

Improving Query-Focused Meeting Summarization with Query-Relevant Knowledge

Tiezheng Yu , Ziwei Ji, Pascale Fung

Center for Artificial Intelligence Research (CAiRE)

Department of Electronic and Computer Engineering

The Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong

{tyuah, zjiad}@connect.ust.hk, pascale@ece.ust.hk

Abstract

Query-Focused Meeting Summarization (QFMS) aims to generate a summary of a given meeting transcript conditioned upon a query. The main challenges for QFMS are the long input text length and sparse query-relevant information in the meeting transcript. In this paper, we propose a knowledge-enhanced two-stage framework called Knowledge-Aware Summarizer (KAS) to tackle the challenges. In the first stage, we introduce knowledge-aware scores to improve the query-relevant segment extraction. In the second stage, we incorporate query-relevant knowledge in the summary generation. Experimental results on the QMSum dataset show that our approach achieves state-of-the-art performance. Further analysis proves the competency of our methods in generating relevant and faithful summaries.¹

1 Introduction

Meetings are an essential part of human collaboration and communication. Especially in recent years, the outbreak of Covid-19 has led people to meet online, where most meetings are automatically recorded and transcribed. Query-Focused Meeting Summarization (QFMS) (Zhong et al., 2021) aims to summarize the given meeting transcript conditioned upon a query, which helps people efficiently catch up to the specific part of the meeting they want to know.

There are two main challenges for QFMS: Firstly, meeting transcripts can be so long that current deep learning models cannot encode them at once. Even for models (Beltagy et al., 2020; Xiong et al., 2022) that accept long text input, the cost of computational complexity is enormous. Secondly, query-relevant content is sparsely scattered in the meeting transcripts, meaning a significant part of the transcripts is noisy information when given a

¹The code is released at: <https://github.com/TysonYu/KA-QFMS>

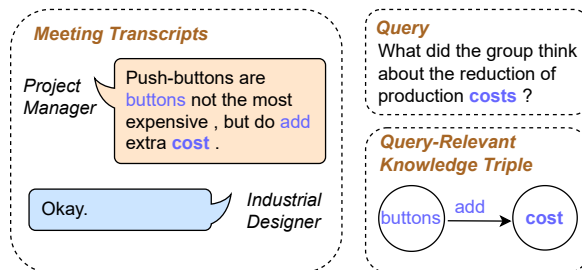


Figure 1: An example of query-relevant knowledge triple extracted from meeting transcripts. The knowledge triple can be used in query-relevant segment extraction as well as summary generation.

particular query. Therefore, the models need to reduce the noisy information’s impact effectively.

In this paper, we focus on the two-stage framework, extracting query-relevant segments from the meeting transcripts and then generating the summary based on the selected content. Compared to the end-to-end approaches (Zhu et al., 2020; Pagnoni et al., 2022) that directly encode the entire meeting transcripts, the two-stage framework is better at keeping computational efficiency and easier to scale up to longer inputs. Specifically, we propose Knowledge-Aware Summarizer (KAS) that incorporates query-relevant knowledge in both stages. In the first stage, we extract knowledge triples from the text segments and introduce knowledge-aware scores to improve segment ranking. In the second stage, the extracted query-relevant knowledge triples are utilized as extra input for the summary generation. We conduct experiments on a QFMS dataset named QMSum (Zhong et al., 2021) and achieve state-of-the-art performance. We further investigate how the different numbers of extracted segments affect the final performance. In addition, we manually evaluate the generation quality regarding fluency, relevance and factual correctness.

