

On Generating Extended Summaries of Long Documents

Sajad Sotudeh[†], Arman Cohan[‡], Nazli Goharian[†]

[†] IR Lab, Georgetown University, Washington D.C., USA
{sajad, nazli}@ir.cs.georgetown.edu,

[‡] Allen Institute for Artificial Intelligence, Seattle, WA, USA
armanc@allenai.org

Abstract

Prior work in document summarization has mainly focused on generating short summaries of a document. While this type of summary helps get a high-level view of a given document, it is desirable in some cases to know more detailed information about its salient points that can't fit in a short summary. This is typically the case for longer documents such as a research paper, legal document, or a book. In this paper, we present a new method for generating extended summaries of long papers. Our method exploits hierarchical structure of the documents and incorporates it into an extractive summarization model through a multi-task learning approach. We then present our results on three long summarization datasets, arXiv-Long, PubMed-Long, and Longsumm. Our method outperforms or matches the performance of strong baselines. Furthermore, we perform a comprehensive analysis over the generated results, shedding insights on future research for long-form summary generation task. Our analysis shows that our multi-tasking approach can adjust extraction probability distribution to the favor of summary-worthy sentences across diverse sections. Our datasets, and codes are publicly available at <https://github.com/Georgetown-IR-Lab/ExtendedSumm>.

Introduction

In the past few years, there has been a significant progress on both extractive (e.g., Nallapati, Zhai, and Zhou 2017; Zhou et al. 2018; Liu and Lapata 2019; Xu et al. 2020; Jia et al. 2020) and abstractive (e.g., See, Liu, and Manning 2017; Cohan et al. 2018; MacAvaney et al. 2019; Zhang et al. 2019; Sotudeh, Goharian, and Filice 2020; Dong et al. 2020) approaches for document summarization. These approaches generate a concise summary of a document, capturing its salient content. However, for a longer document containing numerous details, it is sometimes helpful to read an extended summary, providing details about its different aspects. Scientific papers are examples of such documents; while their

abstracts provide a short summary about their main methods and findings, the abstract does not include details of the methods or experimental conditions. To those who seek more detailed information about a document without having to cover the entire document, an extended or long summary can be desirable (Chandrasekaran et al. 2020; Sotudeh, Cohan, and Goharian 2020; Ghosh Roy et al. 2020).

Many long documents, including scientific papers, follow a certain hierarchical structure where content is organized throughout multiple sections and sub-sections. For example, research papers often describe objectives, problem, methodology, experiments, and conclusions (Collins, Augenstein, and Riedel 2017). Few prior studies have noted the importance of documents' structure in shorter-form summary generation (Collins, Augenstein, and Riedel 2017; Cohan et al. 2018). However, we are not aware of existing summarization methods explicitly approaching modeling the document structure when it comes to generating *extended summaries*.

We approach the problem of generating extended summary by incorporating document's hierarchical structure into the summarization model. Specifically, we hypothesize that integrating the processes of sentence selection and section prediction improves the summarization model's performance over the existing baseline models on extended summarization task. To substantiate our hypothesis, we test our proposed model on three extended summarization datasets, namely, arXiv-Long, PubMed-Long, and Longsumm. We further provide comprehensive analyses over the generated results for two long datasets, demonstrating the qualities of our model over the baseline. Our analysis reveals that the multi-tasking model helps with adjusting sentence extraction probability to the advantage of salient sentences scattered across different sections of the document. Our contributions are threefold:

1. A multi-task learning approach for leveraging document structure in generating extended summaries of long documents.
2. In-depth and comprehensive analyses over the

generated results to explore the qualities of our model in comparison with the baseline model.

3. Collecting two large-scale extended summarization datasets with oracle labels for facilitating ongoing research in extended summarization domain.

Related Work

Scientific document summarization Summarizing scientific papers has garnered vast attention from the research community during recent years, although it has been studied for decades. The characteristics of scientific papers, namely the length, writing style, and discourse structure, lead to special model considerations to overcome the summarization task in scientific domain. Researchers have utilized different approaches to address these challenges. In earlier work, Teufel and Moens (2002) proposed a Naïve Bayes classifier to do content selection over the documents’ sentences with regard to their rhetorical sentence role. More recent works have given rise to the importance of discourse structure and its usefulness in summarizing scientific papers. For example, Collins, Augenstein, and Riedel (2017) used a set of pre-defined section clusters that source sentences are appeared in as a categorical feature to aid the model at identifying summary-worthy sentences. Cohan et al. (2018) introduced large-scale datasets of arXiv and PubMed (collected from public repositories), and used a hierarchical encoder to model the discourse structure of a paper, and then used an attentive decoder to generate the summary. More recently, Xiao and Carenini (2019) proposed a sequence-to-sequence model that incorporates both the global context of the entire document and local context within the specified section. Inspired by the fact that discourse information is important when dealing with long documents (Cohan et al. 2018), we utilize this structure in scientific summarization. Unlike prior works, we integrate sentence selection and sentence section labeling processes through a multi-task learning approach. In a different line of research, the use of citation context information has been shown to be quite effective at summarizing scientific papers (Abu-Jbara and Radev 2011). For instance, Cohan and Goharian (2015, 2018) utilized a citation-based approach, denoting how the paper is cited in the reference papers, to form the summary. Here, we do not exploit any citation context information.

