

An Exploratory Study on Long Dialogue Summarization: What Works and What’s Next

Yusen Zhang^{*♣} Ansong Ni^{*†} Tao Yu[†] Rui Zhang^{*♣}
Chenguang Zhu[◇] Budhaditya Deb[◇] Asli Celikyilmaz[‡]
Ahmed Hassan Awadallah[◇] Dragomir Radev[†]

[♣]Penn State University [†]Yale University [◇]Microsoft Research [‡]Facebook AI Research
{yfbz5488, rmz5227}@psu.edu asli.ca@live.com
{ansong.ni, tao.yu, dragomir.radev}@yale.edu
{chezhu, Budha.Deb, hassanam}@microsoft.com

Abstract

Dialogue summarization helps readers capture salient information from long conversations in meetings, interviews, and TV series. However, real-world dialogues pose a great challenge to current summarization models, as the dialogue length typically exceeds the input limits imposed by recent transformer-based pre-trained models, and the interactive nature of dialogues makes relevant information more context-dependent and sparsely distributed than news articles. In this work, we perform a comprehensive study on long dialogue summarization by investigating three strategies to deal with the lengthy input problem and locate relevant information: (1) extended transformer models such as Longformer, (2) *retrieve-then-summarize* pipeline models with several dialogue utterance retrieval methods, and (3) hierarchical dialogue encoding models such as HMNet. Our experimental results on three long dialogue datasets (QMSum, MediaSum, SummScreen) show that the *retrieve-then-summarize* pipeline models yield the best performance. We also demonstrate that the summary quality can be further improved with a stronger retrieval model and pretraining on proper external summarization datasets.

1 Introduction

Large amount of dialogue data have been produced in meetings, TV series, and interviews (Chen et al., 2021; Zhong et al., 2021; Zhu et al., 2021). Dialogue summarization aims to generate a short summary for long dialogues to help the readers capture important information more efficiently.

A number of existing works on dialogue summarization focus on extracting the main events of a short conversation (Gliwa et al., 2019; Rohde et al., 2021). However, unlike the short dialogues

which contains less than 20 utterances, some tasks for summarizing much longer dialogues have been proposed recently (Chen et al., 2021; Zhong et al., 2021). These datasets are usually derived from meetings and interviews, with hundreds of turns in each dialogue. The length of such dialogues typically exceeds the input limits imposed by recent transformer-based models (Lewis et al., 2020), making it difficult to train an end-to-end summarization model for such tasks. This poses the challenge: *How can we effectively use the current neural summarization models on dialogues that greatly exceed their length limits?*

Additionally, compared with document summarization, dialogues are interactive in nature, makes it more context-dependent and the information in dialogues is more sparsely distributed. Besides, the informal language used in dialogues leads to difficulties in modeling relevance and salience. To solve these issues, hierarchical methods are proposed to model the dialogues at turn level (Zhu et al., 2020a; Rohde et al., 2021). However, generating a short summary that contains all the salient information remains challenging.

In this paper, we systematically investigate these issues on dialog summarization: we first explore the various solutions to the lengthy input problem. Then, we analyze and compare the methods to improve generic summarization models on challenging dialogue datasets. To address the long input issue, we investigate extended transformer models such as Longformer (Beltagy et al., 2020), and several dialogue utterance retrieval methods for a *retrieve-then-summarize* pipeline model, as well as hierarchical dialogue encoding models. For the specific challenges in dialogues, we explore different datasets for pretraining to test the transferability between similar summarization tasks. We evaluate these models on three recent long dialogue summarization datasets: QMSum for meetings (Zhong et al., 2021), MediaSum for inter-

^{*}Equal Contribution.

[‡]The work was done when Asli was at MSR.

views (Zhu et al., 2021), SummScreen for TV series transcripts (Chen et al., 2021). In our experiments, we find that the pipeline method with a dialogue utterance retrieval model yields the best performance, and it can be further improved with a stronger retrieval model. Our experiment results also suggest that pretraining on proper external summarization datasets can effectively improve the performance of dialogue summarization models.

