

Keyphrase Generation Beyond the Boundaries of Title and Abstract

Krishna Garg Jishnu Ray Chowdhury Cornelia Caragea

Computer Science

University of Illinois Chicago

{kgarg8, jraych2, cornelia}@uic.edu

Abstract

Keyphrase generation aims at generating important phrases (keyphrases) that best describe a given document. In scholarly domains, current approaches have largely used only the title and abstract of the articles to generate keyphrases. In this paper, we comprehensively explore whether the integration of additional information from the full text of a given article or from semantically similar articles can be helpful for a neural keyphrase generation model or not. We discover that adding sentences from the full text, particularly in the form of the extractive summary of the article can significantly improve the generation of both types of keyphrases that are either present or absent from the text. Experimental results with three widely used models for keyphrase generation along with one of the latest transformer models suitable for longer documents, Longformer Encoder-Decoder (LED) validate the observation. We also present a new large-scale scholarly dataset FULLTEXTKP for keyphrase generation. Unlike prior large-scale datasets, FULLTEXTKP includes the full text of the articles along with the title and abstract. We release the source code at <https://github.com/kgarg8/FullTextKP>.

1 Introduction

Keyphrases of scientific papers provide important topical information about the papers in a highly concise form and are crucial for understanding the evolution of ideas in a scientific field (Hall et al., 2008; Augenstein et al., 2017). Keyphrases of scientific papers can be also useful for mining and analyzing the literature (Augenstein et al., 2017) and for multiple downstream tasks such as index construction (Ritchie et al., 2006), summarization (Qazvinian et al., 2010), query formulation (Song et al., 2006), recommendation (Augenstein et al., 2017), reviewer matching for paper submissions (Augenstein et al., 2017) and for clustering papers for fast retrieval (Hammouda et al., 2005).

Identifying keyphrases from the scientific documents has been popularly studied under the paradigm of Keyphrase Extraction for over a decade (Mihalcea and Tarau, 2004; Wan and Xiao, 2008a; Florescu and Caragea, 2017; Liu et al., 2009, 2010; Sterckx et al., 2015, 2016; Caragea et al., 2014; Das Gollapalli and Caragea, 2014; Al-Zaidy et al., 2019; Patel and Caragea, 2021). However, one of the major limitations is that these approaches are not able to capture the semantic information in the document, particularly the *absent keyphrases* (keyphrases that are not present in the document).

In this paper, we therefore focus on a newer paradigm of keyphrase generation introduced by Meng et al. (2017). Instead of extracting keyphrases from the document, we model the problem as a neural sequence-to-sequence (seq2seq) approach where we learn to generate keyphrases in an autoregressive way. This approach can not only predict the exact keyphrases that are present in a document (*present keyphrases*) but also those that are semantically relevant but absent from the document (*absent keyphrases*). While many works (Meng et al., 2017; Chen et al., 2019b, 2020; Yuan et al., 2020; Ye et al., 2021b) adopt this approach, their focus is on architectural innovations to improve the generation of both present and absent keyphrases. In contrast, our focus is different from that of the prior works in keyphrase generation: we propose to extend the source of information of the input sequence, which is thus far limited to only title and abstract.

To this end, we curate a large-scale dataset of papers published by ACM¹ which contains not only the title and abstract of the documents but also the full text from the documents. We call this dataset FULLTEXTKP. Owing to the scarcity of such large-scale datasets with full text and the difficulty of parsing and understanding the entire text