Extended summarization While summarization research has been extensively explored in literature, extended summarization has recently gained a huge deal of attention from the research community. Among the first attempts to encourage the ongoing research in this field, Chandrasekaran et al. (2020) set up the Longsumm shared task¹ on producing

extended summaries from scientific documents and provided a extended summarization dataset called Longsumm over which participants were urged to generate extended summaries. To tackle this challenge, researchers used different methodologies. For instance, Sotudeh, Cohan, and Goharian (2020) proposed a multi-tasking approach to jointly learn sentence importance along with its section to be included in the summary. Herein, we aim at validating the multi-tasking model on a variety of extended summarization datasets and provide a comprehensive analysis to guide future research. Moreover, Ghosh Roy et al. (2020) utilized section-contribution pre-computations (training set) to assign weights via a budget module for generating extended summaries. After specifying the section contribution, an extractive summarizer is executed over each section separately to extract salient sentences. Unlike their work, we unify sentence selection and sentence section prediction tasks to effectively aid the model at identifying summary-worthy sentences scattered around different sections. Furthermore, Reddy et al. (2020) proposed a CNN-based classification network for extracting salient sentences. Gidiotis, Stefanidis, and Tsoumakas (2020) proposed to use a divide and conquer (DANCER) approach (Gidiotis and Tsoumakas 2020) to identify the key sections of the paper to be summarized. The PEGASUS abstractive summarizer (Zhang et al. 2019) then runs over each section separately to produce section summaries, which are finally concatenated to form the extended summary. Beltagy, Peters, and Cohan (2020) proposed “Longformer” that utilizes “Dilated Sliding Windows”, enabling the model to achieve better long-range coverage on long documents. With all being mentioned above, to the best of our knowledge, we are the first to conduct quite a comprehensive analysis over the generated summarization results in the extended summarization domain.

Dataset

We use three extended summarization datasets in this research. The first one is Longsumm dataset, which has been provided in the Longsumm 2020 shared task (Chandrasekaran et al. 2020). To further validate the model, we collect two additional datasets called arXiv-Long and PubMed-Long by filtering the instances of arXiv and PubMed corpora to retain those whose abstract contains at least 350 tokens. Also, to measure how our model works on the mixed varied-length scientific dataset, we exploit the arXiv summarization dataset (Cohan et al. 2018).

Longsumm The Longsumm dataset was provided for the Longsumm challenge (Chandrasekaran et al. 2020) whose aim was to generate extended summaries for scientific papers. It consists of two types of summaries:

¹<https://ornlcda.github.io/SDProc/sharedtasks.html>

- **Extractive summaries:** these summaries are coming from the TalkSumm dataset (Lev et al. 2019), containing 1705 extractive summaries of scientific papers according to their video talks in conferences (i.e., ACL, NAACL, etc.). Each summary within this corpus is formed by appending top 30 sentences of the paper.
- **Abstractive summaries:** an add-on dataset containing 531 abstractive summaries from several CS domains such as Machine Learning, NLP, and AI, that are written by NLP and ML researchers on their blogs. The length of summaries in this dataset ranges from 50-1500 words per paper.

In our experiments, we use the extractive set along with 50% of the abstractive set as our training set, containing 1969 papers; and 20% of it as the validation set. Note that these splits are made for the purpose of our internal experiments as the official test set containing 22 abstractive summaries is blind (Chandrasekaran et al. 2020).

arXiv-Long & PubMed-Long. To further test our methods on additional datasets, we construct two extended summarization datasets for our task. For creating the first dataset, we take arXiv summarization dataset introduced by Cohan et al. (2018) and filter the instances whose abstract (i.e., ground-truth summary) contains at least 350 tokens. We call this dataset arXiv-Long. We repeat the same process on the PubMed papers obtained from the Open Access FTP service² and call this dataset PubMed-Long. The motivation is that we are interested in validating our model on extended summarization datasets to investigate its effects compared to the existing works, and 350 is the length threshold that we use to characterize papers with “long” summaries. The resulting sets contain 11,149 instances for arXiv-Long, and 88,035 instances for PubMed-Long datasets. Note that the abstract of papers are used as ground-truth summaries in these two datasets. The overall statistics of the datasets are shown in Table 1. We release these datasets to facilitate future research in extended summarization.³

Methodology

In this section, we discuss our proposed method that aims at jointly learning to predict sentence importance and its corresponding section. Before discussing the details of our summarization model, we investigate the preliminary background that provides a fair basis for implementing our method.

