

Improve Query Focused Abstractive Summarization by Incorporating Answer Relevance

Dan Su*, Tiezheng Yu*, 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

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

Abstract

Query focused summarization (QFS) models aim to generate summaries from source documents that can answer the given query. Most previous work on QFS only considers the query relevance criterion when producing the summary. However, studying the effect of answer relevance in the summary generating process is also important. In this paper, we propose QFS-BART, a model that incorporates the explicit answer relevance of the source documents given the query via a question answering model, to generate coherent and answer-related summaries. Furthermore, our model can take advantage of large pre-trained models which improve the summarization performance significantly. Empirical results on the Debatedpedia dataset show that the proposed model achieves the new state-of-the-art performance.¹

1 Introduction

Query focused summarization (QFS) models aim to extract essential information from a source document(s) and organize it into a summary that can answer a query (Dang, 2005). The input can be either a *single* document that has multiple views or *multiple* documents that contain multiple topics, and the output summary should be focused on the given query. QFS has various applications (e.g., a personalized search engine that provides the user with an overview summary based on their query (Su et al., 2020b)).

Early work on the QFS task mainly focused on generating extractive summaries (Davis et al., 2012; Daumé III and Marcu, 2006; Feigenblat et al., 2017; Xu and Lapata, 2020b), which may contain unreadable sentence ordering and lack cohesiveness.

* The two authors contribute equally.

¹The code is released at: <https://github.com/HLTCHKUST/QFS>

Document: Interrogator Ali Soufan said in an April op-ed article in the New York Times: “It is inaccurate to say that Abu Zubaydah had been uncooperative [and that enhanced interrogation techniques supplies interrogators with previously unobtainable information]. Along with another f.b.i. agent and with several c.i.a. officers present I questioned him from March to June before the harsh techniques were introduced later in August. Under traditional interrogation methods he provided us with important actionable intelligence.”
Query: Are traditional interrogation methods insufficient?
Summary: The same info can be obtained by traditional interrogations.

Table 1: An example of QFS. The input is a document and a corresponding query, and the **highlight** sentence is the answer from our QA module. We observe that the summary and the answers are very correlated.

Other work on abstractive QFS incorporated the query relevance into existing neural summarization models (Nema et al., 2017; Baumel et al., 2018). The closest work to ours was done by (Su et al., 2020a) and (Xu and Lapata, 2020a,b), who leveraged an external question answering (QA) module in a pipeline framework to take into consideration the answer relevance of the generated summary. However, they only used QA as distant supervision to retrieve relevant segments for generating the summary, but did not take into consideration the answer relevance in the generation model. As shown in the Table 1, the query focused summary is correlated to the answer extracted from the QA module.

On the other hand, recent neural summarization models (Paulus et al., 2017; Gehrmann et al., 2018; Zhang et al., 2020) have achieved remarkable performance in *generic* abstractive summarization by taking advantage of large pre-trained language

models (Lewis et al., 2019; Zhang et al., 2020). Yet, how to leverage these models and adapt them to the QFS task remains unexplored.

In this work, we propose QFS-BART, a BART-based (Lewis et al., 2019) framework for abstractive QFS that incorporates explicit answer relevance. We leverage a state-of-the-art QA model (Su et al., 2019) to predict the answer relevance of the given source documents to the query, then further incorporate the answer relevance into the BART-based generation model. We conduct empirical experiments on the Debaterpedia dataset, one of the first large-scale QFS datasets (Nema et al., 2017), and achieve the new state-of-the-art performance on the ROUGE metrics compared to all previously published work.

Our contributions in this work are threefold:

- Our work demonstrates the effectiveness of the answer relevance score in neural abstractive QFS.
- We propose an effective method to incorporate the answer relevance score into the pre-trained language models which can produce more query-relevant summaries.
- Our model reaches the state-of-the-art performance on a single-document QFS dataset (Debaterpedia), and brings substantial improvements over several strong baselines on two multi-document QFS datasets (DUC 2006, 2007).

