

Reinforcement Learning for Abstractive Question Summarization with Question-aware Semantic Rewards

Shweta Yadav*, Deepak Gupta*, Asma Ben Abacha, Dina Demner-Fushman

LHNCBC, U.S. National Library of Medicine, MD, USA

{shweta.shweta, deepak.gupta, asma.benabacha}@nih.gov
ddemner@mail.nih.gov

Abstract

The growth of online consumer health questions has led to the necessity for reliable and accurate question answering systems. A recent study showed that manual summarization of consumer health questions brings significant improvement in retrieving relevant answers. However, the automatic summarization of long questions is a challenging task due to the lack of training data and the complexity of the related subtasks, such as the question focus and type recognition. In this paper, we introduce a reinforcement learning-based framework for abstractive question summarization. We propose two novel rewards obtained from the downstream tasks of (i) question-type identification and (ii) question-focus recognition to regularize the question generation model. These rewards ensure the generation of semantically valid questions and encourage the inclusion of key medical entities/foci in the question summary. We evaluated our proposed method on two benchmark datasets and achieved higher performance over state-of-the-art models. The manual evaluation of the summaries reveals that the generated questions are more diverse and have fewer factual inconsistencies than the baseline summaries. The source code is available here: <https://github.com/shwetanlp/CHQ-Summ>.

1 Introduction

The growing trend in online web forums is to attract more and more consumers to use the Internet for their health information needs. An instinctive way for consumers to query for their health-related content is in the form of natural language questions. These questions are often excessively descriptive and contain more than required peripheral information. However, most of the textual content is not particularly relevant in answering the question

(Kilicoglu et al., 2013). A recent study showed that manual summarization of consumer health questions (CHQ) has significant improvement (58%) in retrieving relevant answers (Ben Abacha and Demner-Fushman, 2019). However, three major limitations impede higher success in obtaining semantically and factually correct summaries: (1) the complexity of identifying the correct question type/intent, (2) the difficulty of identifying salient medical entities and focus/topic of the question, and (3) the lack of large-scale CHQ summarization datasets. To address these limitations, this work presents a new reinforcement learning based framework for abstractive question summarization. We also propose two novel question-aware semantic reward functions: Question-type Identification Reward (QTR) and Question-focus Recognition Reward (QFR). The QTR measures correctly identified question-type(s) of the summarized question. Similarly, QFR measures correctly recognized key medical concept(s) or focus/foci of the summary.

We use the reinforce-based policy gradient approach, which maximizes the non-differentiable QTR and QFR rewards by learning the optimal policy defined by the Transformer model parameters. Our experiments show that these two rewards can significantly improve the question summarization quality, separately or jointly, achieving the new state-of-the-art performance on the MEQSUM and MATINF benchmark datasets. The main contributions of this paper are as follows:

- We propose a novel approach towards question summarization by introducing two question-aware semantic rewards (i) *Question-type Identification Reward* and (ii) *Question-focus Recognition Reward*, to enforce the generation of semantically valid and factually correct question summaries.
- The proposed models achieve the state-of-the-art performance on two question summa-

*These authors contributed equally to this work.

rization datasets over competitive pre-trained Transformer models.

- A manual evaluation of the summarized questions reveals that they achieve higher abstraction levels and are more semantically and factually similar to human-generated summaries.

2 Related Work

In recent years, reinforcement learning (RL) based models have been explored for the abstractive summarization task. Paulus et al. (2017) introduced RL in neural summarization models by optimizing the ROUGE score as a reward that led to more readable and concise summaries. Subsequently, several studies (Chen and Bansal, 2018; Pasunuru and Bansal, 2018; Zhang and Bansal, 2019; Gupta et al., 2020; Zhang et al., 2019b) have proposed methods to optimize the model losses via RL that enables the model to generate the sentences with the higher ROUGE score. While these methods are primarily supervised, Laban et al. (2020) proposed an unsupervised method that accounts for fluency, brevity, and coverage in generated summaries using multiple RL-based rewards. The majority of these works are focused on document summarization with conventional non-semantics rewards (ROUGE, BLEU). In contrast, we focus on formulating the semantic rewards that bring a high-level semantic regularization. In particular, we investigate the question’s main characteristics, i.e., question focus and type, to define the rewards.

