

DialogSum Challenge: Results of the Dialogue Summarization Shared Task

Yulong Chen ^{*1}, Naihao Deng ^{*2}, Yang Liu, Yue Zhang ^{1,3}

¹ School of Engineering, Westlake University

² School of Computer Science and Engineering, University of Michigan

³ Institute of Advanced Technology, Westlake Institute for Advanced Study

yulongchen1010@gmail.com dnaihao@umich.edu

inf.yangl@outlook.com yue.zhang@wias.org.cn

Abstract

We report the results of *DialogSum Challenge*, the shared task on summarizing real-life scenario dialogues at INLG 2022. Four teams participate in this shared task and three submit their system reports, exploring different methods to improve the performance of dialogue summarization. Although there is a great improvement over the baseline models regarding automatic evaluation metrics, such as ROUGE scores, we find that there is a salient gap between model generated outputs and human annotated summaries by human evaluation from multiple aspects. These findings demonstrate the difficulty of dialogue summarization and suggest that more fine-grained evaluation metrics are in need.

1 Introduction

With the power of Pretrained Language Models (PLMs), research on text summarization has made great progresses (Liu and Lapata, 2019). However, previous research focuses on monologue summarization, such as news articles (Paulus et al., 2018; Gehrmann et al., 2018; Liu and Lapata, 2019; Liu et al., 2021), patents (Pilault et al., 2020) and academic papers (Koncel-Kedziorski et al., 2019). However, as an important communicative channel (Bender and Koller, 2020), dialogues receive less attention from the community.

To this end, we propose *DialogSum Challenge* to encourage researchers to investigate different solutions for real-life summarization (Chen et al., 2021b). Different from previous dialogue summarization tasks (Carletta et al., 2005; Gliwa et al., 2019), *DialogSum Challenge* focuses on diverse real-life scenarios such as schooling, work, medication, shopping, leisure, travel with large scale data.

The challenges of *DialogSum* can be stated from three perspectives. First, *DialogSum* include a variety of topics, requiring models to process text with different real-life scenarios. Second, compared with well-structured monologues, dialogues have unique discourse structures and language styles (Grosz et al., 1995b). The structures and use of languages differ from monologues, for instance, the key information spans over the context (Grosz et al., 1995a), which makes dialogues more difficult to encode. Third, compared with monologue summarization, dialogue summaries are written from a different perspective, usually including speakers' intents and actions (Chen et al., 2021a). Therefore, dialogue summarization is abstractive in nature and requires a high level understanding beyond text semantics (Chopra et al., 2016; Khandelwal et al., 2019). Figure 1 shows an example in *DialogSum*. Apart from the research challenges, *DialogSum Challenge* incentivizes summarization systems that can help end users. Summarizing daily spoken dialogues can help archive the important information in business and personal communication. This presumably lifts the burden of manually taking notes, liberating human beings from the tedious work.

Responding to our calls, four teams participate in the shared task, and three of them submit their system reports. The submitted systems typically employ PLMs, such as BART (Lewis et al., 2020) and PEGASUS (Zhang et al., 2020a). In addition, they explore diverse methods to improve the performance, including integrating additional features, modifying the decoding process for better summary generation, multi-task tuning the model with auxiliary tasks and using data from other sources.

To evaluate the performance, we construct a hidden test set that contains 100 manually annotated samples, and evaluate models on hidden and public test sets using both automatic and manual evaluations. For automatic evaluation, we use ROUGE scores and BERTSCORE. For manual

*Equal Contribution.