Background

Extractive Summarization The extractive summarization system aims at extracting salient

²<https://www.ncbi.nlm.nih.gov/pmc/tools/ftp>

³<https://github.com/Georgetown-IR-Lab/ExtendedSumm>

Datasets	# docs	avg. doc. length (tokens)	avg. summ. length (tokens)
arXiv	215K	4938	220
Longsumm	2.2K	5858	920
arXiv-Long	11.1K	9221	574
PubMed-Long	88.0K	5359	403

Table 1: Statistics on arXiv (Cohan et al. 2018), Longsumm (Chandrasekaran et al. 2020), and two extended summarization datasets (arXiv-Long, PubMed-Long), collected by this work.

sentences to be included in the summary. Formally, let P show a scientific paper containing sentences $[s_1, s_2, s_3, \dots, s_m]$, where m is the number of sentences. The extractive summarization is then defined as the task of assigning a binary label ($\hat{y}_i \in \{0, 1\}$) to each sentence s_i within the paper, signifying whether the sentence should be included in the summary.

BERTSUM: BERT for Summarization As our base model we use the BERTSUM extractive summarization model (Liu and Lapata 2019), a BERT-based sentence classification model fine-tuned for summarization.

After BERTSUM outputs sentence representations within the input document, several inter-sentence Transformer layers are stacked upon the BERTSUM to collect document-level features. The final output layer is a linear classifier with Sigmoid activation function to decide whether the sentence should be included or not. The loss function is defined as below:

$$\mathcal{L}_1 = -\frac{1}{N} \sum_{i=1}^n y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i) \quad (1)$$

where N is the output size, \hat{y}_i is the output of the model, and y_i is the corresponding target value. In our experiments, we use this model to extract salient sentences (i.e., those with the positive label) to form the summary. We set this model as the baseline called BERTSUMEXT (Liu and Lapata 2019).

Our model: a section-aware summarizer

Inspired by few prior works that have studied the effect of document’s hierarchical structure in summarization task (Conroy and Davis 2017; Cohan et al. 2018), we define a section prediction task, aiming at predicting the relevant section for each sentence in the document. Specifically, we add an additional linear classification layer on top of BERTSUM sentence representations to predict the relevant section to each sentence. The loss function for the section prediction network is defined as follows:

$$\mathcal{L}_2 = -\sum_{i=1}^S y_i \log(\hat{y}_i) \quad (2)$$

where y_i and \hat{y}_i are the ground-truth and the model scores for each section i in S .

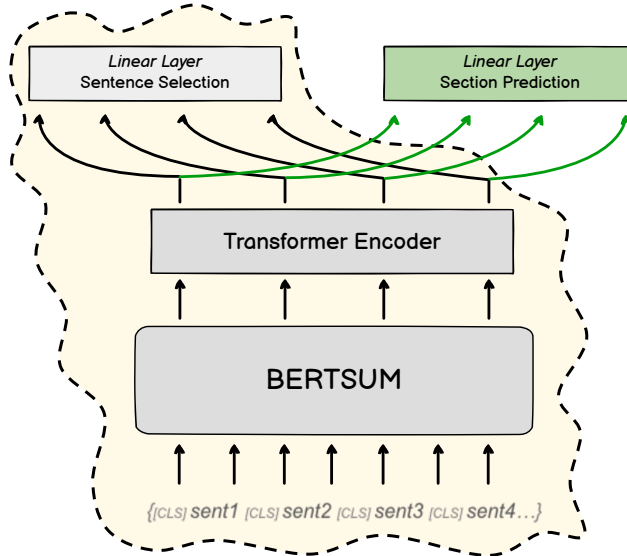


Figure 1: The overview of BERTSUMEXTMULTI model. The baseline model (i.e., BERTSUMEXT) is dash-boarded. The extension to the baseline model is addition of Section Prediction linear layer (specified in green box).