Recently, Ben Abacha and Demner-Fushman (2019) defined the CHQ summarization task and introduced a new benchmark (MEQSUM) and a pointer-generator model. Ben Abacha et al. (2021) organized the MEDIQA-21 shared task challenge on CHQ, multi-document answers, and radiology report summarization. Most of the participating team (Yadav et al., 2021b; He et al., 2021; Sanger et al., 2021) utilized transfer learning, knowledge-based, and ensemble methods to solve the question summarization task. Yadav et al. (2021a) proposed question-aware transformer models for question summarization. Xu et al. (2020) automatically created a Chinese dataset (MATINF) for medical question answering, summarization, and classification tasks focusing on maternity and infant categories. Some of the other prominent works in the abstractive summarization of long and short documents include Cohan et al. (2018); Zhang et al. (2019a); MacAvaney et al. (2019); Sotudeh et al. (2020).

3 Proposed Method

Given a question, the goal of the task is to generate a summarized question that contains the salient information of the original question. We propose a RL-based question summarizer model over the Transformer (Vaswani et al., 2017) encoder-decoder architecture. We describe below the proposed reward functions.

3.1 Question-aware Semantic Rewards

(a) Question-type Identification Reward: Independent of the pre-training task, most language models use maximum likelihood estimation (MLE)-based training for fine-tuning the downstream tasks. MLE has two drawbacks: (1) “*exposure bias*” (Ranzato et al., 2016) when the model expects gold-standard data at each step during training but does not have such supervision when testing, and (2) “*representational collapse*” (Aghajanyan et al., 2021), is the degradation of generalizable representations of pre-trained models during the fine-tuning stage. To deal with the *exposure bias*, previous works used the ROUGE and BLEU rewards to train the generation models (Paulus et al., 2017; Ranzato et al., 2016). These evaluation metrics are based on n-grams matching and might fail to capture the semantics of the generated questions. We, therefore, propose a new question-type identification reward to capture the underlying question semantics.

We fine-tuned a BERT_{BASE} network as a question-type identification model to provide question-type labels. Specifically, we use the [CLS] token representation ($h_{[CLS]}$) from the final transformer layer of BERT_{BASE} and add the feed-forward layers on top of the $h_{[CLS]}$ to compute the final logits

$$l = \mathbf{W}(\tanh(Uh_{[CLS]} + \mathbf{a})) + \mathbf{b}$$

Finally, the question types are predicted using the *sigmoid* activation function on each output neuron of logits l . The fine-tuned network is used to compute the reward $r_{QTR}(Q^p, Q^*)$ as F-Score of question-types between the generated question summary Q^p and the gold question summary Q^* .

(b) Question-focus Recognition Reward: A good question summary should contain the key information of the original question to avoid factual inconsistency. In the literature, ROUGE-based rewards have been explored to maximize the coverage of the generated summary, but it does not guarantee to preserve the key information in the

question summary. We introduce a novel reward function called question-focus recognition reward, which captures the degree to which the key information from the original question is present in the generated summary question. Similar to QTR, we fine-tuned the BERT_{BASE} network for question-focus recognition to predict the focus/foci of the question. Specifically, given the representation matrix ($\mathbf{H} \in \mathcal{R}^{n \times d}$) of n tokens and d dimensional hidden state representation obtained from the final transformer layer of BERT_{BASE}, we performed the token level prediction using a linear layer of the feed-forward network. For each token representation (h_i), we compute the logits $l_i \in \mathcal{R}^{|C|}$, where ($|C|$) is the number of classes and predict the question focus as follows: $f_i = \text{softmax}(\mathbf{W}h_i + \mathbf{b})$. The fine-tuned network is used to compute the reward $r_{QFR}(Q^p, Q^*)$ as F-Score of question-focus between the generated question summary Q^p and the gold question summary Q^* .