2 Related Work

Long Sequence Summarization Recent summarization models are based on Transformer (Vaswani et al., 2017) that has a quadratic time and memory complexity with respect to the input length, preventing it from being used for longer sequences. To address this issue, Beltagy et al. (2020) used the sliding window and global attention, while Zaheer et al. (2020) used a combination of random, sliding window and global attention mechanism to reduce the quadratic complexity to close-linear. Previous benchmarks for long sequence summarization mostly focus on documents instead of dialogues: PUBMED and ARXIV (Cohan et al., 2018) consists of scientific papers which are typically very long; BILLSUM (Kornilova and Eidelman, 2019) is a corpus of U.S. Congressional bills and their summaries; BIGPATENT (Sharma et al., 2019) contains 1.3 million U.S. patent files and human-written summaries.

Dialogue Summarization Dialogue summarization aims to generate concise summaries for dialogues, such as meetings (McCowan et al., 2005; Janin et al., 2003; Zhong et al., 2021; Shang et al., 2018; Zhu et al., 2020a), TV series (Chen et al., 2021), interviews (Zhu et al., 2021), and chat (Gliwa et al., 2019; Zhao et al., 2020; Chen and Yang, 2021). Some summarization datasets (not limited to dialogues) contain queries asking for summarizing specific parts of dialogues (Zhong et al., 2021; Nema et al., 2017), while others only need to summarize whole dialogues (Chen et al., 2021; Gliwa et al., 2019; Hermann et al., 2015). As for dialogue summarization models, Zhu et al. (2020b) described a hierarchical model for both inner- and cross-utterance attention, while Chen and Yang (2020) proposed a multi-view decoder to leverage different extracted views of dialogues, such as topic view and stage view.

	QMSum	SummScreen	MediaSum*
Source	Meeting	TV Series	Interviews
Query-based	YES	NO	NO
# examples	1.8k	26.9k	463.6k
# input tokens	9069.8	6612.5	1553.7
# summary tokens	69.6	337.4	14.4
# speakers	9.2	28.3	6.5

Table 1: Comparison between three long dialogue summarization datasets we mainly study in this work. Numbers in the table are averaged across all samples. (*: MediaSum is only used for pretraining)

3 Methodology

In this section, we will introduce the dataset used to evaluate and pretrain the model, two types of summary models, and the details of the experiment setup.

3.1 Datasets

To explore the problems in long dialogue summarization, we leverage three different long dialogue summarization tasks as main datasets:

QMSum (Zhong et al., 2021) is a query-based multi-domain meeting summarization dataset annotated by humans. It contains 1,808 queries together with 232 long meeting transcripts, with topics as software product, academics, and committee meetings. QMSum also contains annotated gold spans which could be used as the gold labels for training the retrievers;

MediaSum (Zhu et al., 2021) is a large-scale media interview dataset consisting of 463.6K transcripts collected from NPR and CNN. Because MediaSum contains short summaries, i.e. only a short sentence representing the topic, we only use this dataset for pretraining and analysis. Due to the huge size of this dataset, 20k samples are randomly extracted for pretraining;

SummScreen (Chen et al., 2021) is a dialogue summarization dataset consisting of 26.9k pairs of TV series transcripts and human-annotated summaries. It comes with two sources for recaps, and in this work, we choose one of them, i.e. “Forever Dreaming”, for which we call SummScreen-FD as our benchmark.

Tab. 1 shows the statistics for these three long dialogue datasets. Additionally, we also consider **CNN/Dailymail** (Hermann et al., 2015) (CNN/DM), **XSum** (Narayan et al., 2018), and **SAMSum** (Gliwa et al., 2019) as datasets for pretraining in our experiments.

