (Wu et al., 2021) incorporated commonsense knowledge from ConceptNet (Speer et al., 2017) for dialogue summarization.

- **S-BART** (Chen and Yang, 2021) incorporated discourse relations between utterances, and the connections between speakers and actions within utterances to generate abstractive conversation summarization.

B Implementation Details of Commonsense Generation

To generate commonsense, we use COMET and PARA-COMET. Each commonsense model has different choices in terms of model architecture.

³<https://huggingface.co/facebook/bart-large-xsum>

For COMET, we use BART version among several available versions. Publicly available checkpoints were used for both COMET⁴ and PARACOMET⁵. GPT-2 version was used for PARACOMET. For inference, we use beam search with beam size 5 and 10 for each COMET and PARACOMET, the default setting provided in the public repository. All this procedure is done on one GeForce RTX 3090 GPU.

To investigate the overhead, we measure the time required to generate commonsense inferences in SAMSum. SAMSum, consisted of 14732 samples within the train subset, took 18.3 hours to generate all the needed commonsense inferences. In other words, it took about 4.4719 seconds per dialogue to generate the commonsense. Note that SAMSum has an average of 11.2 turns per dialogue, so that this number could vary depending on how long the given dialogue is.

C Automatic Evaluation Metrics

The following metrics are used for the evaluation of baselines and our models:

- **ROUGE** measures the number of overlapping textual units between generated summary and a set of reference summaries.
- **BERTScore** computes the similarity scores by aligning generated and reference summaries on a token-level based on the output of the BERT-based model. Token alignments are computed greedily with the objective of maximizing the cosine similarity between contextualized token embeddings. We report the F1 score.

D Human Evaluation Metrics

In general, the gold-standard method for evaluating text generation is still human evaluation, where human annotators assess the quality of generated texts. We adopt the following human evaluation metrics:

- **Informativeness**: How well does the generated summary captures the key ideas of the source dialogue?
- **Factual Consistency**: How consistent is the generated summary with respect to the source

dialogue? Does the generated summary contain only statements entailed by the source dialogue?

E Commonsense Selection Methods

We consider two different methods in addition to our similarity-based method to filter commonsense inferences: (i) Random: any random commonsense inferences from 25 possible candidates are chosen for each utterance; (ii) NLI-based: deploy a pre-trained language model that is fine-tuned on a natural language inference (NLI) task, to determine whether a commonsense inference does not contradict with the utterance/sentence.

We use random selection method as a baseline to compare whether filtering helps gain additional performance.

NLI based method is also used by previous works (Gabriel et al., 2021; West et al., 2022) to measure the quality of commonsense inferences. Given a pair of $\{u_i, c_i^r\}$ or $\{y_i, z_i^r\}$, we acquire the probability of ENTAILMENT and CONTRADICT. Then we measure the score as:

$$\text{NLI Score}(u_i, c_i^r) = P(\text{ENTAILMENT}) - P(\text{CONTRADICT}) \quad (7)$$

$$\text{NLI Score}(y_i, z_i^r) = P(\text{ENTAILMENT}) - P(\text{CONTRADICT}) \quad (8)$$

where the commonsense inference with the highest NLI Score is selected. As a result, we obtain the input commonsense $\mathcal{C} = \{c_1, c_2, \dots, c_n\}$ aligned with dialogue \mathcal{D} for additional context and the target commonsense $\mathcal{Z} = \{z_1, z_2, \dots, z_m\}$ aligned with summary \mathcal{Y} for additional supervision.

For NLI-based selection, we use RoBERTa-Large (Liu et al., 2019) which is fine-tuned on MNLI (Williams et al., 2018) to score commonsense candidates. Note that we do not have any label telling which commonsense inference is most plausible to be chosen when given an utterance, therefore, we measure the NLI scores in a zero-shot manner.