2 Related Work

Abstractive summarization models aim to generate short, concise and readable text that extracts the salient information from a document. In the past few years, significant achievements (See et al., 2017; Liu and Lapata, 2019; Lewis et al., 2019; Dong et al., 2019) have been made in the *generic* abstractive summarization task which is attributed to the advanced neural architectures and the availability of large-scale datasets (Sandhaus, 2008; Hermann et al., 2015; Grusky et al., 2018).

QFS is a more complex task that aims to generate a summary according to the query and its relevant document(s). Nema et al. (2017) proposed an encode-attend-decode system with an additional query attention mechanism and diversity-based attention mechanism to generate a more query-relevant summary. Baumel et al. (2018) incorpo-

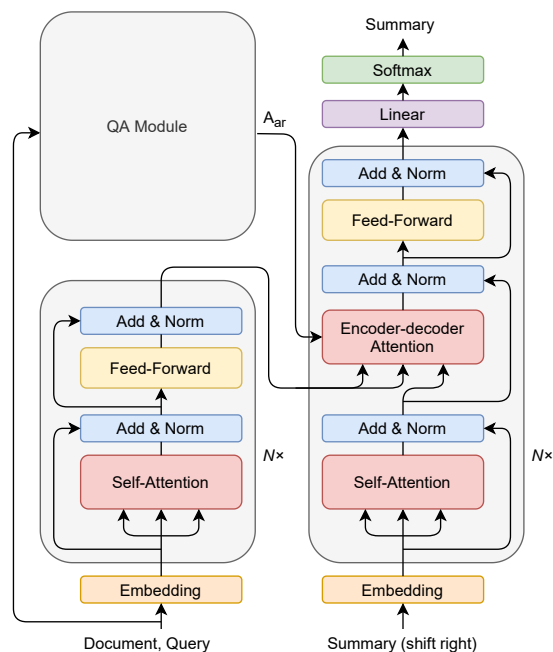


Figure 1: The framework of QFS-BART. The QA module calculates the answer relevance scores, and we incorporate the scores as explicit answer relevance attention to the encoder-decoder attention.

rated query relevance into a pre-trained abstractive summarizer to make the model aware of the query, while Xu and Lapata (2020a) discovered a new type of connection between generic summaries and QFS queries, and provided a universal representation for them which allows *generic* summarization data to be further exploited for QFS. Su et al. (2020b), meanwhile, built a query model for paragraph selection based on the answer relevance score and iteratively summarized paragraphs to a budget. Although Xu and Lapata (2020a) and Su et al. (2020b) utilized QA models for sentence- or paragraph- level answer evidence ranking, they did not make use of answer relevance to query-focused generation.

To the best of our knowledge, we are the first to leverage explicit answer relevance to abstractive QFS. In addition, our approach can be easily combined with pre-trained Transformers (Song et al., 2019; Dong et al., 2019; Lewis et al., 2019; Xiao et al., 2020), which have shown great success for the *generic* abstractive summarization task.

3 Methodology

In this section, we present our approach to incorporating the answer relevance into QFS. First, we introduce the method of generating answer relevance

scores. Then, we describe our answer relevance attention in the Transformer-based model. Third, we introduce our QFS-BART model in which the decoder is composed of a stack of answer relevance decoding layers, as shown in Figure 1.

3.1 Answer Relevance Generation

In recent years, neural models (Yang et al., 2019a; Su et al., 2019) have shown remarkable achievements in QA tasks. In order to apply QA models to the QFS task, we use HLTC-MRQA (Su et al., 2019) to generate the answer relevance score for each word in context. The reason for choosing HLTC-MRQA is twofold: 1) it shows robust generalization and transferring ability on different datasets, and 2) the model shows great performance in QA tasks and significantly outperforms the BERT-large baseline by a large margin. The HLTC-MRQA is introduced as follows.