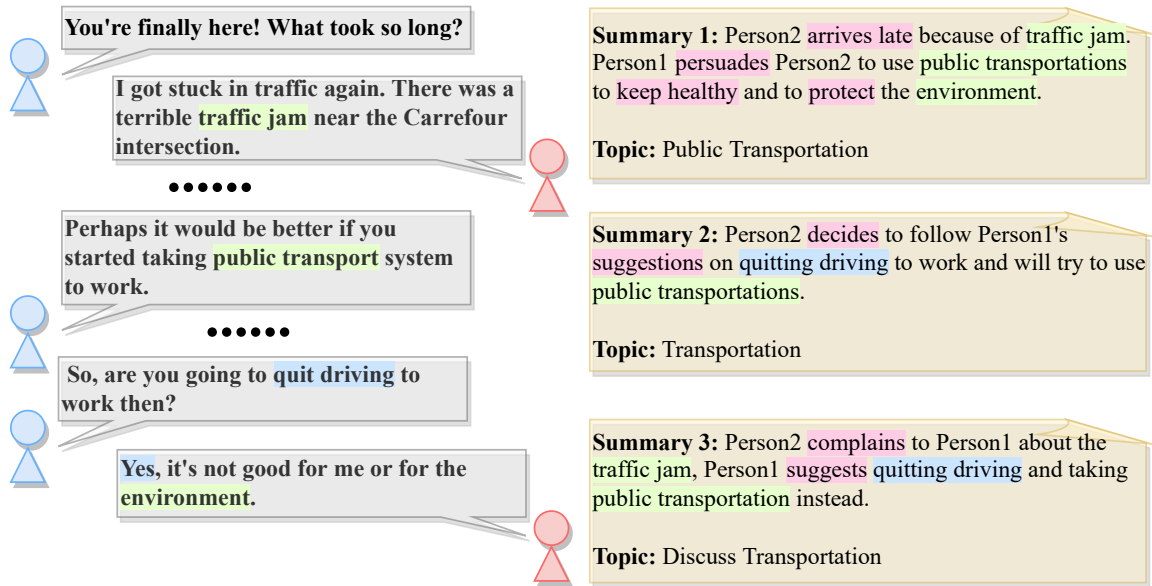


Figure 1: An example in the public test set of the *DialogSum* dataset. Tokens highlighted in red represent tokens that only appear in the summary but not in the dialogue text, requiring the model to summarize with a high level semantic understanding. Tokens highlighted in blue represent the information that spans across turns. Tokens highlighted in green show the corresponding information in the dialogue text and the summary. Note that such information scatters in various places in the dialogue.

evaluation, we follow *DialogSum Challenge* (Chen et al., 2021b) and evaluate model outputs from multiple aspects. Results show that tuning models on CNN/Daily News corpus (Hermann et al., 2015) or AMI dataset (Carletta et al., 2005), and incorporating topics in summary generation process can improve the model performance. However, there are still rooms for models to improve the metric scores, as well as the quality of the generated summaries in terms of identifying the interlocutors' intents, capturing the discourse relation, etc. Besides, we observe the mismatch between BERTSCORE and human scores, which aligns with the findings by Hanna and Bojar (2021).

Full details of the shared task description and logistics, as well as the dataset can be found at <https://cylnlp.github.io/dialogsum-challenge/>

2 Task

Given a piece of dialogue text as input, the task is to ask a model to generate a summary that conveys the key information of the dialogue. The output summary should be concise, coherent, consistent and be written from a listener's perspective.

3 Data

Data Sources We use the train, dev and public test data from *DialogSum*. Additionally, we collect 100 summaries as hidden test set from the same website where *DialogSum* crawls the data ¹. We follow the exact same procedure as the annotation for the original *DialogSum* dataset (Chen et al., 2021a). We remove the non-English characters, correct typos and grammatical errors, and filter out duplicated dialogues based on text similarity. The annotators are instructed to write the summaries for each dialogue by: (1) conveying the salient information in the dialogue and; (2) keeping the summary short and; (3) writing from the observer perspective and in formal language. Additionally, we ask annotators to keep tense consistency, preserve important discourse relations, explicitly describe emotion and speaker's intents. Also, annotators are instructed to provide a short topic for each dialogue. Table 3 shows the statistics for the data in *DialogSum Challenge*.

¹<http://tingroom.com>

¹We compute the ROUGE scores between two dialogues and filter out dialogues that have more than 80% ROUGE-1 scores.

Model	Public Test Set				Hidden Test Set			
	R1	R2	RL	BERTSCORE	R1	R2	RL	BERTSCORE
Human	53.35	26.72	50.84	92.63	-	-	-	-
GoodBai	47.61	21.66	45.48	92.72	49.66	26.03	48.55	91.69
UoT	47.29	21.65	45.92	92.26	49.75	25.15	46.50	91.76
IITP-CUNI	47.26	21.18	45.17	92.70	45.89	21.88	43.16	91.13

Table 1: Scores by automatic metrics for each submission and human results. We embolden the top scores among models, as well as the human score if it is the highest among all the scores.