Our contributions in this work are threefold: (1) Our work demonstrates the effectiveness of leveraging query-relevant knowledge in QFMS. (2) We

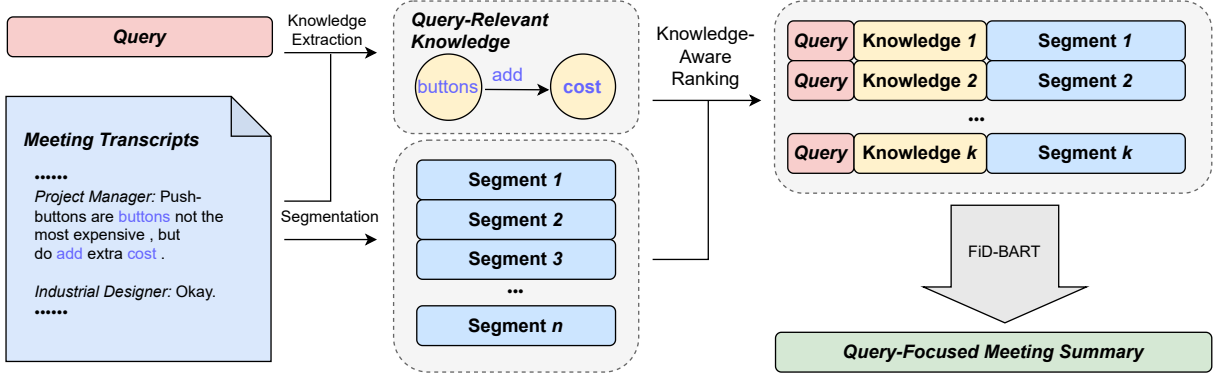


Figure 2: An overview of our proposed framework. We first extract query-relevant knowledge triples based on the query and meeting transcripts and select top- k segments through knowledge-aware ranking. Then we generate the query-focused meeting summary by FiD-BART from query, knowledge and selected meeting transcripts.

propose KAS, a two-stage framework incorporating query-relevant knowledge for QFMS. (3) Experimental results show that our approach achieves state-of-the-art performance on a QFMS dataset (QMSum). Further analysis and human evaluation indicate the advantage of our method.

2 Related Works

Existing QFMS methods can be divided into two categories (Vig et al., 2022): two-stage approaches and end-to-end approaches. Two-stage approaches (Zhong et al., 2021; Vig et al., 2022; Zhang et al., 2022) first extract query-relevant snippets and then generate the summary upon the extracted snippets, while end-to-end approaches (Zhu et al., 2020; Zhong et al., 2022; Pagnoni et al., 2022) directly generate summaries from the whole meeting transcripts. However, both types of methods have their disadvantages. For example, most of the two-stage approaches select query-relevant content based on utterance, which ignores the contextual information between multiple utterances. As for the end-to-end approaches, the computational and memory requirements will increase rapidly when the input text becomes longer, making models challenging to adapt to long meeting transcripts. To our knowledge, we are the first to incorporate query-relevant knowledge in QFMS and demonstrate its effectiveness. Besides, we include multiple utterances in each segment to preserve the contextual information, and our approach can easily extend to process long input text.

3 Methodology

This section presents our two-stage framework. Firstly, we introduce the extractor, which extracts

query-relevant segments and knowledge from the source document. Then, a generative model synthesizes the query, extracted segments, and knowledge into the final summary.

3.1 Knowledge-aware Extractor

Meeting Transcripts Segmentation To preserve the contextual information between utterances, we split the meeting transcripts into segments, and each segment could contain multiple utterances. To do so, given an input meeting transcripts T , we will separate it into n segments $S = \{S_1, S_2, \dots, S_n\}$ where each S_i is fewer than l tokens. Specifically, we feed meeting transcripts to the segment utterance by utterance until it reaches l .