Based on XLNet (Yang et al., 2019b), HLTC-MRQA is fine-tuned on multiple QA datasets with an additional multilayer perceptron (MLP). Given a context that contains n words, the model outputs a distribution $s \in (0, 1)$ for each word’s probability of being the start word of the answer and a probability distribution $e \in (0, 1)$ to be the end word of answer. To generate the answer relevance score r for each word, we calculate it by summing two distributions:

$$r = s + e, \quad (1)$$

where $r \in (0, 2)$.

3.2 Answer Relevance Attention

Scaled dot-product attention (Vaswani et al., 2017) is the core-component of the Transformer-based model:

$$Attention(Q, K, V) = softmax\left(\frac{QK^T}{\sqrt{d}}V\right), \quad (2)$$

where d is the dimension of the query matrix Q , key matrix K and value matrix V . The Transformer encoder is constructed by self-attention layers, where all of the keys, values and queries come from the input sequence. This makes each token in the input attend to all other tokens. The Transformer decoder layer is a combination of a self-attention layer and encoder-decoder attention layer. In the encoder-decoder attention layer, the query comes from the decoder’s self-attention layer, and the key and value come from the output of the encoder.

This allows every generated token to attend to all tokens in the input sequence.

In this work, we propose to incorporate the word-level answer relevance score as additional explicit encoder-decoder attention in the transformer decoder. Given a document with n tokens, we generate a summary with a maximum length of m tokens. Let $x^l \in \mathbb{R}^{n*d}$ denotes the output of the l -th transformer encoder layer and $y^l \in \mathbb{R}^{m*d}$ denotes the output of the l -th transformer decoder layer’s self-attention layer. The encoder-decoder attention $\alpha^l \in \mathbb{R}^{m*n}$ can be computed as:

$$\alpha^l = softmax\left(\frac{(y^l W_Q)(x^l W_K)}{\sqrt{d_k}} + A_{ar}\right), \quad (3)$$

where W_Q and $W_K \in \mathbb{R}^{d_k*d_k}$ are parameter weights and $A_{ar} \in \mathbb{R}^{m*n}$ is our explicit answer relevance score. Since the original answer relevance score is an n -dimensional vector, we repeat it m times to generate an m by n attention matrix, which means our answer relevance attention is equal to all generated tokens.

3.3 QFS-BART

Generative pre-trained models (Dong et al., 2019; Lewis et al., 2019; Raffel et al., 2019) have shown remarkable performance in natural language generation (NLG), including text summarization. We choose to combine our answer relevance attention with BART (Lewis et al., 2019), a denoising autoencoder built with a sequence-to-sequence model, for two reasons: 1) BART achieves state-of-the-art performance on several summarization datasets (i.e. CNN/DailyMail (Hermann et al., 2015) and Xsum (Narayan et al., 2018)). 2) BART follows the standard Transformer encoder-decoder architecture, and we can easily combine the answer relevance as explicit attention to the encoder-decoder attention layers. In detail, we incorporate the same answer relevance attention for all Transformer decoder layers.

Domain adaption for natural language processing tasks is widely studied (Blitzer et al., 2007; Daumé III, 2009; Liu et al., 2020; Yu et al., 2021). Hua and Wang (2017) first studied the adaptation of neural summarization models and showed that the models were able to select salient information, even when trained on out-of-domain data. Inspired by this, we leverage a two-stage fine-tuning method for our QFS-BART. In the first stage, we directly fine-tune the original BART model with the Xsum

Models	ROUGE-1	ROUGE-2	ROUGE-L
<i>Without Pre-training</i>			
Transformer (Vaswani et al., 2017)	28.16	17.48	27.28
Transformer (CONCAT)	41.72	33.62	41.25
Transformer (ADD)	41.10	33.35	40.72
SD2* (Nema et al., 2017)	41.26	18.75	40.43
CSA Transformer* (Xie et al., 2020)	46.44	37.38	45.85
<i>With Pre-training</i>			
RSA Word Count* (Baumel et al., 2018)	53.09	16.10	46.18
QR-BERTSUM-TL* (Laskar et al., 2020)	57.96	45.20	57.05
BART-FT	57.98	43.62	56.30
QFS-BART	59.02	44.59	57.44