Model	Public Test Set				Hidden Test Set			
	R1	R2	RL	BERTSCORE	R1	R2	RL	BERTSCORE
TCS_WITM	47.02	21.20	44.90	90.13	50.32	25.59	47.40	91.81

Table 2: Scores by automatic metrics for the submission from TCS_WITM. The model submitted by TCS_WITM predicts 3 summaries based on the 3 topics in the public test set. We take the highest score among the 3 summaries to calculate the scores.

Train	Dev	Test _{public}	Test _{hidden}	Total
12,460	500	500	100	13,560

Table 3: Number of dialogues in each split for *DialogSum Challenge*.

4 Evaluation Set-Ups

Automatic Evaluation We adopt two metrics, ROUGE (Lin, 2004) and BERTSCORE (Zhang et al., 2020b) for automatic evaluation. We use RoBERTa (Liu et al., 2019) large as the backbone to calculate BERTSCORE.

Manual Evaluation Furthermore, we conduct manual evaluation from various aspects, including standard summarization metrics Kryscinski et al. (2019, 2020), coreference information, intent identification, discourse relation following Chen et al. (2021a), as well as objectiveness (whether the summary is unsusceptible to subjectivity such as subjective assumptions in the dialogues). Besides, annotators give an overview score for the predicted summary.

5 Submissions

5.1 IITP-CUNI

The model submitted by Indian Institute of Technology Patna and Charles University employs a multi-task learning set-up to improve model performance. In their experiments, they explore several auxiliary tasks including extractive summarization to classify whether a given sentence belongs to the summary or not, novelty detection (Ghosal et al.,

2022) of whether the given summaries correspond to the same dialogue, as well as a masked language modeling (Devlin et al., 2019) task to recover summaries. They find that the BART (Lewis et al., 2020) large model tuned with the auxiliary task of the extractive summarization task with data from AMI (Carletta et al., 2005) corpus performs the best.

5.2 UoT

The participants from the University of Tübingen use the pre-trained BART model which is further tuned on CNN/Daily News corpus (Hermann et al., 2015). Besides, they penalize generating longer summaries in the decoding part of the model, and post-process the summaries to resolve generation errors (e.g. replacing speakers’ names who do not appear in the dialogue with #Person_1# or #Person_2#, and fixing duplicated labels such as #Person_1#Person_1# to #Person_1#).

They also explore methods such as intermediate task transfer learning with training on common-sense reasoning task or other summarization task first and then tune the model on *DialogSum*, transforming dialogue structures to news with structures similar to what BART model is trained on, as well as data augmentation by using data from SAM-Sum (Gliwa et al., 2019) dataset. However, they do not find any improvement using these techniques.

5.3 TCS_WITM

The model from TCS research adopts the pre-trained PEGASUS (Zhang et al., 2020a) large model which is further fine-tuned on CNN/Daily

Dialogue:

#Person1#: Excuse me, could you tell me how to get to the school clinic? I've lost my way.

#Person2#: Yes. Go straight ahead till you come to the traffic lights, turn left there and it's the first turning on the right.

#Person1#: Straight ahead to the traffic lights, left and then right.

#Person2#: That's it. It'll take you about five minutes.

#Person1#: Thank you very much.

Gold:

Summary: #Person1# is lost on the way to the school clinic. #Person2# shows #Person1# the correct direction.

IITP-CUNI:

Prediction: #Person2# shows #Person1# the way to the school clinic.

UoT:

Prediction: #Person2# shows #Person1# the way to the school clinic.

TCS_WITM:

Prediction: #Person2# tells #Person1# how to get to the school clinic.

GoodBai:

Prediction: #Person2# tells #Person1# how to get to the school clinic.