The entire extractive network is then trained to optimize both tasks (i.e., sentence selection and section prediction) in a multi-task setting:

$$\mathcal{L}_{\text{Multi}} = \alpha \mathcal{L}_1 + (1 - \alpha) \mathcal{L}_2 \quad (3)$$

where \mathcal{L}_1 is the binary cross-entropy loss from sentence selection task, \mathcal{L}_2 is the categorical cross-entropy loss from section prediction network, and α is the weighting parameter that balances the learning procedure between the sentence and section prediction tasks.

Experimental Setup

In this section, we give details about the pre-processing steps on the datasets and parameters that we used for the experimented models.

For our baseline, we used the pre-trained BERTSUM model and implementation provided by the authors (Liu and Lapata 2019).⁴ The BERTSUMEXTMULTI is that of the model used in (Sotudeh, Cohan, and Goharian 2020), but without post-processing module at inference time, which utilizes trigram-blocking (Liu and Lapata 2019) to hinder repetitions in the final summary. We intentionally removed the post-processing part as the model could attain higher scores in the absence of this module throughout our experiments. In order to obtain ground-truth section labels associated with each sentence, we utilized the external sequential-sentence package⁵ by Cohan et al. (2019). To provide oracle labels for source sentences in our datasets, we use a greedy

labelling approach (Liu and Lapata 2019) with slight modification for labelling up top 30, 15, and 25 sentences for Longsum, arXiv-Long, and PubMed-Long datasets, respectively, since these numbers of oracle sentences yielded the highest oracle scores.⁶ For the joint model, we tuned α (loss weighting parameter) at 0.5 as it resulted in the highest scores throughout our experiments. In all our experiments, we pick the checkpoint that achieves the best average of ROUGE-2 and ROUGE-L scores on the validation intervals as our best model for inference.

Results

In this section, we present the performance of the baseline and our model over the validation and test sets of the extended summarization datasets. We then discuss our proposed model’s performance compared to baseline over a mix of varied-length summarization dataset (i.e., arXiv). As the evaluation metrics, we report the summarization systems’ performance in terms of ROUGE-1 (F1), ROUGE-2 (F1), and ROUGE-L (F1) metrics.

As we see in Table 2, we notice that having section predictor model incorporated into summarization model (i.e., BERTSUMEXTMULTI model) performs fairly well compared to the baseline model. This is a particularly important finding since it characterizes the importance of injecting documents’ structure when summarizing a scientific paper. While the score gap is relatively higher in arXiv-Long and Longsum datasets, it is similar in PubMed-Long dataset.

As observed in Table. 3, it is noticeable that

⁴<https://github.com/nlpyang/PreSumm>

⁵https://github.com/allenai/sequential_sentence_classification

⁶The modification was made to assure that the oracle sentences are sampled from diverse sections.

Model	Dataset	Validation			Test		
		RG-1(%)	RG-2(%)	RG-L(%)	RG-1(%)	RG-2(%)	RG-L(%)
BERTSUMEXT	Longsumm	43.2	12.4	16.8	–	–	–
BERTSUMEXTMULTI		43.3	13.0*	17.0	53.1	16.8	20.3
BERTSUMEXT	arXiv-Long	47.1	18.2	20.8	47.2	18.4	21.1
BERTSUMEXTMULTI		47.8*	18.9*	21.3*	47.8*	19.2*	21.5*
BERTSUMEXT	PubMed-Long	49.1	24.3	25.7	49.1	24.5	25.8
BERTSUMEXTMULTI		48.9	24.1	25.5	48.9	24.1	25.5

Table 2: ROUGE (F1) results of the baseline (i.e., BERTSUMEXT) and our proposed model (i.e., BERTSUMEXTMULTI) on extended summarization datasets. * shows the statistically significant improvement (paired t-test, $p < 0.01$). The validation set for Longsumm refers to our internal validation set (20% of the abstractive set) as there was no official validation set provided for this dataset.

	RG-1	RG-2	RG-L	F-Measure average
<i>Other systems</i>				
Summaformers (Ghosh Roy et al. 2020)	49.38	16.86	21.38	29.21
Wing	50.58	16.62	20.50	29.23
IITBH-IITP (Reddy et al. 2020)	49.03	15.74	20.46	28.41
Auth-Team (Gidiotis, Stefanidis, and Tsoumakas 2020)	50.11	15.37	19.59	28.36
CIST_BUPT (Li et al. 2020)	48.99	15.06	20.13	28.06
<i>This work</i>				
BERTSUMEXTMULTI	53.11	16.77	20.34	30.07

Table 3: ROUGE (F1) results of our multi-tasking model on the blind test set of Longsumm shared task containing 22 abstractive summaries (Chandrasekaran et al. 2020), along with the performance of other participants’ systems. We only show top 5 participants in this table.