Knowledge-aware Ranking The knowledge-aware ranking approach will select top- k segments according to a combination of semantic search scores and knowledge-aware scores. We apply Multi-QA MPNet (Song et al., 2020) to calculate semantic search scores. Multi-QA MPNet is trained on 215M question-answer pairs from various sources and domains, including Stack Exchange, MS MARCO (Nguyen et al., 2016), WikiAnswers (Fader et al., 2014) and many more. Given the query and the segments, the model outputs 768-dimension vectors to represent them. The semantic search score for each segment is computed according to the cosine similarity:

$$Score_{se} = \frac{MPNet(Q) \cdot MPNet(S_i)}{\|MPNet(Q)\| \|MPNet(S_i)\|} \quad (1)$$

To compute the knowledge-aware score for each segment, we first use OpenIE (Angeli et al., 2015) to extractive knowledge triples. Then, to filter out

Model	ROUGE-1	ROUGE-2	ROUGE-L	Average
<i>Two-stage</i>				
MaRGE (Xu and Lapata, 2021)	31.99	8.97	27.93	22.96
DYLE (Mao et al., 2022)	34.42	9.71	30.10	24.74
SUMM ^N (Zhang et al., 2022)	34.03	9.28	29.48	24.26
RelReg (Vig et al., 2022)	34.91	11.91	30.73	25.85
RelReg-W (Vig et al., 2022)	36.45	12.81	32.28	27.18
<i>End-to-end</i>				
LED (Beltagy et al., 2020)	31.60	7.80	20.50	19.97
DialogLM (Zhong et al., 2022)	34.50	9.92	30.27	24.90
SegEnc-W (Vig et al., 2022)	37.80	13.43	33.38	28.20
BART-LS (Xiong et al., 2022)	37.90	12.10	33.10	27.70
SOCRATIC (Pagnoni et al., 2022)	38.06	13.74	33.51	28.44
KAS (Ours)	38.80	14.01	34.26	29.02

Table 1: Results on QMSum test set. The previous works can be divided into two-stage and end-to-end approaches.

the triples irrelevant to the query, we only keep the triples that contain overlapping words with the query. The knowledge-aware score is obtained by L2 normalizing the number of the remaining triples m_i in each segment:

$$Score_{ka} = \frac{m_i}{\sqrt{\sum_{i=1}^n |m_i|^2}} \quad (2)$$

Finally, we calculate the ranking score by summing the semantic search score and the knowledge-aware score:

$$Score_{rank} = Score_{se} + Score_{ka} \quad (3)$$

The segment with the top- k ranking score will be selected for the next stage of summary generation. We denote the remaining segments as $S = \{S_1, S_2, \dots, S_k\}$.

3.2 Generator

We choose BART-large (Lewis et al., 2020), a transformer-based (Vaswani et al., 2017) generative pre-trained language model, as our backbone model for the generator because of its remarkable performance on text summarization benchmarks. Following the idea of Fusion-in-Decoder (FiD) and its applications in generation tasks (Izcard and Grave, 2021; Su et al., 2022; Vig et al., 2022), we employ FiD-BART for encoding multiple segments independently in the encoder and fuse information from all segments in the decoder jointly through the encoder-decoder attention.

To incorporate the extracted knowledge in the summary generation, we use the knowledge as extra inputs other than the query and segments. In detail, we remove the stop words in the knowledge triples and then merge all the remaining words of knowledge triples in each segment as a set of knowledge phrases. The concatenation of each segment S_i , knowledge phrases K_i and the query Q will be processed by FiD-BART encoder:

$$h_{enc}^i = Encoder(Q \oplus K_i \oplus S_i) \quad (4)$$

Finally, the decoder performs encoder-decoder attention over the concatenation of all segments' encoder outputs. In this way, the computational complexity grows linearly with the number of segments rather than quadratically, while jointly processing all segments in the decoder enables the model to aggregate information from multiple segments.

4 Experiments

4.1 Experimental Setup

We choose the top 12 text segments in the segment selection, with each fewer than 512 tokens. For the summary generator, inspired by the effectiveness of pre-finetuning models on relevant datasets to transfer task-related knowledge in the abstractive summarization (Yu et al., 2021; Vig et al., 2022), we initialize our BART-large model using the checkpoint² pre-finetuned on WikiSum (Liu et al., 2018).