¹<https://dl.acm.org/>

Method	Snippets from title and abstracts and beyond
Title	Performance of Power Detector Sensors of DTV Signals in IEEE 802.22 WRANs
Abstract	Sensing is the most important component in any cognitive radio system. The IEEE 802.22 Working Group (WG) is formulating the first worldwide standard for cognitive radios to operate in the television (TV) bands...
Citations	The IEEE 802.22 functional requirements document [2] states that spectrum sensing is required ...
Non-Citations	The keep-out region is a region around the primary user (e.g. DTV transmitter) where... Given these simulation scenarios we can evaluate various spectrum sensing techniques.
Summary	Of course, identification of which TV channels... This paper will describe ... in the IEEE 802.22 WG to evaluate spectrum sensing techniques... Calculating the base station keep-out region ...
Ret-Aug	The Bernoulli nonuniform sampling is further extended to matrix formulation, which allows the application of spectrum sensing for cognitive radio signal detection. - From Paper: 'Noise Enhancement and SNR Equivalence in Bernoulli Nonuniform Sampling'
Keyphrases	Cognitive Radio , IEEE 802.22 , Spectrum Sensing , Power Detector , Keep-out Region , False alarm probability , Misdetection probability

Table 1: A sample from FULLTEXTKP dataset showing PRESENT and ABSENT gold keyphrases in the different semantic spaces of the article. Refer to the Methods section for the description of the methods.

of the documents, the integration of additional information that goes beyond the title and abstract for keyphrase generation has been largely ignored.

In this paper, we provide innovative ways of using certain parts of the documents that could be rich source of information, e.g., using citation sentences from the content of the document or an extractive summary of the document (as illustrated in Table 1). Interestingly, through comprehensive experiments, we find that the citation contexts or provenance that had been a rich source of information for the keyphrase extraction task (Caragea et al., 2014; Das Gollapalli and Caragea, 2014) are not the richest sources of information in the keyphrase generation task. In contrast, we observe that the semantic information from the document could be best assimilated if we summarize the document and use the summary of the document instead of the document itself. To summarize each document, we use an unsupervised summarization approach and show that the extractive summary not only contains the highest topical information useful for keyphrase generation but can also fit the computational budget. Remarkably, the approach based on the incorporation of summary outperforms methods that use just title and abstract by a wide margin, oftentimes yielding 2-3 times improvements in performance in both present and absent keyphrases.

Our experiments are performed with three widely used models for keyphrase generation as well as, to our knowledge, for the first time in keyphrase generation, with a transformer model suited for long documents, Longformer Encoder-

Decoder (LED) (Beltagy et al., 2020).

Overall, we make the following contributions:

1. We explore the benefit of integrating additional information from different data sources (not just the title and abstract) into neural seq2seq models for keyphrase generation (i.e., predicting both present and absent keyphrases). The different data sources include random sentences from the article body,² sentences from the summary of the body, citation sentences, non-citation sentences, and sentences from other related documents in the training set.
2. We show that the sentences from the body of the article that form an extractive summary are great sources of finding good content for keyphrase generation.
3. We present a new dataset, FULLTEXTKP, of $\sim 140,000$ articles, each with its full text.

2 Related Work

Keyphrase Generation. Recent approaches to keyphrase generation have been dominated by neural seq2seq models because they provide a mechanism to also generate absent keyphrases. Meng et al. (2017) originally proposed an RNN-based seq2seq model along with CopyNet (CopyRNN) for keyphrase generation. Chen et al. (2018) improved CopyRNN by taking correlations between the predicted keyphrases into account (CorrRNN).

²Body includes all the sections of a scientific article besides the title and abstract

Yuan et al. (2020) extended the CopyRNN model to dynamically generate a variable number of diverse keyphrases. Chen et al. (2020) proposed an exclusive hierarchical decoding framework to capture the hierarchical compositionality of keyphrases (ExHiRD). Swaminathan et al. (2020) explored Generative Adversarial Networks for keyphrase generation. Ye et al. (2021b) modeled the target keyphrases as a set instead of a predefined sequence to improve the biased matching of the generated and target keyphrases. These methods use only title and abstract as the input sequence. Our work can aid all of these methods to further enhance the performance of keyphrase generation by augmenting the source texts with information from the body of the documents.