BERTSUMEXTMULTI approach performs top among the state-of-the-art long summarization methods on the blind test set of LongSum challenge (Chandrasekaran et al. 2020). While this model improves ROUGE-1 quite significantly over the other state-of-the-art, it stays competitive on ROUGE-2 and ROUGE-L metrics. In terms of ROUGE (F1) F-Measure average, BERTSUMEXTMULTI model ranks first by a huge margin compared to the other systems.

To test the model on mixed varied-length summarization datasets, we trained and tested it on arXiv (Cohan et al. 2018) dataset, which contains a mix of varying length abstracts as ground-truth summaries. Table 4 shows that our model can achieve competitive performance on this dataset. While the model does not yield any improvement on arXiv dataset, our hypothesis was to investigate if our model is superior to existing models on longer-form datasets –such as those we have used in this research, which we validated by presenting the evaluation results on long summarization datasets.

Analysis

In order to gain insights into how our multi-tasking approach works on different long datasets, we perform an extensive analysis in this section to explore the qualities of our multi-tasking system (i.e., BERTSUMEXTMULTI) over the baseline (i.e., BERTSUMEXT). Specifically, we perform two types of analyses: 1) quantitative analysis; 2) qualitative analysis.

For the first part, we choose to use two metrics: RG_{diff} which denotes the average ROUGE (F1) difference (i.e., gap) between the baseline and our model ⁷. Positive values indicate the improvement, while negative values denote the decline in scores. Similarly, F_{diff} is the average difference of F1 score between the baseline and our model. We create three bins sorted by RG_{diff} : IMPROVED which contains the reports whose average ROUGE (F1) score is improved by the multi-tasking model; TIED including those that the multi-tasking model leaves unchanged in terms of modifying average ROUGE (F1) score; and DECLINED containing those whose average ROUGE (F1) score has decreased by the joint model.

For the qualitative analysis section, we specifically

⁷The average is defined on ROUGE-1 (F1), ROUGE-2

Model	Dataset	Validation			Test		
		RG-1(%)	RG-2(%)	RG-L(%)	RG-1(%)	RG-2(%)	RG-L(%)
BERTSUMEXT	arXiv	43.6	16.6	20.2	44.0	16.8	20.4
BERTSUMEXTMULTI		43.4	16.5	19.8	43.5	16.5	20.0

Table 4: ROUGE (F1) results of the baseline (i.e., BERTSUMEXT) and our proposed model (i.e., BERTSUMEXTMULTI) on arXiv summarization dataset.

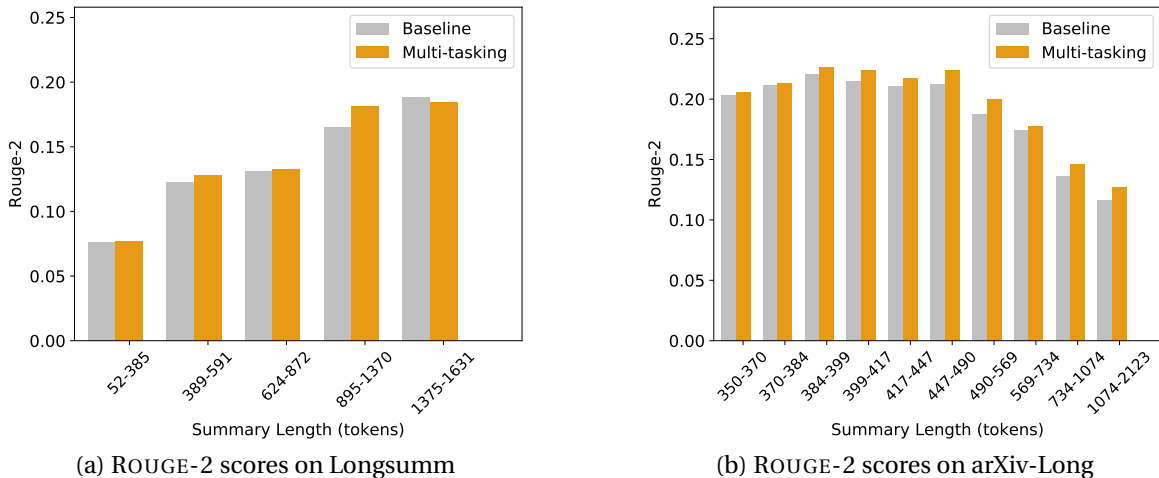


Figure 2: Bar charts exhibiting the correlation of ground-truth summary length (in tokens) with the baseline (i.e., BERTSUMEXT) and our multi-tasking model’s (i.e., BERTSUMEXTMULTI) performance. The diagrams are shown for Longsum and arXiv-Long datasets’ test set. Each bin contains 31 summaries for Longsum, and 196 summaries for arXiv-Long. As denoted, the multi-tasking model generally outperforms the baseline on later bins which include longer-form summaries.

aim at comparing the methods in terms of section distribution since that is where our method’s improvements are expected to come from. Furthermore, we conduct an additional length analysis over the results generated by the baseline versus our model.