As shown in Table 7, using the similarity-based selection method consistently outperforms other methods, regardless of the type of commonsense knowledge model. Since NLI-based method is more intuitive compared to similarity-based methods, and was used in previous works, one might ask

⁴<https://github.com/allenai/comet-atomic-2020>

⁵<https://github.com/skgabriel/paracommet>

Generation Model	Selection Model	SAMSum				DialogSum			
		R-1	R-2	R-L	B-S	R-1	R-2	R-L	B-S
COMET	Random	53.04	27.17	48.49	71.34	46.05	20.46	40.61	70.84
	NLI-based	53.21	28.02	48.85	71.53	45.26	19.94	40.04	70.54
	Similarity-based	53.24	28.10	48.90	71.71	46.31	20.95	41.10	71.71
PARA-COMET	Random	52.95	27.62	48.51	71.45	45.59	20.16	40.23	70.65
	NLI-based	52.99	28.22	48.61	71.69	45.14	20.01	39.98	70.99
	Similarity-based	53.73	28.81	49.50	71.92	46.20	20.39	40.83	71.32

Table 7: Performance of SICK++ by varying the commonsense related methods.

why NLI-based method does not show good performance. We conjecture this due to the complexity of each task. Measuring the relation of inclusion is more complex in nature compared to simply measuring the semantic similarity. Our methodology uses a zero-shot setting, therefore it is harder to reach the standards without supervision. The outperforming choice of commonsense selection method could differ when trained with labeled data, and we leave this to future work.

Also, one might conjecture that using the top-1 selected commonsense inference with the similarity-based method is a copy of the utterance, resulting in inferences with similarity value 1.0 only selected. However, we found that the mean value of the top-1 commonsense inferences are 0.535799, and standard deviation 0.176364. This shows that the commonsense inferences isn't a copy except for a few bad cases. Considering both diversity and quality is important, and we also leave this to future work.

F Choose of Commonsense Relations from COMET

In prior work such as [Chakrabarty et al. \(2022\)](#) and [Li et al. \(2021a\)](#), it is conventional to selectively use a subset of the COMET relations, depending on the characteristics of a target domain and task. In our work, the social-interaction relations such as *xIntent* and *xWant* are most preferred with the best performance as they are strongly relevant to human-human interaction in dialogue.

Dialogue	<p>Frank: Son, will you come home this weekend? Son: Not sure yet. Son: Something happened? Frank: Of course not. Frank: Your mother miss you. Son: I miss her too. Frank: So will you come? Son: I will try. Frank: Good, I will tell your mother that you will come. Son: Oh, dad.. ok I will come.</p>	Commonsense	<p>Frank has to go to work.. son is not sure yet. Person asks to son what happened. Frank doesn't want to be rude. your mother misses you.. son misses his mother. Frank is too shy to ask.. son will try son will come. Person asks if he can come.</p>
Gold Summary	Son is coming to see his parents' this weekend.		
BART-xsum	Son will come home this weekend.		
SICK	Son will come home this weekend. He misses his mother.		
Dialogue	<p>Julie: <file photo> Emily: <3 Julie Love, I'm sending tons of kisses ;*;*;*; Emily: <emoji> Julie: Merry Christmas and a lovely mood throughout the whole year, darling. Emily: Thank you, for you too <3 Julie: Thanks:*) Julie: <file photo> <file photo></p>	<p>Julie sent a photo. to show love. Emily sent a photo. Julie gives a hug Person is thanked. Julie gets a hug. Julie sent a photo.</p>	
Gold	Emily and Julie wish Merry Christmas to each other.		
SICK++	Julie and Emily are exchanging Christmas greeting.		
BART-xsum	Julie sends Emily tons of kisses.		
Dialogue	<p>Stewart: Can you believe he even said that about the forests Stewart: Raking? Really? Shari: Yes... I can believe that this is an ignorant man... Shari: He proves it daily.. This is just one more example! Stewart: He just has no clue... Stewart: I mean, there are so many people dead and all he can think to do is criticize the forestry department? With a totally inappropriate suggestion? Stewart: I can't wait to vote for anyone else but him... Shari: I know what you mean.. Half my friends voted for him just to see what would happen! Well, guess what? Stewart: Yeah, we couldn't go another 4 years with a Democrat.. ...</p>	<p>the forest to be healthy. to think about the situation. Shari doesn't want to be ignorant. Shari wants to be helpful. he has no clue... Shari thinks it's inappropriate. to vote for someone else. Shari votes for him They want to get rid of him.</p>	
Gold	Stewart and Shari find the current president ignorant and incompetent. They hope he gets voted out. Stewart is going to see what possibilities there are of volunteering in the upcoming elections.		
SICK++	Stewart and Shari don't like the fact that the current president raked the forests. They think he's an ignorant man. Shari and Stewart don't want to vote for him, but they have to make the best of it now.		
BART-xsum	Stewart and Shari don't like the way the president is behaving. They are going to vote for anyone else but him.		