²<https://github.com/salesforce/query-focused-sum>

model	ROUGE-1	ROUGE-2	ROUGE-L
KAS ($k=4$)	36.46	12.06	31.75
KAS ($k=8$)	37.86	13.75	33.30
KAS ($k=12$)	38.80	14.01	34.26

Table 2: Results Our KAS on the QMSum test set with different number of k . k denotes the number of segments that are selected in the first stage.

model	R-1	R-2	R-L	Entity F-1
KAS	38.80	14.01	34.26	35.59
w/o KA in ranking	37.71	12.96	32.94	34.03
w/o KA in generator	37.52	13.57	33.29	35.06

Table 3: Ablation results on the QMSum test set without knowledge-Aware (KA). Our KA modules improve the performance of KAS on both stages.

See Appendix A and B for more details of the experimental setup and baselines.

We evaluate the models on the QMSum (Zhong et al., 2021) dataset, which consists of 1,808 query-summary pairs over 232 meetings from product design, academic, and political committee domains. We report ROUGE³ (Lin, 2004) as the automatic evaluation results.

4.2 Main Results

We compare our proposed model with strong baselines and previous state-of-the-art models. As shown in Table 1, our approach outperforms both the two-stage and end-to-end methods by a large margin on all evaluation metrics. We also investigate how the number of selected segments in the first stage affects the final performance. As shown in Table 2, when the number of selected segments grows, the quality of the final summary also improves. Previous end-to-end approaches usually outperform two-stage approaches because they encode the entire meeting transcripts, which conduct lots of computing resources. Our approach performs better by encoding 12 segments of input (6,144 tokens) in the second stage generation than the state-of-the-art end-to-end approach that directly encodes 16384 tokens. Therefore, we strike a balance between performance and efficiency, which is essential in real-world applications. Besides, our approach can easily extend to processing long input text since the ranking method is unsupervised.

³<https://github.com/pltrdy/files2rouge>

model	Fluency	Relevance	Factual Correctness
SegEnc-W	4.68	4.60	3.76
KAS	4.79	4.68	4.47
Ground truth	4.92	4.82	4.60

Table 4: Human evaluation results on QMSum test set.

4.3 Ablation Study

We conduct an ablation study to investigate the contribution of the knowledge-aware modules by removing the knowledge-based scoring in the ranking (w/o KA in ranking) and the knowledge input in the summary generation (w/o KA in generator), respectively. Besides, we evaluate the entity-level factual consistency (Nan et al., 2021) of the summary to test the effectiveness of our knowledge-aware modules in keeping the knowledge entities in generated summaries. Specifically, we report the F-1 scores of the entity overlap (Entity F-1) between the source and the generated summary. Table 3 presents that the model’s performance decreases in both metrics, especially the Entity F-1, without the knowledge-aware module, which indicates the effectiveness of our method.

4.4 Human Evaluation

We further conduct a human evaluation to assess the models on fluency, relevance, and factual correctness. We randomly select 50 samples from the QMSum test set and ask three annotators to score the summary from one to five, with higher scores being better. We compare SegEnc-W because it is the best publicly available model. Table 4 illustrates our approach slightly outperforms SegEnc-W according to fluency and relevance and achieves a significantly better score than SegEnc-W on factual correctness with p -value <0.05 . Since SegEnc-W and our methods both use BART based generators, it is in line with our expectation that we reach similar fluency and relevance scores. Meanwhile, our proposed query-relevant knowledge can help reduce hallucination and generate more factual correctness summaries. More details of the human evaluation setup are in Appendix C.

5 Conclusion

In this paper, we propose a two-stage framework named KAS that incorporates query-relevant knowledge in both stages. Extensive experimental results on the QMSum dataset show the effectiveness of our method. We further conduct de-

tailed analysis and human evaluation to prove our method’s capacity to generate fluency, relevant and faithful summaries.