Quantitative Analysis

We first perform the quantitative analysis over the long summarization datasets’ test sets in two parts including 1) *Metric analysis* which aims at comparing different bins based on the average ROUGE score difference of the baseline and our model; 2) *Length analysis* that targets at finding the correlation between the summary length on different bins and models’ performance.

Metric analysis Table 5 shows the overall quantities of Longsum and arXiv-Long datasets in terms of average difference of ROUGE and F1 scores. As shown, the multi-tasking approach is able to improve 76 summaries with an average ROUGE (F1) improvement of 2.05%. This is even more when it

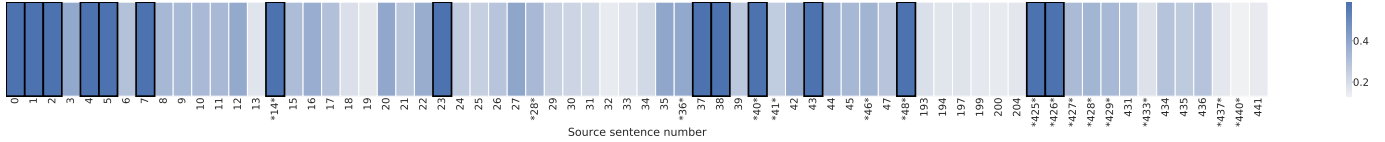
(F1), and ROUGE-L (F1) scores.

comes to evaluating the model on arXiv-Long dataset with average ROUGE improvement of 2.40%.

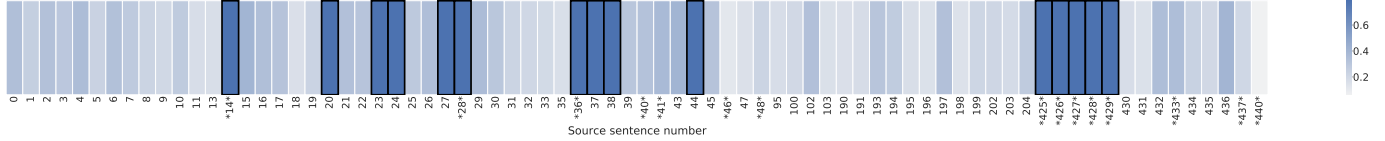
Interestingly, our method can consistently improve F1 measure in general (See total F1 scores in Table. 5). Seemingly, F1 metric directly correlates with ROUGE (F1) metric on arXiv-Long dataset, whereas this is not the case on DECLINED bin of the Longsum dataset. This might be due to the relatively small test set size of Longsum dataset. It has to be mentioned that IMPROVED bin holds relatively higher counts and improved metrics than that of DECLINED bin across both datasets in our evaluation.

Length analysis We analyze the generated results by both models to see if the summary length affects the models’ performance using bar charts in Figure 2. The bar charts are intended to provide the basis for comparing both models on different length bins (x-axis), which are evenly-spaced (i.e., having the same number of papers). It has to be mentioned that we used five bins (each bin with 31 summaries) and ten bins (each bin with 196 summaries) for Longsum and arXiv-Long datasets, respectively.

As shown in Figure 2 (a), for Longsum dataset, as the length of the ground-truth summary increases,



(a) Extraction probability distribution of the baseline model (i.e., BERTSUMEXT) over the source sentences.



(b) Extraction probability distribution of the multi-tasking model (i.e., BERTSUMEXTMULTI) over the source sentences.

Figure 3: Heat-maps showing the extraction probabilities over the source sentences (Paper ID: astro-ph9807040 sampled from arXiv-Long dataset). For simplicity, we have only shown the sentences that gain over 15% extraction probability by the models. The cells bordered in black show the models’ final selection, and oracle sentences are indicated with *.