Figure 2: An example where all the predicted summaries miss the the context #Person1# is lost, while all of the gold summaries contain this context. However, all the predicted summaries successfully capture the event of #Person2# shows #Person1# the direction. ²

News. They also incorporate the topics provided in *DialogSum* dataset and feed the topic together with the dialogue text to their model. The extra information from the topics boost up the model performance compared to the baseline performance of simply feeding dialogues to the model.

6 Results

Table 1 and Table 2 show the results of human agreement and submissions from participants in *DialogSum Challenge* by the automatic metrics of ROUGE-1 (R1), ROUGE-2 (R2), ROUGE-L (RL), and BERTSCORE. We do not compare TCS_WITM with other models because TCS_WITM uses gold topic information.

In general, scores of the submissions are higher than the baseline models in the original *Dialog-*

Sum paper, demonstrating the effort from the participants. Both submissions from UoT and IITP-CUNI tune their models on other dataset then on *DialogSum*. In particular, UoT tunes their model on CNN/Daily News corpus, while IITP-CUNI tunes their model on AMI dataset. The reported results show that the model by UoT outperforms the model by IITP-CUNI. This might be attributed to the different in training size, where AMI (Carletta et al., 2005) has 137 meetings, while CNN/Daily News corpus (Hermann et al., 2015) has 312,000 articles. Thus, the model tuned on CNN/Daily News corpus might have better generalization ability. The model by TCS_WITM which adopts such a method achieves 50.32 in ROUGE-1 score for the hidden test case, showing that generating the summary with the given topic can also help the performance.

However, even the best-performed model underperforms humans by a margin larger than 5.0 in terms of all the ROUGE scores. This indicates that *DialogSum* is challenging and there is still a large room for future improvement.

Although existing works on summarization adopt BERTSCORE (Gabriel et al., 2021), we observe that the BERTSCORE deviates from the human scores. For instance, though GoodBai achieves the best BERTSCORE on the public test set, it is TCS_WITM with a lower BERTSCORE that achieves the best human scores (Overview score in Table 4). Same phenomenon also exists for the hidden test set. This observation aligns with the finding from Hanna and Bojar (2021) that BERTSCORE performance still deviates from human. Thus, the BERTSCORE is still not perfect to serve as the ultimate metric for summarization tasks, and our community might come up with a better automatic metric that aligns with human scores.

7 Human Analysis

We randomly sample 50 examples for the public test set and 20 examples for the hidden test set to conduct manual analysis. As discussed in our proposal, we include the metrics of coreference information, discourse relation, objective description, intent identification, standard summarization metrics and overview scores. We annotate -1, 0, 1 for the metrics of coreference information, discourse relation, objective description, intent identification where 1 means all correct, 0 means partially correct and -1 means all incorrect. We annotate from

Model	Public Test Set						Hidden Test Set					
	CoRef	Dis	Obj	Intent	Summ	Over	CoRef	Dis	Obj	Intent	Summ	Over
Perfect Score	1.00	1.00	1.00	1.00	5.00	5.00	1.00	1.00	1.00	1.00	5.00	5.00
GoodBai	0.96	0.86	1.00	0.72	4.12	3.96	0.90	0.90	1.00	0.70	4.20	4.15
UoT	0.98	0.92	1.00	0.80	4.18	4.08	0.75	0.75	1.00	0.80	4.00	3.70
IITP-CUNI	0.88	0.66	0.96	0.76	3.94	3.64	0.75	0.85	1.00	0.70	3.80	3.70

Table 4: Prediction results by one of the annotators of the *DialogSum* dataset. “CoRef”, “Dis”, “Obj”, “Intent”, “Summ”, “Over” indicates coreference information, discourse relation, objective description, intent identification, standard summarization metrics and overall scores, respectively. We embolden the best scores for each column.

Model	Public Test Set						Hidden Test Set					
	CoRef	Dis	Obj	Intent	Summ	Over	CoRef	Dis	Obj	Intent	Summ	Over
TCS_WITM	0.88	0.90	1.00	0.82	4.20	4.10	0.90	0.80	0.84	0.70	3.95	3.80

Table 5: Prediction results by one of the annotators of the *DialogSum* dataset for TCS_WITM.

1 to 5 for the standard summarization metrics and overview scores. The higher, the better.

Coreference Information Whether the summary aligns interlocutors and their conversation actions or contents.