3.2 Policy Gradient REINFORCE

We cast question summarization as an RL problem, where the “agent” (ProphetNet decoder) interacts with the “environment” (Question-type or focus prediction networks) to take “actions” (next word prediction) based on the learned “policy” p_θ defined by ProphetNet parameters (θ) and observe “reward” (QTR and QFR). We utilized ProphetNet (Qi et al., 2020) as the base model because it is specifically designed for sequence-to-sequence training and it has shown near state-of-the-art results on natural language generation task. We use the REINFORCE algorithm (Williams, 1992) to learn the optimal policy which maximizes the expected reward. Toward this, we minimize the loss function $\mathcal{L}_{RL} = -E_{Q^s \sim p_\theta}[r(Q^s, Q^*)]$, where Q^s is the question formed by sampling the words q_t^s from the model’s output distribution, i.e. $p(q_t^s | q_1^s, q_2^s, \dots, q_{t-1}^s, \mathcal{S})$. The derivative of \mathcal{L}_{RL} is approximated using a single sample along with baseline estimator b :

$$\nabla_\theta \mathcal{L}_{RL} = -(r(Q^s, Q^*) - b) \nabla_\theta \log p_\theta(Q^s) \quad (1)$$

The Self-critical Sequence Training (SCST) strategy (Rennie et al., 2017) is used to estimate the baseline reward by computing the reward with the question generated by the current model using the greedy decoding technique, i.e., $b = r(Q^g, Q^*)$. We compute the final reward as a weighted sum of QTR and QFR as follows:

$$r(Q^p, Q^*) = \gamma_{QTR} \times r_{QTR}(Q^p, Q^*) + \gamma_{QFR} \times r_{QFR}(Q^p, Q^*) \quad (2)$$

We train the network with the mixed loss as discussed in Paulus et al. (2017). The overall network loss is as follows:

$$\mathcal{L} = \alpha \mathcal{L}_{RL} + (1 - \alpha) \mathcal{L}_{ML} \quad (3)$$

where, α is the scaling factor and \mathcal{L}_{ML} is the negative log-likelihood loss and equivalent to $-\sum_{t=1}^{t=m} \log p(q_t^* | q_1^*, q_2^*, \dots, q_{t-1}^*, \mathcal{S})$, where \mathcal{S} is the source question.

4 Experimental Results & Analysis

4.1 Datasets

We utilized two CHQ abstractive summarization datasets: MEQSUM and MATINF¹ to evaluate the proposed framework. The MEQSUM² training set consists of 5, 155 CHQ-summary pairs and the test set includes 500 pairs. We chose 100 samples from the training set as the validation dataset.

For fine-tuning the question-type identification and question-focus recognition models, we manually labeled the MEQSUM dataset with the question type: (‘Dosage’, ‘Drugs’, ‘Diagnosis’, ‘Treatments’, ‘Duration’, ‘Testing’, ‘Symptom’, ‘Usage’, ‘Information’, ‘Causes’) and foci. We use the labeled data to train the question-type identification and question-focus recognition networks. For question-focus recognition, we follow the BIO notation and classify each token for the beginning of focus token (**B**), intermediate of focus token (**I**), and other token (**O**) classes. Since, the gold annotations for question-types and question-focus were not available for the MATINF dataset, we used the pre-trained network trained on the MEQSUM dataset to obtain the silver-standard question-types and question-focus information for MATINF³. The MATINF dataset has 5, 000 CHQ-summary pairs in the training set and 500 in the test set.

4.2 Experimental Setups

We use the pre-trained uncased version⁴ of ProphetNet as the base encoder-decoder model. We use a beam search algorithm with beam size 4 to decode the summary sentence. We train all summarization models on the respective training dataset for 20 epochs. We set the maximum question and summary sentence length to 120 and 20, respectively.