Bin	Dataset	Count	RG _{diff}	F _{diff}
IMPROVED	Longsum	76	2.05	6.16
TIED		4	0	0
DECLINED		74	-1.47	1.95
Total		154	0.31	4.11
IMPROVED	arXiv-Long	1,084	2.40	4.47
TIED		67	0	0.32
DECLINED		801	-1.82	-1.34
Total		1952	0.59	1.94

Table 5: IMPROVED, TIED, and DECLINED bins on the test set of Longsum and arXiv-Long datasets. The numbers show the improvements (positive) and drops (negative) compared to the baseline model (i.e., BERTSUMEXT).

the multi-tasking model generally improves over the baseline consistently on both datasets, except for the last bin on Longsum dataset where it achieves comparable performance. This behaviour is also observed on ROUGE-1 and ROUGE-L for Longsum dataset. The ROUGE improvement is even more noticeable when it comes to analysis over arXiv-Long dataset (See Figure 2 (b)). Thus, the length analysis supports our hypothesis that the multi-tasking model outperforms the baseline more significantly when the summary is of longer-form.

Qualitative analysis

As the results of the qualitative analysis on the IMPROVED bin is observed, we found out that the multi-tasking model can effectively sample sentences from diverse sections when the ground-truth summary is also sampled from diverse sections. It improves significantly over the baseline when the extractive model can detect salient sentences from

important sections.

By investigating the summaries from DECLINED bin, we noticed that in declined summaries, while our multi-tasking approach can adjust extraction probability distribution to diverse sections, it has difficulty picking up salient sentences (i.e., positive sentences) from the corresponding section; thus, it leads to relatively lower ROUGE score. This might be improved if two networks (i.e., sentence selection and section prediction) are optimized in a more elegant way such that the extractive summarizer can further select salient sentences from the specified sections when they could be identified. For example, the improved multi-tasking methods can involve task prioritization (Guo et al. 2018) to dynamically balance the learning process between two tasks during training, rather than using a fixed α parameter.

In the cases where the F1 score and ROUGE (F1) were not consistent with each other, we observed that adding non-salience sentences to the final summary hurts the final ROUGE (F1) scores. In other words, while the multi-tasking approach can achieve a higher F1 score compared to the baseline since it chooses different non-salient (i.e., negative) sentences than baseline, the overall ROUGE (F1) scores drop slightly. Having conditional decoding length (i.e., sentences) might help with this as done in (Mao et al. 2020).

Fig. 3 shows the extraction probabilities that each model outputs on the source sentences. It is observable that the baseline model picks most of the sentences (47%) from the beginning of the paper, while the multi-tasking approach (b) can effectively distract probability distribution to summary-worthy sentences that are all around different sections of the paper, and pick those with higher confidence. Our model achieves the overall F1 score of 53.33% on this sample paper, while the baseline’s F1 score is 33.33%.

Conclusion & Future Work

In this paper, we approach the problem of generating extended summaries, given a long document. Our proposed model is a multi-task learning approach that unifies sentence selection and section prediction processes, extracting summary-worthy sentences. We further collect two large-scale extended summary datasets (arXiv-Long and PubMed-Long) from scientific papers. Our results on three datasets show the efficacy of the joint multi-task model in the extended summarization task. While it achieves fairly competitive performance with the baseline on one of three datasets, it consistently improves over the baseline in the other two evaluation datasets. We further performed extensive quantitative and qualitative analyses over the generated results by both models. These evaluations revealed our model's qualities compared to the baseline. Based on the error analysis, it could be noticed that the performance of this model highly depends on the multi-tasking objectives. Future studies could fruitfully explore this issue further by optimizing the multi-task objectives in a way that both sentence selection and section prediction tasks can benefit.