3.2 Models

3.2.1 Retrieve-then-summarize Pipeline

Dialogues tend to be relatively long, and most existing summarization models cannot process such long inputs. The two-stage *retrieve-then-summarize* pipeline first retrieves the most relevant subtext in the dialogue and then feeds to a summarizer. We experiment with the following retrievers:

- **TF-IDF** (Jones, 1972) Based on bag-of-words representation, TF-IDF measures term frequency (TF) and normalizes them with inverted document frequency (IDF);
- **BM25** (Robertson and Zaragoza, 2009) Similar to TF-IDF but accounts for document length and term saturation;
- **Locator**¹ The utterance locator model proposed by Zhong et al. (2021) using convolution neural networks with BERT (Devlin et al., 2019).

For TF-IDF and BM25, we limit the number of retrieved utterances to be at most 10% of the whole dialogue, while we directly use the utterances predictor by Locator in its setting. After retrieval, we use the BART-large model fine-tuned on the output of the various retrievers to produce the summary.

3.2.2 End-to-end Summarization Models

To study how current state-of-the-art neural summarizers perform on long dialogue summarization, we choose the following three models:

BART (Lewis et al., 2020) is a transformer-based encoder-decoder model which obtains a number of state-of-the-art results on various text generation tasks. We use this model as our baseline summarization model for studying its ablations under different settings. The maximum number of input tokens is 1,024 so we truncate the input when it exceeds such limit.²

HMNet (Zhu et al., 2020a) is a hierarchical network for dialogue summarization. It models the structure of the dialogue, using a token level encoder to encode each sentence and a turn level encoder for aggregating each turn. We use HMNet as a representative for the hierarchical type of models and compare it with other baselines. Due

¹We obtained the locator output from the original authors.

²We also tried to extend the positional embeddings to 2,048 for BART to accept longer input but found the results to be worse in our case.

to the limitation of the memory cost, we constrain the maximum number of tokens to be 8,192 for HMNet, which is 8x as large as BART mentioned above.

Longformer (Beltagy et al., 2020) adapts the self-attention mechanism from full attention matrix to sliding window attention + global attention, which is more memory efficient. Longformer can accept up to 16K tokens and has shown improvement over long document summarization using its long-encoder-decoder (LED) variant. We allow the maximum input of 4,096 tokens for Longformer and cutoff the rest of the input, as we found further increasing such limit yields no improvements.

To incorporate queries in QMSum for these end-to-end models, we simply append the queries to the front of the meeting transcripts, as it is a standard practice for query-based summarization and also question answering (Devlin et al., 2019).

3.3 Experiment Setup

For a fair comparison between all models, we fit all of the models into the same RTX 8000 GPU with 48 GiB of GPU memory. We adopt the fairseq³ implementation for BART, and the original code base for both Longformer⁴ and HMNet⁵. We inherit the hyperparameters for all those models for fine-tuning in our experiments.⁶ Our most expensive experiments are fine-tuning for HMNet and Longformer, which take around 8 hours, while the runtime for BART model is less than one hour. We use ROUGE (Lin, 2004) as our main evaluation metric and `pyrouge` library⁷ as the ROUGE implementation throughout all experiments.

4 Result and Analysis

Here we demonstrate our findings in four corresponding subsections. We also show some concrete examples and perform qualitative analysis in § 4.5

4.1 Dealing with Long Dialogues

We compare several methods for addressing the long input issue for dialogue summarization, including different utterance retrieval methods describe in § 3.2.1 for a retrieve-then-summarize framework, heuristics for shortening the dialogue

³<https://github.com/pytorch/fairseq>

⁴<https://github.com/allenai/longformer>

⁵<https://github.com/microsoft/HMNet>

⁶For more implementation details, please refer to our experiment code: <https://github.com/chatc/LongDialSumm>.