¹Since the dataset was in Chinese, we translated it to English using Google Translate.

²<https://github.com/abachaa/MeQSum>

³<https://github.com/WHUIR/MATINF>

⁴<https://huggingface.co/microsoft/prophetnet-large-uncased>

	Models	MEQSUM			MATINF*		
		R-1	R-2	R-L	R-1	R-2	R-L
Baselines	Seq2Seq (Sutskever et al., 2014)	25.28	14.39	24.64	17.77	5.10	21.48
	Seq2Seq + Attention (Bahdanau et al., 2015)	28.11	17.24	27.82	19.45	6.45	23.77
	Pointer Generator (PG) (See et al., 2017)	32.41	19.37	36.53	23.31	7.01	26.61
	SOTA (Ben Abacha and Demner-Fushman, 2019)	44.16	27.64	42.78	–	–	–
	SOTA* (Ben Abacha and Demner-Fushman, 2019)	40.00	24.13	38.56	24.58	7.30	28.08
	Transformer (Vaswani et al., 2017)	25.84	13.66	29.12	22.25	5.89	26.06
	BertSumm (Liu and Lapata, 2019)	26.24	16.20	30.59	31.16	11.94	34.70
	T ₅ _{BASE} (Raffel et al., 2019)	38.92	21.29	40.56	39.66	21.24	41.52
	PEGASUS (Zhang et al., 2019a)	39.06	20.18	42.05	40.05	23.67	43.30
	BART _{LARGE} (Lewis et al., 2019)	42.30	24.83	43.74	42.52	23.13	43.98
	MINILM (Wang et al., 2020)	43.13	26.03	46.39	35.60	18.08	38.70
	ProphetNet (Qi et al., 2020)	43.87	25.99	46.52	46.94	27.77	48.43
	ProphetNet + ROUGE-L	44.33	26.32	46.90	48.17	28.13	48.66
Joint Learning	ProphetNet + Q-type	44.40	26.63	47.05	47.19	28.02	48.70
	ProphetNet + Q-focus	44.62	26.61	47.28	47.14	28.06	48.64
	ProphetNet + Q-type + Q-focus	44.67	26.72	47.34	47.18	28.04	48.65
Proposed Approach	ProphetNet + QTR	44.60	26.69	47.38	47.51	28.40	48.94
	ProphetNet + QFR	45.36	27.33	47.96	47.53	28.29	49.11
	ProphetNet + QTR + QFR	45.52	27.54	48.19	47.73	28.54	49.33

Table 1: Comparison of the proposed models and various baselines. SOTA* denotes the method trained on the same data that we used. MATINF* denotes a translated English subset of the original Chinese MATINF dataset.

Summary Label	MEQSUM				MATINF			
	M1	M2	M3	M4	M1	M2	M3	M4
Semantics Preserved (PC/FC)	14/19.5	9.5/29	18/28	19.5/29	6/32.5	9.5/33	13.5/34	14/35
Factual Consistent (PC/FC)	11/25	7.5/35	9.5/36.5	10/38	5.5/35	7/36	7.5/41	9/42.5
Incorrect	23	11	12.5	11	10.5	11.5	11.5	10
Acceptable	18.5	10	12.5	12.5	15	10.5	8.5	9.5
Perfect	8.5	29	25	26.5	24.5	28	30	30.5

Table 2: Results of the manual evaluation of the summaries generated by ProphetNet (M1), M1+QTR (M2), M1+QFR (M3), and M1+QTR+QFR (M4). For *Semantic Preserved* and *Factual Consistent*, we report the partially correct (PC) and fully correct (FC) numbers.