References

- Abu-Jbara, A.; and Radev, D. R. 2011. Coherent Citation-Based Summarization of Scientific Papers. In *ACL*.
- Beltagy, I.; Peters, M. E.; and Cohan, A. 2020. Longformer: The Long-Document Transformer. *ArXiv* abs/2004.05150.
- Chandrasekaran, M. K.; Feigenblat, G.; Hovy, E.; Ravichander, A.; Shmueli-Scheuer, M.; and de Waard, A. 2020. Overview and Insights from the Shared Tasks at Scholarly Document Processing 2020: CL-SciSumm, LaySumm and LongSumm. In *SDP*.
- Cohan, A.; Beltagy, I.; King, D.; Dalvi, B.; and Weld, D. S. 2019. Pretrained Language Models for Sequential Sentence Classification. In *EMNLP/IJCNLP*.
- Cohan, A.; Dernoncourt, F.; Kim, D. S.; Bui, T.; Kim, S.; Chang, W.; and Goharian, N. 2018. A Discourse-Aware Attention Model for Abstractive Summarization of Long Documents. In *NAACL-HLT*.
- Cohan, A.; and Goharian, N. 2015. Scientific Article Summarization Using Citation-Context and Article's Discourse Structure. In *EMNLP*.
- Cohan, A.; and Goharian, N. 2018. Scientific document summarization via citation contextualization and scientific discourse. *International Journal on Digital Libraries* 19(2-3): 287–303.
- Collins, E.; Augenstein, I.; and Riedel, S. 2017. A Supervised Approach to Extractive Summarisation of Scientific Papers. In *CoNLL*.
- Conroy, J. M.; and Davis, S. 2017. Section mixture models for scientific document summarization. *International Journal on Digital Libraries* 19: 305–322.
- Dong, Y.; Wang, S.; Gan, Z.; Cheng, Y.; Cheung, J.; and jing Liu, J. 2020. Multi-Fact Correction in Abstractive Text Summarization. In *EMNLP*, volume abs/2010.02443.
- Ghosh Roy, S.; Pinnaparaju, N.; Jain, R.; Gupta, M.; and Varma, V. 2020. Summaformers @ LaySumm 20, LongSumm 20. In *SDP*.
- Gidiotis, A.; Stefanidis, S.; and Tsoumakas, G. 2020. AUTH @ CLSciSumm 20, LaySumm 20, LongSumm 20. In *SDP*.
- Gidiotis, A.; and Tsoumakas, G. 2020. A Divide-and-Conquer Approach to the Summarization of Academic Articles. *ArXiv* abs/2004.06190.
- Guo, M.; Haque, A.; Huang, D.-A.; Yeung, S.; and Fei-Fei, L. 2018. Dynamic Task Prioritization for Multitask Learning. In *ECCV*.
- Jia, R.; Cao, Y.; Tang, H.; Fang, F.; Cao, C.; and Wang, S. 2020. Neural Extractive Summarization with Hierarchical Attentive Heterogeneous Graph Network. In *EMNLP*.
- Lev, G.; Shmueli-Scheuer, M.; Herzig, J.; Jerbi, A.; and Konopnicki, D. 2019. TalkSumm: A Dataset and Scalable Annotation Method for Scientific Paper Summarization Based on Conference Talks. *ACL*.
- Li, L.; Xie, Y.; Liu, W.; Liu, Y.; Jiang, Y.; Qi, S.; and Li, X. 2020. CIST@CL-SciSumm 2020, LongSumm 2020: Automatic Scientific Document Summarization. In *SDP*.
- Liu, Y.; and Lapata, M. 2019. Text Summarization with Pretrained Encoders. In *EMNLP/IJCNLP*.
- MacAvaney, S.; Sotudeh, S.; Cohan, A.; Goharian, N.; Talati, I.; and Filice, R. W. 2019. Ontology-Aware Clinical Abstractive Summarization. In *SIGIR*.
- Mao, Y.; Qu, Y.; Xie, Y.; Ren, X.; and Han, J. 2020. Multi-document Summarization with Maximal Marginal Relevance-guided Reinforcement Learning. In *EMNLP*.
- Nallapati, R.; Zhai, F.; and Zhou, B. 2017. SummaRuNNer: A Recurrent Neural Network Based Sequence Model for Extractive Summarization of Documents. In *AAAI*.
- Reddy, S.; Saini, N.; Saha, S.; and Bhattacharyya, P. 2020. IITBH-IITP@CL-SciSumm20, CL-LaySumm20, LongSumm20. In *SDP*.
- See, A.; Liu, P. J.; and Manning, C. D. 2017. Get To The Point: Summarization with Pointer-Generator Networks. In *ACL*.
- Sotudeh, S.; Cohan, A.; and Goharian, N. 2020. GUIR @ LongSumm 2020: Learning to Generate Long Summaries from Scientific Documents. In *SDP*.

Sotudeh, S.; Goharian, N.; and Filice, R. 2020. Attend to Medical Ontologies: Content Selection for Clinical Abstractive Summarization. In *ACL*.

Teufel, S.; and Moens, M. 2002. Summarizing Scientific Articles: Experiments with Relevance and Rhetorical Status. *Computational Linguistics* 28: 409–445.

Xiao, W.; and Carenini, G. 2019. Extractive Summarization of Long Documents by Combining Global and Local Context. In *EMNLP/IJCNLP*.

Xu, J.; Gan, Z.; Cheng, Y.; and jing Liu, J. 2020. Discourse-Aware Neural Extractive Text Summarization. In *ACL*.

Zhang, J.; Zhao, Y.; Saleh, M.; and Liu, P. J. 2019. PEGASUS: Pre-training with Extracted Gap-sentences for Abstractive Summarization. In *ICML*.

Zhou, Q.; Yang, N.; Wei, F.; Huang, S.; Zhou, M.; and Zhao, T. 2018. Neural Document Summarization by Jointly Learning to Score and Select Sentences. In *ACL*.