Discourse Relation Whether the summary captures important relations between main events, identifying discourse relations and using appropriate phrases to express such relations.

Objective Description Whether the summary employs objective languages to describe dialogues.

Intent Identification Whether the summary captures the interlocutors’ intents.

Standard Summarization Metrics (Kryscinski et al., 2019, 2020) Whether the summary is fluent, consistent, relevant and coherent. However, in practice, we find that summaries generated by PLMs are mostly fluent, sometimes better than human annotated summary. And we have already evaluate consistent and coherent with more fine-grained metrics (Coreference Information, Objective Description, etc). Thus, we focus on relevance and judge whether the generated summary is informative and relevant.

Overview Scores Overview score of the summary with the aforementioned metrics.

Table 4 reports the scores from the aforementioned metrics. There is not a universal model which performs the best across all of these metrics, instead, each model excels at certain metrics. Overall, TCS_WITM achieves the best overview score on the public test set, while GoodBai achieves the best overview score on the hidden test set.

Dialogue:

#Person1#: What time is it, Tom?

#Person2#: Just a minute. It’s ten to nine by my watch.

#Person1#: Is it? I had no idea it was so late. I must be off now.

#Person2#: What’s the hurry?

#Person1#: I must catch the nine-thirty train.

#Person2#: You’ve plenty of time yet. The railway station is very close. It won’t take more than twenty minutes to get there.

Gold:

Summary: #Person1# is catching a train. Tom asks #Person1# not to hurry.

IITP-CUNI:

Prediction: #Person1# and Tom are in a hurry to catch the nine-thirty train.

Figure 3: An error example where the model fails to distinguish the intent between #Person1# and Tom (#Person2#).²

8 Error Analysis

Table 1, 2, 4, 5 show that the submitted models underperform human beings. Here we analyze some examples where the models make mistakes or fail to capture certain information.

Figure 3 shows an example where the model from IITP-CUNI makes a factual error and fails to reason about who is in the hurry. In order to capture the correct relationship, the model needs to reason that “Tom” is the name of #Person2#, and Tom (#Person2#) is asking #Person1# not to hurry by saying “You’ve plenty of time yet”. However,

²We only include one of the gold summaries for demonstration purpose.

the model from IITP-CUNI fails in such reasoning processes. This suggests that reasoning about information across the dialogue discourse is challenging.

Figure 2 shows an example where all of the model predictions deviate from the gold summaries. All of the 3 summaries annotated by human beings include the context of #Person1# being lost on the way.² In contrast, none of the model predictions include this context. This is plausible as the majority of the dialogue is dedicated to #Person2# showing #Person1# the correct direction, and the model might only capture such salient information in the dialogue. However, the general pattern when human beings summarize is to lay out the cause (context) of an event before telling the event, which is demonstrated in the gold summaries. Thus, there is still a large room for improvement for the model to generate human-like summaries.

Appendix A gives more examples of models making mistakes in terms of the metrics from § 7.

9 Conclusion

We host *DialogSum Challenge* of summarizing daily dialogue conversation at INLG 2021. Our dataset possesses characteristics distinguished from the existing datasets and poses new challenge to the summarization community. There are 4 teams who submit their models during the challenges. An overview of their methods is provided in this report. We evaluate their predictions by automatic metrics and human analysis. Results show that there are still rooms for models to improve the ROUGE scores, as well as the quality of the generated summaries in terms of identifying interlocutors' intents, capturing the discourse relation, etc. Besides, we observe the mismatch between BERTSCORE and human scores, which aligns with the findings by [Hanna and Bojar \(2021\)](#). Therefore, we advocate to our community to explore automatic metrics that can better align with human scores.

Acknowledgement

Yue Zhang is the corresponding author. We would like to thank Yi Zhang for her help in the human analysis as well as the collection of the hidden test set. This work receives a support from the Tencent AI Lab Rhino- Bird Focused Research Program.