Table 8: Successful examples of generated summaries with SICK from DialogSum.

Dialogue	<p>Person1: Are you familiar with American-styled accounting?</p> <p>Person2: I am afraid not.</p> <p>Person2: I haven't worked in an American company so far.</p> <p>Person1: What are the most fundamental concepts underlying the accounting process?</p> <p>Person2: The first is accounting entity, and the second is going concern.</p> <p>Person2: The third is measuring unit.</p> <p>Person2: The fourth is accounting period, and the fifth is objectivity.</p>	Commonsense	<p>Person1 asks PersonY if they are familiar with accounting.</p> <p>Person2 is too afraid.</p> <p>Person2 is too young to work to learn about accounting.</p> <p>Person2 is not qualified.</p> <p>Person2 doesn't know how to measure.</p> <p>Person2 has to be objective.</p>
Gold	Person2 tells Person1 about the fundamental concepts of the accounting process.		
SICK++	Person2 tells Person1 the most fundamental concepts underlying the accounting process.		
BART-xsum	Person1 asks Person2 about American-styled accounting.		
Dialogue	<p>Person1 Oh, it's getting late.</p> <p>Person1 I've got to run.</p> <p>Person1 It was nice talking to you, karren.</p> <p>Person2 Thanks, Tim.</p> <p>Person2 Nice meeting you, too.</p> <p>Person1 I guess we'll see each other around.</p> <p>Person2 Yeah, I hope so.</p> <p>Person2 Well, take it easy.</p> <p>Person1 You too.</p>	<p>Person1 has to go to work..</p> <p>to be running.</p> <p>Person1 calls back.</p> <p>to talk to Tim.</p> <p>to meet PersonY.</p> <p>Person1 calls PersonY.</p> <p>Person2 asks Person2 if they are sure.</p> <p>Person2 has to work.</p> <p>to talk to PersonY.</p>	
Gold	Tim and Karren say goodbye.		
SICK++	Tim and Karren say goodbye to each other.		
BART-xsum	Tim and Karren meet each other for the first time.		
Dialogue	<p>Person1 Taxi!</p> <p>Person2 Where to, sir?</p> <p>Person1 I'd like to go to the railway station please.</p> <p>Person2 Please hop in.</p> <p>Person1 Is it a long run to the station?</p> <p>Person2 It'll take about 20 minutes.</p> <p>Person1 The streets are heavy with traffic at this time of a day, are they?</p> <p>Person2 Yes, they are.</p> <p>Person1 Is it the rush hour now?</p> <p>Person2 Yes, it is.</p> <p>Person2 Are you in a hurry sir?</p> <p>Person1 No, I'm not.</p> <p>Person1 Would you please drive slowly and carefully?</p> <p>Person2 Yes, sir.</p>	<p>Person1 calls a taxi.</p> <p>Person2 asks for directions.</p> <p>to go to the train.</p> <p>PersonY asks PersonY to get in..</p> <p>to go to the station.</p> <p>PersonY asks how long it will take.</p> <p>the traffic is heavy.</p> <p>Person2 doesn't know what they are.</p> <p>Person1 has to go to work.</p> <p>Person2 doesn't know if it is.</p> <p>Person2 asks PersonY to hurry up.</p> <p>No, I'm not.</p> <p>Person1 asks Person2 to slow down.</p> <p>Person2 is asked a question.</p>	
Gold	Person1 takes a taxi to the railway station in the rush hour.		
SICK++	Person1 takes a taxi to the railway station.		
BART-xsum	Person1 calls a taxi to go to the railway station. Person2 tells him it'll take about 20 minutes and drives slowly and carefully.		

Table 9: Successful examples of generated summaries with SICK from DialogSum.

Error Type	Dialogue	Commonsense
Copying Utterance	#Person2#: Have a good day! #Person2#: Well, take it easy. #Person1#: Were you born in Los Angeles?	have a good day. to take it easy. born in Los Angeles.
Factual Consistency	#Person2#: I'm afraid not. #Person2#: But I'm not sleepy, darling. #Person2#: I haven't worked in an American company so far.	Person2 is too afraid. Person2 is sleepy. Person2 is too young to work.
Not Informative	#Person2#: I'm afraid not. #Person1#: No, not much. #Person2#: I've heard this one before.	Person2 is too afraid. Person1 says no. Person2 thinks.

Table 10: Failed examples of generated summaries with SICK.