⁷<https://github.com/bheinzerling/pyrouge>

Methods	R-1	R-2	R-L
Retrieve-then-summarize			
Random	31.1	7.9	20.9
TF-IDF	32.5	8.5	21.4
BM25	32.9	9.0	22.0
Locator	29.9	7.6	19.6
Gold span	36.6	14.0	25.5
End-to-end (Cutoff at max # tokens)			
BART-large(1024)	32.6	8.7	21.6
Longformer-large(4096)	31.6	7.8	20.5

Table 2: Comparison of different methods for addressing the length of the dialogues on QMSum. All "retrieve-then-summarize" pipelines use BART-Large as a backend for summarization. "Gold span" denotes the annotated relevant turns in QMSum.

as well as baseline methods to establish reasonable bounds. From Tab. 2, we can see that even in the query-based dialogue summarization with QMSum, randomly selecting utterances still presents a strong baseline. Over different modeling choices, the retrieve-then-summarize framework generally works better than end-to-end learning with dialogue cutoff at maximum input length. We do not observe an advantage of using Longformer over the BART model. This raises the question on whether all utterances in the dialogue are needed to produce a good summary or irrelevant utterances would add more noise. Moreover, we notice that all these methods present a non-trivial gap with the summarization performance on the gold span, which uses relevant utterances annotated by humans. This suggests that there is plenty room for improvement if a better utterance retrieval method is developed.

4.2 Robustness to Input Length

As we discussed, some dialogues (e.g., QMSum) contain more than 20k tokens. They exceed the input limitation of most existing summarization models. In this section, we further analyze the performance of summarization models as the input length changes. To compare the robustness between two types of models (mainly BART and HMNet), we divide the test dialogues by the number of tokens. As we can see in Fig. 1, the performance of the BART model decreases sharply when the dialogue input becomes longer while the HMNet shows the opposite effect. This could be the result of their unique properties: BART is pretrained on the datasets with a limited length (*i.e.*, 1,024) and the input has to be truncated to fit the limitation, while HMNet obtains more information when the

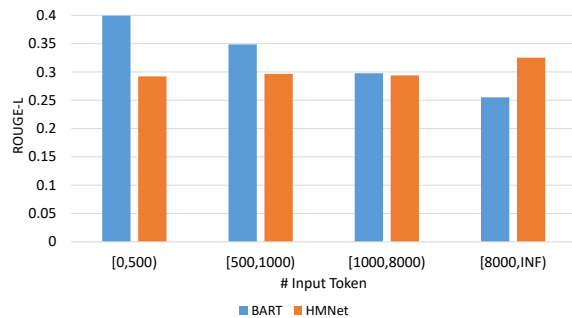


Figure 1: The ROUGE-L F1 scores of BART and HMNet on QMSum stratified by the number of input dialogue tokens.

input is longer. However, the overall performance of HMNet is worst than BART.

4.3 Incorporating Queries

Certain dialogue summarization tasks, such as QMSum, require generating a summary based on a specific question about the dialogue (e.g., opinion of a speaker or conclusion to a topic). In this section, we study the influence of incorporating queries in dialogue summarization. Tab. 4 shows the performance of two models, BART and HMNet, on QMSum with and without queries at the beginning of the input. For the input to the two models, we use the gold relevant text spans given a query in QMSum to avoid the influences of retrieval models. The results show that encoding queries has a large impact on both types of models, especially for BART, even if the gold utterances are given.

4.4 Transfer Ability between Different Tasks

Pretraining has been shown effective for document summarization by introducing external knowledge from other similar tasks (Hermann et al., 2015; Fabbri et al., 2019). We hypothesize that it is especially important for dialogue summarization because the dataset size is usually small. Therefore, we study the transfer learning between different dialogue summarization tasks via pretraining. Tab. 3 shows the performance of BART-large models that are pretrained using different datasets and later fine-tuned on QMSum and SummScreenFD. The results show that BART-large pretrained on CNN/Dailymail dataset (BART-CNN) yields the best performance after finetuning, though CNN/Dailymail consists of News articles and is not in dialogue format. We also note that pretraining on external datasets can also hurt the performance, and thus such pretraining datasets need to