References

- Emily M. Bender and Alexander Koller. 2020. [Climbing towards NLU: On meaning, form, and understanding in the age of data](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5185–5198, Online. Association for Computational Linguistics.
- Jean Carletta, Simone Ashby, Sebastien Bourban, Mike Flynn, Mael Guillemot, Thomas Hain, Jaroslav Kadlec, Vasilis Karaiskos, Wessel Kraaij, Melissa Kronenthal, et al. 2005. The ami meeting corpus: A pre-announcement. In *International workshop on machine learning for multimodal interaction*, pages 28–39. Springer.
- Yulong Chen, Yang Liu, Liang Chen, and Yue Zhang. 2021a. [DialogSum: A real-life scenario dialogue summarization dataset](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 5062–5074, Online. Association for Computational Linguistics.
- Yulong Chen, Yang Liu, and Yue Zhang. 2021b. [DialogSum challenge: Summarizing real-life scenario dialogues](#). In *Proceedings of the 14th International Conference on Natural Language Generation*, pages 308–313, Aberdeen, Scotland, UK. Association for Computational Linguistics.
- Sumit Chopra, Michael Auli, and Alexander M. Rush. 2016. [Abstractive sentence summarization with attentive recurrent neural networks](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 93–98, San Diego, California. 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.
- Saadia Gabriel, Antoine Bosselut, Jeff Da, Ari Holtzman, Jan Buys, Kyle Lo, Asli Celikyilmaz, and Yejin Choi. 2021. [Discourse understanding and factual consistency in abstractive summarization](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 435–447, Online. Association for Computational Linguistics.
- Sebastian Gehrmann, Yuntian Deng, and Alexander Rush. 2018. [Bottom-up abstractive summarization](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4098–4109, Brussels, Belgium. Association for Computational Linguistics.

- Tirthankar Ghosal, Tanik Saikh, Tameesh Biswas, Asif Ekbal, and Pushpak Bhattacharyya. 2022. Novelty detection: A perspective from natural language processing. *Computational Linguistics*, 48(1):77–117.
- 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.
- Barbara J. Grosz, Aravind K. Joshi, and Scott Weinstein. 1995a. **Centering: A framework for modeling the local coherence of discourse**. *Computational Linguistics*, 21(2):203–225.
- Barbara J Grosz, Aravind K Joshi, and Scott Weinstein. 1995b. Centering: A framework for modelling the local coherence of discourse.
- Michael Hanna and Ondřej Bojar. 2021. **A fine-grained analysis of BERTScore**. In *Proceedings of the Sixth Conference on Machine Translation*, pages 507–517, Online. Association for Computational Linguistics.
- Karl Moritz Hermann, Tomáš Kociský, 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.
- Urvashi Khandelwal, Kevin Clark, Dan Jurafsky, and Lukasz Kaiser. 2019. **Sample efficient text summarization using a single pre-trained transformer**. *ArXiv preprint*, abs/1905.08836.
- Rik Koncel-Kedziorski, Dhanush Bekal, Yi Luan, Mirella Lapata, and Hannaneh Hajishirzi. 2019. **Text Generation from Knowledge Graphs with Graph Transformers**. 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 2284–2293, Minneapolis, Minnesota. Association for Computational Linguistics.
- Wojciech Kryscinski, Nitish Shirish Keskar, Bryan McCann, Caiming Xiong, and Richard Socher. 2019. **Neural text summarization: A critical evaluation**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 540–551, Hong Kong, China. Association for Computational Linguistics.
- Wojciech Kryscinski, Bryan McCann, Caiming Xiong, and Richard Socher. 2020. **Evaluating the factual consistency of abstractive text summarization**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9332–9346, Online. Association for Computational Linguistics.
- 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.
- Yang Liu and Mirella Lapata. 2019. **Text summarization with pretrained encoders**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3730–3740, Hong Kong, China. Association for Computational Linguistics.
- Yang Liu, Sheng Shen, and Mirella Lapata. 2021. **Noisy self-knowledge distillation for text summarization**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 692–703, Online. Association for Computational Linguistics.
- 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*, abs/1907.11692.
- Romain Paulus, Caiming Xiong, and Richard Socher. 2018. **A deep reinforced model for abstractive summarization**. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Jonathan Pilault, Raymond Li, Sandeep Subramanian, and Chris Pal. 2020. **On extractive and abstractive neural document summarization with transformer language models**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9308–9319, Online. Association for Computational Linguistics.
- Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter J. Liu. 2020a. **PEGASUS: pre-training with extracted gap-sentences for abstractive summarization**. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 11328–11339. PMLR.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020b. [Bertscore: Evaluating text generation with BERT](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.