We first fine-train the proposed network by minimizing only the maximum likelihood (ML) loss. Next, we initialize our proposed model with the fine-trained ML weights and train the network with the mixed-objective learning function (Eq. 3). We performed experiments on the validation dataset by varying the α , γ_{QTR} and γ_{QFR} in the range of (0, 1). The scaling factor (α) value 0.95, was found to be optimal (in terms of Rouge-L) for both the datasets. The values of $\gamma_{QTR} = 0.4$ and $\gamma_{QFR} = 0.6$ were found to be optimal on the validation sets of both datasets. To update the model parameters, we used Adam (Kingma and Ba, 2015) optimization algorithm with the learning rate of $7e - 5$ for ML training and $3e - 7$ for RL training. We obtained the optimal hyper-parameters values based on the performance of the model on the validation sets of MEQSUM and MATINF in the respective experiments. We used a cosine annealing learning rate (Loshchilov and Hutter, 2017) decay schedule, where the learning rate decreases linearly from the initial learning set in the optimizer

to 0. To avoid the gradient explosion issue, the gradient norm was clipped within 1. For all the baseline experiments, we followed the official source code of the approach and trained the model on our datasets. We implemented the approach of Ben Abacha and Demner-Fushman (2019) to evaluate the performance on both datasets. All experiments were performed on a single NVIDIA Tesla V100 GPU having GPU memory of 32GB. The average runtimes (each epoch) for the proposed approaches M_2 , M_3 and M_4 were 2.7, 2.8 and 4.5 hours, respectively. All the proposed models have 391.32 million parameters.

4.3 Results

We present the results of the proposed question-aware semantic rewards on the MEQSUM and MATINF datasets in Table-1. We evaluated the generated summaries using the ROUGE (Lin, 2004) metric⁵. The proposed model achieves new state-of-the-art performance on both datasets by

⁵<https://pypi.org/project/py-rouge/>

Original Question-I: who makes bromocriptine i am wondering what company makes the drug bromocriptine... i have on my pituitary gland ... i have to buy them...

Reference: who manufactures bromocriptine?

Generated Summary

ProphetNet: what is bromocriptine?

Proposed Approach: what company makes bromocriptine and how much does it cost?

Original Question-II: Have been on methadone for four years. I am interested in the rapid withdrawal under anesthesia, but do not have a clue where I can find a doctor or hospital who does this. I also would like to know the approximate cost and if or what insurance companies pay for this.

Reference: how can I find a physician (s) or hospital (s) who specialize in rapid methadone withdrawal under anesthesia, and the cost and insurance benefits for the procedure?

Generated Summary

ProphetNet: what is the treatment for rapid withdrawal of methadone under anesthesia?

Proposed Approach: where can i find physician (s) who specialize in rapid withdrawal of methadone?

Table 3: Correct/Incorrect summaries generated on MEQSUM. Example-I shows a perfect summary over ProphetNet. The second example shows an incorrect summary with a partially extracted focus (*‘under anesthesia’*) and two missing types (*‘cost’*, *‘procedures’*).

outperforming competitive baseline Transformer models. We also compare the proposed model with the joint learning baselines, where we regularize the question summarizer with the additional loss obtained from the question-type (Q-type) identification and question-focus (Q-focus) recognition model. To make a fair comparison with the proposed approach, we train these joint learning-based models with the same weighted strategy shown in Eq. 3. The results reported in Table 1 show the improvement over the ProphetNet on both datasets.

In comparison to the benchmark model on MEQSUM, our proposed model obtained an improvement of 9.63%. A similar improvement is also observed on the MATINF dataset. Furthermore, the results show that individual QTR and QFR rewards also improve over ProphetNet and ROUGE-based rewards. These results support two major claims: (1) question-type reward assists the model to capture the underlying question semantics, and (2) awareness of salient entities learned from the question-focus reward enables the generation of fewer incorrect summaries that are unrelated to the question topic. The proposed rewards are model-independent and can be plugged into any

pre-trained Seq2Seq model. On the downstream tasks of question-type identification and question-focus recognition, the pre-trained BERT model achieves the F-Score of 97.10% and 77.24%, respectively, on 10% of the manually labeled MEQSUM pairs.