	QMSum			SummScreen-FD		
	ROUGE-1	ROUGE-2	ROUGE-L	ROUGE-1	ROUGE-2	ROUGE-L
BART-Large	36.56	14.05	25.54	27.12	4.88	16.82
+ XSum	34.90	13.49	24.90	27.17	4.59	17.02
+ MediaSum	34.23	13.06	25.21	27.73	5.03	17.09
+ CNN/DM	39.88	15.94	28.02	28.86	5.55	17.39
+ CNN/DM-SAMSum	35.46	12.52	24.62	28.15	5.41	17.25
+ CNN/DM-MediaSum	36.79	13.69	25.94	28.68	5.31	17.42

Table 3: The performance of BART-large models that are pretrained on various summarization datasets.

	ROUGE-1	ROUGE-2	ROUGE-L
BART-CNN			
w/o Query	34.48	11.5	23.11
w/ Query	39.88	15.94	28.02
HMNet			
w/o Query	35.1	10.1	30.8
w/ Query	36.8	10.9	31.9

Table 4: The performance comparison between BART and HMNet models on the query-based meeting summarization QMSum dataset.

be carefully chosen.

We also analyze the performance of BART-large by pretraining it on more than one dataset to test if BART-large can be further improved. We use the BART-large model pretrained on CNN/DM (BART-CNN) as baseline model since BART-CNN yields the best performance compared with the others. And then pretrain the same BART-CNN model on SAMSum and MediaSum separately. However, Tab. 3 shows that after pretraining BART-CNN on these two datasets, ROUGE scores decrease sharply on QMSum dataset, and lightly on SummScreen-FD dataset except for ROUGE-L. This result demonstrates that pretraining on multiple dataset may not further improve the performance of the pretrained models.

4.5 Case Study

We exam several summaries generated by BART-large model pretrained on three different datasets. We found that the BART-CNN model yields the best output with the least number of syntax errors and the closest content to the desired ones, while the output of BART-MediaSum model is usually shorter than Gold resulting in incomplete generation, and BART-XSum model usually predicts summaries with errors and duplication. This could be the result of data bias of pretraining datasets — Summaries in MediaSum and XSum are shorter than CNN/DM. However, despite the better performance of BART-CNN model, these cut-off models

fail to predict some part of the gold summary when the number of tokens in input dialogue is larger than the maximum input length of the model. For concrete examples, please refer to Appendix A.

5 Conclusion and Future Work

We first explore the lengthy input problem of dialogue summarization through experiments on transformers and retrieval models. We conclude that the retrieve-summarize pipeline results in the best performance. Then, the experiments demonstrate the important role of queries and robustness to input length for different types of models. We found that adding a single query sentence in the input greatly improves ROUGE scores on QMSum. Additionally, BART performs worse when the input is beyond 512 tokens, even with extended positional embeddings; on the contrary, the hierarchical model performs better for longer inputs. We also test the transferability of summarization datasets by pretraining the language model on similar tasks. We conclude that the BART-large model pretrained on CNN/DM yields the best performance on both QMSum and SummScreen-FD.

For future work on solving the long input problem, we found that using an utterance retrieval model for summarization is a promising direction, yet modeling relevance between query and dialogue utterances remains a challenging task. And for the summarization models, it is worth exploring methods to 1) pretrain on valuable datasets for dialogue summarization, 2) better fuse the queries into the neural models, and 3) make the model robust to the input length (like HMNet) and maintain the high performance in the meantime (like BART).

Acknowledgments

The authors would like to thank Ming Zhong, Da Yin, Yang Liu for their discussions and anonymous reviewers for their helpful comments. This work is supported in part by a grant from Microsoft Research.