Dialogue:
#Person1#: How long does it take to get to downtown from here?
#Person2#: It is 15 minutes ' drive.
#Person1#: What companies do we have in our neighborhood?
#Person2#: Mitsubishi, HP, IBM and many other famous corporations.
#Person1#: Does the 7th floor belong to our company too?
#Person2#: Yes, our company has two floors, the 6th and 7th floor in this building.
#Person1#: It ' s such a large firm. Do we have our own staff restaurant?
#Person2#: Yes, at the end of the hall.

Gold:
Summary: #Person2# tells #Person1# information about their company and its surroundings.

UoT:
Prediction: #Person2# asks #Person2# about the distance to downtown and the companies in their neighborhood.

Figure 4: An error example for predicting the wrong coreference information.

A More Error Examples

Figure 4, 5, 6, 7, 8 show examples of predicting the wrong coreference information, wrong discourse relation, summarizing with description that is not objective, wrong intent of the interlocutors and with a low standard summarization score as well as a low overview score, respectively.

Dialogue:

#Person1#: Hi! How are things going with you?
#Person2#: I am doing well. What's up with you?
#Person1#: Believe it or not, the company I worked for closed down, so I'm out of a job.
#Person2#: The same with me. Have you given much thought to what you want to do next?
#Person1#: I am not being all that particular right now because I just need to keep a roof over my head.
#Person2#: How about interviews? Have you been on any of those yet?
#Person1#: I wish that I could get the opportunity to be interviewed.
#Person2#: How about the electrician program that they have listed over there?
#Person1#: I read about that, and the position sounded great!
#Person2#: Let's go see how we can apply for those positions.

Gold:

Summary: #Person1# and #Person2# are both unemployed. #Person2# suggests applying for the electrician program and #Person1# agrees.

IITP-CUNI:

Prediction: #Person1# tells #Person2# #Person1# is out of a job because #Person1# needs to keep a roof over #Person1#'s head. #Person2# suggests applying for electrician positions.

Figure 5: An error example for predicting the wrong discourse relation.

Dialogue:

#Person1#: John dates her seven times a week.
#Person2#: Really? That's a straws in the wind.
#Person1#: I think so. Maybe he's fallen for her.
#Person2#: Yeah. They suit each other. A perfect match between a man and a girl.
#Person1#: Right.

Gold:

Summary: #Person1# and #Person2# think that John and the girl are a perfect match.

IITP-CUNI:

Prediction: #Person1# and #Person2# talk about John and the girl he loves.

Figure 6: An error example for summarization with description that is not objective.

Dialogue:

#Person1#: Can I help you?
#Person2#: I need some stamps for this letter.
#Person1#: What kind of stamps do you want?
#Person2#: How much do I need for this letter?
#Person1#: I must weigh it first. Err... It's five grams over weigh, Do you want to send it as an ordinary or registered letter?
#Person2#: I want it registered. How much is it then?
#Person1#: Registration plus overnight... err... seven dollars in all.
#Person2#: Here's a 10 - dollar bill.
#Person1#: Now, your receipt, and the change.
#Person2#: Thanks. Good-bye.

Gold:

Summary: #Person2# wants to send a letter. #Person1# says it's five grams overweight plus overnight so seven dollars in all.

GoodBai:

Prediction: #Person1# helps #Person2# buy some stamps for a registered letter.

Figure 7: An error example for predicting the wrong intent of the interlocutors.

Gold:

Summary: #Person1# and #Person2# are talking about the terrorist attacks on 9-11, which was nightmarish for #Person1#'s family.

UoT:

Prediction: #Person1# and #Person2# talk about where they were for the terrorist attacks on 9-11. #Person1# was at home with #Person1#'s parents in New York City and #Person1# didn't see the crash itself but

Figure 8: An example with a low standard summarization score and a low overview score. The prediction generated by the model seems unfinished.