External Information. Some of the previous methods try to incorporate external information to enhance the performance of keyphrase extraction. For example, MAUI (Medelyan et al., 2009) uses semantic information from Wikipedia, CeKE (Caragea et al., 2014) and CiteTextRank (Gollapalli and Caragea, 2014) use information from citation networks, and SingleRank and ExpandRank (Wan and Xiao, 2008a) use information from local textual neighborhoods. However, these methods are limited to extraction-based methods whereas we focus on generating keyphrases to have a holistic or semantic understanding of the document when the natural language text is fed to the neural model.

Several works in keyphrase generation started to use external information to generate keyphrases. For example, Chen et al. (2019a), Santosh et al. (2021) proposed methods where they augmented a set of additional keyphrases from semantically similar documents and use hidden state representations built from these keyphrases, Diao et al. (2020) used cross-document attention networks, Ye et al. (2021a) used information from retrieved document-keyphrase pairs that are similar to the source document, Shen et al. (2021) used a phrase bank by pooling all keyphrases from all the articles. Similar to these works, one of our methods (Retrieval-Augmentation) retrieves relevant sentences from semantically similar documents and directly appends them to the title and abstract.

Datasets. There have been numerous efforts in the direction of curating large-scale datasets for either keyphrase extraction or keyphrase generation. In the non-scholarly domain, the datasets such as KPTime (Gallina et al., 2019), JPTimes (Gallina

et al., 2019), KPCrowd (Marujo et al., 2013), DUC-2001 (Wan and Xiao, 2008b), were curated from news articles, OpenKP (Xiong et al., 2019) was mined from webpages from a search engine. In the scholarly domain, there has been ongoing research with the datasets such as Krapivin (Krapivin et al., 2008), Inspec (Hulth, 2003), SemEval (Kim et al., 2010), NUS (Nguyen and Kan, 2007), KP20k (Meng et al., 2017) and OAGK (Cano and Bojar, 2018). The limitation of these datasets is that they are either very small (of the magnitude of a few hundreds or thousands) or they are bounded to title and abstract. For instance, the widely used KP20k dataset (Meng et al., 2017) consists of only the title and abstract of the research articles and the average words per document is as low as 176. To leverage and understand the benefit of additional data, we mined a new dataset based on the articles primarily from the ACM database which contains not just the titles and abstracts but also the full text of the scholarly documents. In our dataset, the typical length of full papers is 5000-10000 tokens.

3 Methods

We use four different models, three of which are popular models in the Keyphrase Generation task, viz. catSeq (Yuan et al., 2020), One2Set (Ye et al., 2021b), ExHiRD (Chen et al., 2020); and the fourth model is based on one of the latest Transformer models, i.e., Longformer Encoder-Decoder (LED) (Beltagy et al., 2020), suitable for long documents. The baseline in each model takes the title and abstract (T+A) as input and predicts a sequence of keyphrases. We explore different extensions (also illustrated in Table 1) for each model by concatenating different types of data to the title and abstract (T+A). During the concatenation of additional information, we use a delimiter <sep> between the title and abstract and between all the additional sentences as follows: Title <sep> Abstract <sep> Sent₁ <sep> Sent₂ <sep> ... Sent_k. We describe different types of data (input sequence) below:

T+A+RANDOM: For this setup, we concatenate k randomly chosen sentences (from the body) to the title and abstract in the original order of their occurrence in the body.

T+A+CITATIONS: For this setup, we collect all the sentences (from the body) that have cited some other article. We call such sentences as *citation sentences*. We randomly select k sentences from the entire pool and concatenate them to the title and

abstract in the order of their occurrence in the body of the given article. The incorporation of citation sentences is inspired by Naïve-Bayes or TextRank-based approaches (Gollapalli and Caragea, 2014; Caragea et al., 2014) or construction of Citation-Graph (Viswanathan et al., 2021) which enhanced their keyphrase extraction performance through the integration of citation information.

T+A+NON-CITATIONS: For this setup, we collect all the sentences (from the body) that do not cite any other article. We randomly select k sentences from the entire pool and concatenate them to the title and abstract in the order of their occurrence in the body of the given article.