References

- Iz Beltagy, Matthew E Peters, and Arman Cohan. 2020. **Longformer: The long-document transformer**. *ArXiv preprint*, abs/2004.05150.
- 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.
- Jiaao Chen and Diyi Yang. 2021. **Structure-aware abstractive conversation summarization via discourse and action graphs**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1380–1391, Online. Association for Computational Linguistics.
- Mingda Chen, Zewei Chu, Sam Wiseman, and Kevin Gimpel. 2021. **Summscreen: A dataset for abstractive screenplay summarization**. *ArXiv preprint*, abs/2104.07091.
- Arman Cohan, Franck Dernoncourt, Doo Soon Kim, Trung Bui, Seokhwan Kim, Walter Chang, and Nazli Goharian. 2018. **A discourse-aware attention model for abstractive summarization of long documents**. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 615–621, New Orleans, Louisiana. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. **BERT: Pre-training of deep bidirectional transformers for language understanding**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Alexander Fabbri, Irene Li, Tianwei She, Suyi Li, and Dragomir Radev. 2019. **Multi-news: A large-scale multi-document summarization dataset and abstractive hierarchical model**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1074–1084, Florence, Italy. Association for Computational Linguistics.
- 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.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. **Teaching machines to read and comprehend**. In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 1693–1701.
- Adam Janin, Don Baron, Jane Edwards, Dan Ellis, David Gelbart, Nelson Morgan, Barbara Peskin, Thilo Pfau, Elizabeth Shriberg, Andreas Stolcke, et al. 2003. **The icisi meeting corpus**. In *2003 IEEE International Conference on Acoustics, Speech, and Signal Processing, 2003. Proceedings.(ICASSP'03).*, volume 1, pages I–I. IEEE.
- Karen Sparck Jones. 1972. **A statistical interpretation of term specificity and its application in retrieval**. *Journal of documentation*.
- Anastassia Kornilova and Vladimir Eidelman. 2019. **BillSum: A corpus for automatic summarization of US legislation**. In *Proceedings of the 2nd Workshop on New Frontiers in Summarization*, pages 48–56, Hong Kong, China. 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, Barcelona, Spain. Association for Computational Linguistics.
- Iain McCowan, Jean Carletta, Wessel Kraaij, Simone Ashby, S Bourban, M Flynn, M Guillemot, Thomas Hain, J Kadlec, Vasilis Karaiskos, et al. 2005. **The ami meeting corpus**. In *Proceedings of the 5th International Conference on Methods and Techniques in Behavioral Research*, volume 88, page 100. Cite-seer.
- 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.
- Preksha Nema, Mitesh M. Khapra, Anirban Laha, and Balaraman Ravindran. 2017. **Diversity driven attention model for query-based abstractive summarization**. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1063–1072, Vancouver, Canada. Association for Computational Linguistics.

- Stephen Robertson and Hugo Zaragoza. 2009. The probabilistic relevance framework: Bm25 and beyond. *Foundations and Trends in Information Retrieval*, 3(4):333–389.
- Tobias Rohde, Xiaoxia Wu, and Yinhan Liu. 2021. [Hierarchical learning for generation with long source sequences](#). *ArXiv preprint*, abs/2104.07545.
- Guokan Shang, Wensi Ding, Zekun Zhang, Antoine Tixier, Polykarpos Meladianos, Michalis Vazirgiannis, and Jean-Pierre Lorré. 2018. [Unsupervised abstractive meeting summarization with multi-sentence compression and budgeted submodular maximization](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 664–674, Melbourne, Australia. Association for Computational Linguistics.
- Eva Sharma, Chen Li, and Lu Wang. 2019. [BIG-PATENT: A large-scale dataset for abstractive and coherent summarization](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2204–2213, Florence, Italy. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Manzil Zaheer, Guru Guruganesh, 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](#). *ArXiv preprint*, abs/2007.14062.
- Lulu Zhao, Weiran Xu, and Jun Guo. 2020. [Improving abstractive dialogue summarization with graph structures and topic words](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 437–449, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Ming Zhong, Da Yin, Tao Yu, Ahmad Zaidi, Mutethia Mutuma, Rahul Jha, Ahmed Hassan Awadallah, Asli Celikyilmaz, Yang Liu, Xipeng Qiu, and Dragomir Radev. 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, Online. Association for Computational Linguistics.
- Chenguang Zhu, Yang Liu, Jie Mei, and Michael Zeng. 2021. [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.
- Chenguang Zhu, Ruochen Xu, Michael Zeng, and Xuedong Huang. 2020a. [A hierarchical network for abstractive meeting summarization with cross-domain pretraining](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 194–203, Online. Association for Computational Linguistics.
- Chenguang Zhu, Ruochen Xu, Michael Zeng, and Xuedong Huang. 2020b. [A hierarchical network for abstractive meeting summarization with cross-domain pretraining](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 194–203, Online. Association for Computational Linguistics.