6 Limitations

Our method is trained and tested on the only publicly available QFMS dataset named QMSum. QMSum consists of three different domains (Academic, Committee and Product), which makes the evaluation robust. However, QMSum only contains 1,808 samples, which is relatively small. We hope larger QFMS datasets be proposed to accelerate the development of this field.

In the first stage of our approach, we extract a fixed length (6,144 tokens) of the meeting transcripts as the second stage input text. Therefore, the model’s performance could be affected since some query-relevant information could be cut off in the first stage.

7 Ethics Statement

Although our approach can generate query-focused meeting summaries and achieve a much better factual correctness score than other models, we can not avoid generating hallucinated content from the generative models. We recommend that when our approach is deployed in real-world applications, additional post-processing should be carried out to remove unreliable summaries. In addition, we will indicate that the summary we provide is for reference only, and users need to check the original meeting transcripts to get accurate information.

References

- Gabor Angeli, Melvin Jose Johnson Premkumar, and Christopher D Manning. 2015. Leveraging linguistic structure for open domain information extraction. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 344–354.
- Iz Beltagy, Matthew E Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*.
- Anthony Fader, Luke Zettlemoyer, and Oren Etzioni. 2014. Open question answering over curated and extracted knowledge bases. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1156–1165.
- Gautier Izacard and Édouard Grave. 2021. Leveraging passage retrieval with generative models for open domain question answering. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 874–880.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- 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.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Peter J Liu, Mohammad Saleh, Etienne Pot, Ben Goodrich, Ryan Sepassi, Lukasz Kaiser, and Noam Shazeer. 2018. Generating wikipedia by summarizing long sequences. *arXiv preprint arXiv:1801.10198*.
- Ziming Mao, Chen Henry Wu, Ansong Ni, Yusen Zhang, Rui Zhang, Tao Yu, Budhaditya Deb, Chenguang Zhu, Ahmed Awadallah, and Dragomir Radev. 2022. Dyle: Dynamic latent extraction for abstractive long-input summarization. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1687–1698.
- Feng Nan, Ramesh Nallapati, Zhiguo Wang, Cicero dos Santos, Henghui Zhu, Dejiao Zhang, Kathleen Mckeown, and Bing Xiang. 2021. Entity-level factual consistency of abstractive text summarization. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2727–2733.
- Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. 2016. Ms marco: A human generated machine reading comprehension dataset. *choice*, 2640:660.
- Artidoro Pagnoni, Alexander R Fabbri, Wojciech Kryściński, and Chien-Sheng Wu. 2022. Socratic pre-training: Question-driven pretraining for controllable summarization. *arXiv preprint arXiv:2212.10449*.
- Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. 2020. Mpnet: Masked and permuted pre-training for language understanding. *Advances in Neural Information Processing Systems*, 33:16857–16867.
- Dan Su, Xiaoguang Li, Jindi Zhang, Lifeng Shang, Xin Jiang, Qun Liu, and Pascale Fung. 2022. Read before generate! faithful long form question answering with machine reading. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 744–756.

- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Jesse Vig, Alexander Richard Fabbri, Wojciech Kryściński, Chien-Sheng Wu, and Wenhao Liu. 2022. Exploring neural models for query-focused summarization. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 1455–1468.
- Wenhan Xiong, Anchit Gupta, Shubham Toshniwal, Yashar Mehdad, and Wen-tau Yih. 2022. Adapting pretrained text-to-text models for long text sequences. *arXiv preprint arXiv:2209.10052*.
- Yumo Xu and Mirella Lapata. 2021. Generating query focused summaries from query-free resources. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6096–6109.
- Tiezheng Yu, Zihan Liu, and Pascale Fung. 2021. Adaptsum: Towards low-resource domain adaptation for abstractive summarization. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5892–5904.
- Yusen Zhang, Ansong Ni, Ziming Mao, Chen Henry Wu, Chenguang Zhu, Budhaditya Deb, Ahmed Awadallah, Dragomir Radev, and Rui Zhang. 2022. Summn: A multi-stage summarization framework for long input dialogues and documents: A multi-stage summarization framework for long input dialogues and documents. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1592–1604.
- Ming Zhong, Yang Liu, Yichong Xu, Chenguang Zhu, and Michael Zeng. 2022. Dialoglm: Pre-trained model for long dialogue understanding and summarization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 11765–11773.
- Ming Zhong, Da Yin, Tao Yu, Ahmad Zaidi, Mutethia Mutuma, Rahul Jha, Ahmed Hassan, Asli Celikyilmaz, Yang Liu, Xipeng Qiu, et al. 2021. Qmsum: A new benchmark for query-based multi-domain meeting summarization. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5905–5921.
- Chenguang Zhu, Ruochen Xu, Michael Zeng, and Xuedong Huang. 2020. A hierarchical network for abstractive meeting summarization with cross-domain pretraining. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 194–203.

A Experimental Details

We use the off-the-shelf model Multi-QA MPNet⁴ for semantic searching. We initialize the generator from the off-the-shelf BART-large model through Hugging Face. We use learning rates $6e^{-5}$ following (Lewis et al., 2020) and Adam optimizer (Kingma and Ba, 2014) to fine-tune KAS. We use a batch size of one and train all models on one RTX 3090 for ten epochs. During decoding, we use beam search with a beam size of five and decode until an end-of-sequence token is emitted. The final results are the average test set performance on the best three checkpoints in the validation set. In addition, the QMSum dataset has MIT License, so we can freely use it.

B Baselines Details

We conduct both two-stage and end-to-end baselines for comparison. MaRGE (Xu and Lapata, 2021) used **Masked ROUGE** extractor in the first stage. DYLE (Mao et al., 2022) is a dynamic latent extraction approach for abstractive long-input summarization. SUMM^N (Zhang et al., 2022) is a multi-stage summarization framework for long input dialogues and documents. RelReg (RElevance REGression) (Vig et al., 2022) is similar to MaRGE, which trains a relevance prediction model directly on QFS data using the original, non-masked query. RelReg-W (Vig et al., 2022) uses the same framework as RelReg but initializes the generator from the WikiSum pre-trained checkpoint. LED (Beltagy et al., 2020) is an encoder-decoder transformer-based model that employs efficient attention to process long input text. DialogLM (Zhong et al., 2022) is a pre-trained neural encoder-decoder model for long dialogue understanding and summarization, and it uses a hybrid attention approach by combining sparse local attention with dense global attention. SegEnc-W (Vig et al., 2022) is a fusion-in-decoder method that initialized the model from WikiSum (Liu et al., 2018). BART-LS (Xiong et al., 2022) used pooling-augmented blockwise attention to improve the efficiency and pre-trains the model with a masked-span prediction task with spans of varying lengths. SOCRATIC (Pagnoni et al., 2022) uses a question-driven, unsupervised pretraining objective to improve controllability in summarization tasks, and it

⁴<https://huggingface.co/sentence-transformers/multi-qa-mpnet-base-cos-v1>

follows the SegEnc framework.

C Human Evaluation

In Table 4, we conduct a human evaluation of the fluency, relevancy and factual correctness of the generated summaries from the QMSum dataset. In detail, we randomly sample 50 summaries from KAS, SegEnc-W and the ground truth summary from the same meeting transcripts for comparison. Assessments are scored on a scale of one to five, with higher scores being better. Fluency means the summary is grammatically correct and contextually coherent. Relevancy means the summary is relevant to the query. Factual correctness means the summary is faithful according to the information given in the meeting transcripts and does not contain hallucinations. We assign each summary to three annotators and take the average score as the final result. In total, we used three annotators from Hong Kong, and all annotators voluntarily participated in the human evaluation. All annotators agree to have their evaluation results included as part of this paper’s results.