Table 2: ROUGE-F1 scores for Debatepedia QFS dataset. Results with * mark are taken from the corresponding papers. The previous work can be divided into two categories: 1) training the models from scratch, and 2) using pre-trained models and fine-tuning on a QFS dataset.

dataset, and in the second stage, we fine-tune our QFS-BART model with QFS datasets. All the parameters in the model are initialized from the first stage. In order to make the model capture both query relevance and answer relevance, the input text is formatted in the following way:

[CLS] document [SEP] query.

The answer relevance attention score for the document is generated by the QA model, and we take the maximum number in the document as the attention score for all the words in the query.

4 Experimental Setup

Datasets We use multiple QA datasets, including SQuAD (Rajpurkar et al., 2016), NewsQA (Trischler et al., 2016), TriviaQA (Joshi et al., 2017), SearchQA (Dunn et al., 2017), HotpotQA (Yang et al., 2018) and NaturalQuestions (Kwiatkowski et al., 2019) to train HLTC-MRQA, following Su et al. (2019). We evaluate our model on the Debatepedia dataset (Nema et al., 2017) and DUC2005-7 dataset (in Appendix).

Training Details For all the experiments, we use the BART-large version to implement our models. We use a mini-batch size of 32 and train all the models on one V100 16G. During decoding, we use beam search with the beam size of 4. We decode until an end-of-sequence token is emitted and early stop when the generated summary reaches to 48 tokens.

5 Results & Analysis

We compare our proposed QFS-BART model with the following models: 1) **Transformer** does not consider the queries in the Debatepedia dataset. 2)

Document: Interrogator Ali Soufan said in an April op-ed article in the New York Times: “It is inaccurate to say that Abu Zubaydah had been uncooperative [and that enhanced interrogation techniques supplies interrogators with previously unobtainable information]. Along with another f.b.i. agent and with several c.i.a. officers present I questioned him from March to June before the harsh techniques were introduced later in August. Under traditional interrogation methods he provided us with important actionable intelligence.”

Query: Are traditional interrogation methods insufficient?

BART-FT: Al Qaeda detainee Abu Zubaydah has been cooperative under traditional interrogation.

QFS-BART: The same info can be obtained by traditional interrogation.

Gold: The same info can be obtained by traditional interrogations.

Table 3: A example taken from Debatepedia test set. The generated summary from **QFS-BART** is almost the same as the gold summary.

Transformer (CONCAT) concatenates the query and the document. 3) **Transformer (ADD)** adds the query encoded vector to the document encoder. 4) **SD2** adds a query attention model and a new diversity-based attention model to the encode-attend-decode paradigm. 5) **CSA Transformer** combines conditional self-attention (CSA) with Transformer. 6) **RAS Word Count** incorporates query relevance into a pre-trained abstractive summarization model. 7) **QR-BERTSUM-TL** presents a transfer learning technique with the Transformer-based BERTSUM model (Liu and Lapata, 2019). 8) **BART-FT** concatenates the document and query, and directly fine-tunes on the

Debatepedia dataset.

We adopt ROUGE score (Lin, 2004) as the evaluation metric. As shown in Table 2, QFS-BART significantly outperforms the models without pre-training. Compared with the models utilizing pre-training, ours improves the ROUGE-1 and ROUGE-L scores by a large margin.

5.1 Case Study

We present a case study comparing between the strong baseline BART-FT model, our QFS-BART model and the gold summary, shown in Table 3. It’s clear that the baseline model tends to copy spans from the document which are not directly related to the query and the QFS-BART model produces a more query- and answer- related summary.