A Generated Cases

Tab. 5 shows some concrete sample summaries generated by BART-large model pretrained on various datasets.

	QMSum	SummScreen-FD
MediaSum	Transcribers are working on transcribing the data from the corpus . The next step is to insure that the data are clean first , and then channelized . The transcribers are also working on ensuring that the mark-up is consistent all the way throughout .	Sydney and Will are sent to a secret CIA project to find out if their father is alive or dead . Meanwhile , Sydney and Vaughn are sent to a secret CIA facility to find out what Irina is up to .
XSum	The transcribers have transcribed about thirty-five hours of transcripts from the corpus . The next step is to insure that the data is clean first , and then channelized . The transcribers are working on is to insure that the data is clean first , and then channelized . The transcribers are also incorporating additional conventions that Liz requested in terms of having a systematic handling of numbers , acronyms and acronyms which I had n't been specific about .	Sydney and Will are shocked to learn that Sydney's father , who was killed in Madagascar , is alive and working for the CIA . Will is also shocked to learn that Sydney 's mother , who was killed in the Rambaldi experiment , is alive . Will is also shocked to learn that Sydney 's father is a scientist . Will is also shocked to learn that Sydney 's mother is a scientist . Will is also shocked to learn that Sydney 's mother is a scientist
CNN	The team was working on transcribing the data , and the next step was to ensure that the data was clean first , and then channelized . The team was working on ensuring that the data was spell-checked , that the mark-up was consistent all the way throughout , and that they incorporated additional conventions that Liz requested in terms of having a systematic handling of numbers , acronyms , and acronyms which they had n't been specific about .	Sydney and Will investigate the death of her father , who was killed in a Russian KGB operation in 1982 . They discover that the Rambaldi device was a Russian spy device , which was used to test the IQ of children . Sydney 's father was a KGB agent , and she is now a KGB agent . She is also a double agent , and she is working for the CIA . She is also working for the CIA to find out who is behind the death of her father . Meanwhile , Irina is worried about her father 's death , and she is worried about her relationship with Vaughn .
Gold	Efforts by speaker fe008 are in progress to ensure that transcripts are clean (i.e . spell checked) , channelized , and conform to set conventions regarding the coding of numbers , acronyms , and explicit comments (e.g . door slams , coughs , and laughter) . Subsequent efforts by speaker fe008 will be to tighten up boundaries on the time bins . Inter-annotator agreement was reported to be very good .Speaker mn014 's multi-channel speech/non-speech segmenter is in use .	Sydney races to find a cure for Vaughn , but in order to find the antidote , Sydney must make a deal with Sark that could endanger Sloane 's life . Meanwhile , Will continues his research for Vaughn and discovers some disturbing inconsistencies involving 20-year - old standardized IQ tests . Sydney finds out that Vaughn has a girlfriend .

Table 5: Sample output summaries of various pretrained models on QMSum and SummScreen. The summary S of row X , column Y indicates that BART-large model which is pretrained on X dataset generates summary S from test set of Y . The errors and duplication are marked in red. The out-of-boundary contents are marked in grey. Tokens marked in brown indicate the keywords emerged in Gold summary.