Manual Evaluation: Two annotators, experts in medical informatics, performed an analysis of 50 summaries randomly selected from each test set. In MATINF, nine out of the 50 samples contained translation errors. We thus randomly replaced them. In both datasets, we annotated each summary with two labels *‘Semantics Preserved’* and *‘Factual Consistent’* to measure (1) whether the semantics (i.e., question intent) of the source question was preserved in the generated summary and (2) whether the key entities/foci were present in the generated summary. In the manual evaluation of the quality of the generated summaries, we categorize each summary into one of the following categories: *‘Incorrect’*, *‘Acceptable’*, and *‘Perfect’*. We report the human evaluation results (average of two annotators) on both datasets in Table-2. The results show that our proposed rewards enhance the model by capturing the underlying semantics and facts, which led to higher proportions of perfect and acceptable summaries. The error analysis identified two major causes of errors: (1) Wrong question types (e.g. the original question contained multiple question types or has insufficient type-related training instances) and (2) Wrong/partial focus (e.g. the model fails to capture the key medical entities).

5 Conclusion

In this work, we present an RL-based framework by introducing novel question-aware semantic rewards to enhance the semantics and factual consistency of the summarized questions. The automatic and human evaluations demonstrated the efficiency of these rewards when integrated with a strong encoder-decoder based ProphetNet transformer model. The proposed methods achieve state-of-the-art results on two-question summarization benchmarks. In the future, we will explore other types of semantic rewards and efficient multi-rewards optimization algorithms for RL.

Acknowledgements

This research was supported by the Intramural Research Program of the National Library of Medicine, National Institutes of Health.

Ethics / Impact Statement

Our project involves publicly available datasets of consumer health questions. It does not involve any direct interaction with any individuals or their personally identifiable data and does not meet the Federal definition for human subjects research, specifically: “a systematic investigation designed to contribute to generalizable knowledge” and “research involving interaction with the individual or obtains personally identifiable private information about an individual.”

References

- Armen Aghajanyan, Akshat Shrivastava, Anchit Gupta, Naman Goyal, Luke Zettlemoyer, and Sonal Gupta. 2021. [Better fine-tuning by reducing representational collapse](#). In *International Conference on Learning Representations*.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. [Neural machine translation by jointly learning to align and translate](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Asma Ben Abacha and Dina Demner-Fushman. 2019. [On the role of question summarization and information source restriction in consumer health question answering](#). *AMIA Summits on Translational Science Proceedings*, 2019:117.
- Asma Ben Abacha and Dina Demner-Fushman. 2019. [On the summarization of consumer health questions](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 2228–2234. Association for Computational Linguistics.
- Asma Ben Abacha, Yassine Mrabet, Yuhao Zhang, Chaitanya Shivade, Curtis Langlotz, and Dina Demner-Fushman. 2021. [Overview of the MEDIQA 2021 shared task on summarization in the medical domain](#). In *Proceedings of the 20th Workshop on Biomedical Language Processing*, pages 74–85, Online. Association for Computational Linguistics.
- Yen-Chun Chen and Mohit Bansal. 2018. Fast abstractive summarization with reinforce-selected sentence rewriting. *arXiv preprint arXiv:1805.11080*.
- 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. *arXiv preprint arXiv:1804.05685*.
- Deepak Gupta, Hardik Chauhan, Ravi Tej Akella, Asif Ekbal, and Pushpak Bhattacharyya. 2020. Reinforced multi-task approach for multi-hop question generation. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 2760–2775.
- Yifan He, Mosha Chen, and Songfang Huang. 2021. [damo_nlp at MEDIQA 2021: Knowledge-based pre-processing and coverage-oriented reranking for medical question summarization](#). In *Proceedings of the 20th Workshop on Biomedical Language Processing*, pages 112–118, Online. Association for Computational Linguistics.
- Halil Kilicoglu, Marcelo Fiszman, and Dina Demner-Fushman. 2013. Interpreting consumer health questions: The role of anaphora and ellipsis. In *Proceedings of the 2013 Workshop on Biomedical Natural Language Processing*, pages 54–62.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Philippe Laban, Andrew Hsi, John Canny, and Marti A Hearst. 2020. The summary loop: Learning to write abstractive summaries without examples. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, volume 1.
- 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: Proceedings of the ACL-04 workshop*, volume 8. Barcelona, Spain.
- 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.
- Ilya Loshchilov and Frank Hutter. 2017. [SGDR: stochastic gradient descent with warm restarts](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Sean MacAvaney, Sajad Sotudeh, Arman Cohan, Nazli Goharian, Ish Talati, and Ross W Filice. 2019. Ontology-aware clinical abstractive summarization. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1013–1016.
- Ramakanth Pasunuru and Mohit Bansal. 2018. Multi-reward reinforced summarization with saliency and entailment. *arXiv preprint arXiv:1804.06451*.