6 Conclusions

In this work, we propose QFS-BART, an abstractive summarization model for query focused summarization. We use a generalizing QA model to make explicit answer relevance scores for all words in the document and combine them to the encoder-decoder attention. We also leverage pre-trained model (e.g. BART) and two-stage fine-tuning method which further improve the summarization performance significantly. Experimental results show the proposed model achieves state-of-the-art performance on Debatepedia dataset and outperforms several comparable baselines on DUC 2006-7 datasets.

Acknowledgments

We want to thank the anonymous reviewers for their constructive feedback. This work is partially funded by ITS/353/19FP of the Innovation Technology Commission, the Hong Kong SAR Government.

References

Tal Baumel, Matan Eyal, and Michael Elhadad. 2018. Query focused abstractive summarization: Incorporating query relevance, multi-document coverage, and summary length constraints into seq2seq models. *arXiv preprint arXiv:1801.07704*.

John Blitzer, Mark Dredze, and Fernando Pereira. 2007. Biographies, bollywood, boom-boxes and blenders: Domain adaptation for sentiment classification. In *Proceedings of the 45th annual meeting of the association of computational linguistics*, pages 440–447.

Hoa Trang Dang. 2005. Overview of duc 2005. In *Proceedings of the document understanding conference*, volume 2005, pages 1–12.

Hal Daumé III. 2009. Frustratingly easy domain adaptation. *arXiv preprint arXiv:0907.1815*.

Hal Daumé III and Daniel Marcu. 2006. Bayesian query-focused summarization. In *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics*, pages 305–312.

Sashka T Davis, John M Conroy, and Judith D Schlesinger. 2012. Occams—an optimal combinatorial covering algorithm for multi-document summarization. In *2012 IEEE 12th International Conference on Data Mining Workshops*, pages 454–463. IEEE.

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 *Advances in Neural Information Processing Systems*, pages 13063–13075.

Matthew Dunn, Levent Sagun, Mike Higgins, V Ugur Guney, Volkan Cirik, and Kyunghyun Cho. 2017. Searchqa: A new q&a dataset augmented with context from a search engine. *arXiv preprint arXiv:1704.05179*.

Guy Feigenblat, Haggai Roitman, Odellia Boni, and David Konopnicki. 2017. Unsupervised query-focused multi-document summarization using the cross entropy method. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 961–964.

Sebastian Gehrmann, Yuntian Deng, and Alexander M Rush. 2018. Bottom-up abstractive summarization. *arXiv preprint arXiv:1808.10792*.

Max Grusky, Mor Naaman, and Yoav Artzi. 2018. Newsroom: A dataset of 1.3 million summaries with diverse extractive strategies. *arXiv preprint arXiv:1804.11283*.

Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. *Advances in neural information processing systems*, 28:1693–1701.

Xinyu Hua and Lu Wang. 2017. A pilot study of domain adaptation effect for neural abstractive summarization. *arXiv preprint arXiv:1707.07062*.

Mandar Joshi, Eunsol Choi, Daniel S Weld, and Luke Zettlemoyer. 2017. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. *arXiv preprint arXiv:1705.03551*.

- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, et al. 2019. Natural questions: a benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:453–466.
- Md Tahmid Rahman Laskar, Enamul Hoque, and Jimmy Huang. 2020. Query focused abstractive summarization via incorporating query relevance and transfer learning with transformer models. In *Canadian Conference on Artificial Intelligence*, pages 342–348. Springer.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Yang Liu and Mirella Lapata. 2019. Text summarization with pretrained encoders. *arXiv preprint arXiv:1908.08345*.
- Zihan Liu, Yan Xu, Tiezheng Yu, Wenliang Dai, Ziwei Ji, Samuel Cahyawijaya, Andrea Madotto, and Pascale Fung. 2020. Crossner: Evaluating cross-domain named entity recognition. *arXiv preprint arXiv:2012.04373*.
- Rada Mihalcea and Paul Tarau. 2004. Texttrank: Bringing order into text. In *Proceedings of the 2004 conference on empirical methods in natural language processing*, pages 404–411.
- 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. *arXiv preprint arXiv:1808.08745*.
- Preksha Nema, Mitesh Khapra, Anirban Laha, and Balaraman Ravindran. 2017. Diversity driven attention model for query-based abstractive summarization. *arXiv preprint arXiv:1704.08300*.
- Romain Paulus, Caiming Xiong, and Richard Socher. 2017. A deep reinforced model for abstractive summarization. *arXiv preprint arXiv:1705.04304*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2019. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. Squad: 100,000+ questions for machine comprehension of text. *arXiv preprint arXiv:1606.05250*.
- Evan Sandhaus. 2008. The new york times annotated corpus. *Linguistic Data Consortium, Philadelphia*, 6(12):e26752.
- Abigail See, Peter J Liu, and Christopher D Manning. 2017. Get to the point: Summarization with pointer-generator networks. *arXiv preprint arXiv:1704.04368*.
- Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. 2019. Mass: Masked sequence to sequence pre-training for language generation. *arXiv preprint arXiv:1905.02450*.
- Dan Su, Yan Xu, Genta Indra Winata, Peng Xu, Hyeonday Kim, Zihan Liu, and Pascale Fung. 2019. Generalizing question answering system with pre-trained language model fine-tuning. In *Proceedings of the 2nd Workshop on Machine Reading for Question Answering*, pages 203–211.
- Dan Su, Yan Xu, Tiezheng Yu, Farhad Bin Siddique, Elham Barezi, and Pascale Fung. 2020a. CAiRE-COVID: A question answering and query-focused multi-document summarization system for COVID-19 scholarly information management. In *Proceedings of the 1st Workshop on NLP for COVID-19 (Part 2) at EMNLP 2020*, Online. Association for Computational Linguistics.
- Dan Su, Yan Xu, Tiezheng Yu, Farhad Bin Siddique, Elham J Barezi, and Pascale Fung. 2020b. Caire-covid: A question answering and multi-document summarization system for covid-19 research. *arXiv preprint arXiv:2005.03975*.
- Adam Trischler, Tong Wang, Xingdi Yuan, Justin Harris, Alessandro Sordani, Philip Bachman, and Kaheer Suleman. 2016. Newsqa: A machine comprehension dataset. *arXiv preprint arXiv:1611.09830*.
- 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 *Advances in neural information processing systems*, pages 5998–6008.
- Dongling Xiao, Han Zhang, Yukun Li, Yu Sun, Hao Tian, Hua Wu, and Haifeng Wang. 2020. Ernie-gen: An enhanced multi-flow pre-training and fine-tuning framework for natural language generation. *arXiv preprint arXiv:2001.11314*.
- Yujia Xie, Tianyi Zhou, Yi Mao, and Weizhu Chen. 2020. Conditional self-attention for query-based summarization. *arXiv preprint arXiv:2002.07338*.
- Yumo Xu and Mirella Lapata. 2020a. Abstractive query focused summarization with query-free resources. *arXiv preprint arXiv:2012.14774*.
- Yumo Xu and Mirella Lapata. 2020b. Coarse-to-fine query focused multi-document summarization. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3632–3645.

Yumo Xu and Mirella Lapata. 2020c. Query focused multi-document summarization with distant supervision. *arXiv preprint arXiv:2004.03027*.

Wei Yang, Yuqing Xie, Aileen Lin, Xingyu Li, Luchen Tan, Kun Xiong, Ming Li, and Jimmy Lin. 2019a. End-to-end open-domain question answering with bertserini. *arXiv preprint arXiv:1902.01718*.

Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Ruslan Salakhutdinov, and Quoc V Le. 2019b. Xlnet: Generalized autoregressive pretraining for language understanding. *arXiv preprint arXiv:1906.08237*.

Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W Cohen, Ruslan Salakhutdinov, and Christopher D Manning. 2018. Hotpotqa: A dataset for diverse, explainable multi-hop question answering. *arXiv preprint arXiv:1809.09600*.

Tiezheng Yu, Zihan Liu, and Pascale Fung. 2021. Adaptsum: Towards low-resource domain adaptation for abstractive summarization. *arXiv preprint arXiv:2103.11332*.

Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter Liu. 2020. Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. In *International Conference on Machine Learning*, pages 11328–11339. PMLR.

A Adapting QFS-BART to DUC 2005-7

DUC 2005-7 are datasets for the multi-document query focused summarization (QFS) task. As shown in the Table 4, the documents and summaries of the DUC datasets are extremely longer than those in the Debatepedia (Nema et al., 2017) dataset. We thus need to adapt the QFS-BART model to handle the multi-document scenario and produce longer output.

Datasets	Document(s)	Query	Summary
Debatepedia	66.40	11.16	9.97
DUC 2005	20058.12	26.60	243.56
DUC 2006	14330.14	23.30	246.84
DUC 2007	10759.17	21.57	243.94

Table 4: Average length of the input documents, queries and output summaries for the Debatepedia and DUC 2005-7 datasets. For the DUC datasets, we add up the lengths of all the documents related the same query

In this paper, we introduce a two-step architecture: 1) Retrieve answer-related sentences given the query, rank them by the confidence score (generated from Equation 4) and concatenate them. 2) Use our QFS-BART to produce an abstractive summary.

A.1 Answer Retrieving

We split documents into paragraphs and feed each paragraph to the QA model to get answer-related sentences. Then the sentences are ranked by the confidence score.

Document Segmentation The QA model selects one answer span given an input document, and the sentences that contain the span will be chosen as the answer-related sentences. Since we only retain the answer-related sentences as input to the next step, we set the maximum paragraph length to 300 words to avoid missing too much information in this step. Specifically, we feed text to the paragraph sentence by sentence until it reaches the maximum length.

Answer Relevance Ranking The paragraphs are fed to the QA model to generate answer-related sentences and the corresponding answer relevance scores. We align each sentence with a confidence score from the corresponding answer span. The confidence score is defined as:

$$confidence_score = P_{start} + P_{end}, \quad (4)$$

where P_{start} and P_{end} is two probability distributions over the tokens in the context. $P_{start}(i)/P_{end}(i)$ the probability of the i -th token is the start/end of the answer span in context.

A.2 Summary Generation

We use the answer-related sentences and their answer relevance scores as the input to the QFS-BART model. The DUC 2005 dataset is used as a development set to optimize the model, and we evaluate the performance on the DUC 2006-7 dataset. We compare our QFS-BART with the following models.

Models	DUC 2006			DUC 2007		
	1	2	SU4	1	2	SU4
LEAD	32.1	5.3	10.4	33.4	6.5	11.3
TEXTRANK	34.2	6.4	11.4	35.8	7.7	12.7
HLTC-MRQA	39.1	8.3	13.5	40.6	9.6	14.7
BART-CAQ*	38.3	7.1	12.9	40.5	9.2	14.4
BART-FT	38.9	8.5	13.9	40.4	10.0	15.1
QFS-BART	39.4	8.6	14.1	39.22	9.39	14.34

Table 5: ROUGE-F1 scores for DUC 2006-7 dataset. Results with * mark are taken from the corresponding papers.

LEAD (Xu and Lapata, 2020c) returns all leading sentences of the most recent document up to 250 words.

TEXTRANK (Mihalcea and Tarau, 2004) is a graph-based ranking model that incorporate two unsupervised methods for keyword and sentence extraction.

HLTC-MRQA truncates the ranked answer related sentences from our first step as the extractive summary.

BART-CQA (Su et al., 2020b) uses QA models for paragraph selection and iteratively summarizes paragraphs to 250 words.

We adopt ROUGE-F1 score (Lin, 2004) as the evaluation metric. As shown in Table 5, HLTC-MRQA significantly outperforms the LEAD and TEXTRANK baselines, which indicates the effectiveness of our answer retrieval. However, QFS-BART does not perform well on DUC 2006-7 datasets. We conjecture that the model can not converge to the task well with limited training samples (DUC 2005 contains only 300 samples).