- Romain Paulus, Caiming Xiong, and Richard Socher. 2017. A deep reinforced model for abstractive summarization. *arXiv preprint arXiv:1705.04304*.
- Weizhen Qi, Yu Yan, Yeyun Gong, Dayiheng Liu, Nan Duan, Jiusheng Chen, Ruofei Zhang, and Ming Zhou. 2020. ProphetNet: Predicting future n-gram for sequence-to-SequencePre-training. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2401–2410, Online. Association for Computational Linguistics.
- 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*.
- Marc’Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. 2016. Sequence level training with recurrent neural networks. In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*.
- Steven J Rennie, Etienne Marcheret, Youssef Mroueh, Jerret Ross, and Vaibhava Goel. 2017. Self-critical sequence training for image captioning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7008–7024.
- Mario Sanger, Leon Weber, and Ulf Leser. 2021. WBI at MEDIQA 2021: Summarizing consumer health questions with generative transformers. In *Proceedings of the 20th Workshop on Biomedical Language Processing*, pages 86–95, Online. Association for Computational Linguistics.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. Get to the point: Summarization with pointer-generator networks. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1073–1083. Association for Computational Linguistics.
- Sajad Sotudeh, Nazli Goharian, and Ross W Filice. 2020. Attend to medical ontologies: Content selection for clinical abstractive summarization. *arXiv preprint arXiv:2005.00163*.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in Neural Information Processing Systems 27*, pages 3104–3112. Curran Associates, Inc.
- 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.
- Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. 2020. Minilm: Deep self-attention distillation for task-agnostic compression of pre-trained transformers. *arXiv preprint arXiv:2002.10957*.
- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256.
- Canwen Xu, Jiaxin Pei, Hongtao Wu, Yiyu Liu, and Chenliang Li. 2020. Matinf: A jointly labeled large-scale dataset for classification, question answering and summarization. *arXiv preprint arXiv:2004.12302*.
- Shweta Yadav, Deepak Gupta, Asma Ben Abacha, and Dina Demner-Fushman. 2021a. Question-aware transformer models for consumer health question summarization. *arXiv preprint arXiv:2106.00219*.
- Shweta Yadav, Mourad Sarroui, and Deepak Gupta. 2021b. NLM at MEDIQA 2021: Transfer learning-based approaches for consumer question and multi-answer summarization. In *Proceedings of the 20th Workshop on Biomedical Language Processing*, pages 291–301, Online. Association for Computational Linguistics.
- Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter J Liu. 2019a. Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. *arXiv preprint arXiv:1912.08777*.
- Shiyue Zhang and Mohit Bansal. 2019. Addressing semantic drift in question generation for semi-supervised question answering. 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 2495–2509, Hong Kong, China. Association for Computational Linguistics.
- Yuhao Zhang, Derek Merck, Emily Bao Tsai, Christopher D Manning, and Curtis P Langlotz. 2019b. Optimizing the factual correctness of a summary: A study of summarizing radiology reports. *arXiv preprint arXiv:1911.02541*.