T+A+SUMMARY: In this method, we summarize the body of the article using a state-of-the-art unsupervised summarization algorithm, *PacSum* (Zheng and Lapata, 2019) and append k sentences from the summary to T+A of the article. Zheng and Lapata showed that TF-IDF-based *PacSum* is better than TextRank and other baselines, and is also highly competitive with even BERT-based *PacSum*. *PacSum* can be replaced with other summarization algorithms but we choose TF-IDF-based *PacSum* simply as a representative of a powerful and efficient model for unsupervised summarization to test the efficacy of summarizing the body.

The steps to obtain the extractive summary of the documents using *PacSum* are as follows. Consider a document as a list of sentences and each sentence as a list of words. The entire document can be visualized as a directed graph in which the individual sentences are the nodes and the edges are weighted by similarity. We first calculate the tf-idf to understand the relevance of a word to a sentence in the collection of sentences (document), quite analogous to computing tf-idf of a word of a document in a corpus of documents. We then compute the similarity between every pair of sentences. Next, we compute a threshold based on which we normalize the similarity scores. Subsequently, we compute forward and backward edge scores for the directed edges. Afterward, we compute the degree centrality of a sentence (node) in the document (graph) as the weighted average of forward and backward edge scores. Finally, we rank the sentences based on their centrality values, select the k highest ranked sentences, and concatenate them to the title and abstract. We describe the algorithm in Appendix C.

T+A+RETRIEVAL-AUGMENTATION: In this

method, we retrieve and augment k semantically similar sentences from the training corpus to the T+A of each article. To this end, we first create a set of all the sentences from the titles and abstracts of the articles in the entire *training* dataset. Next, we embed each sentence in the set using SPECTER (Cohan et al., 2020). We treat these embeddings as *key embeddings* representing the corresponding sentences (*values*). Given a target article (*query*), we then embed its title and abstract using SPECTER. This embedding serves as a *query embedding* to search related sentences from other articles. Subsequently, we compute the dot-product *similarity* of the query embedding with all the key embeddings using FAISS (Johnson et al., 2019). Last, we select k sentences (*values*) corresponding to the top k most similar key embeddings and concatenate the sentences to the title and abstract of the query article. We provide the algorithm in Appendix D.

4 Experimental Setup

4.1 FULLTEXTKP dataset

To evaluate the performance of models leveraging different types of information, we construct a new dataset, which we call FULLTEXTKP. Our FULLTEXTKP dataset consists of research papers that are published by ACM³ and are available in the ACM digital library in its International Conference Proceedings Series (ICPS).

We used only the articles which have at least the five fields, viz., title, abstract, keywords, full text and references. We lowercased the text and constructed numerous regexes to remove the escape sequences, html tags, urls, emails, etc. We also replaced all the numbers and roman numerals with a <digit> token. Further, we used PunktSentenceTokenizer⁴ to segment the document into sentences and NLTK’s word_tokenizer⁴ to tokenize the sentences into tokens. We constructed suitable regex to extract the citation sentences and we considered all other sentences as non-citation sentences. The articles without any citation sentences and duplicates were further removed from the collection. Thus, we were able to collect 142,844 articles. We split the FULLTEXTKP dataset into 80/10/10 for train, test and validation sets. Statistical information about

³The full ACM dataset is available from the ACM Digital Library by request.

⁴https://www.nltk.org/_modules/nltk/tokenize/punkt.html

Total papers in the corpus	142,844
Training set size	114,271
Validation set size	14,287
Test set size	14,286
% of present keyphrases	55.8
% of absent keyphrases	44.2
Average keyphrases per paper	4.3

Table 2: FULLTEXTKP dataset statistics

the dataset is shown in Table 2.

The construction of this dataset addresses the sparsity of large-scale datasets for keyphrase extraction/generation from scientific papers and is aimed at enabling deep learning modeling. Currently, there exist only a couple of large-scale datasets for this task, e.g., KP20K (Meng et al., 2017), and OAGK (Cano and Bojar, 2018). Unlike these datasets, which contain only the title and abstract of each paper, our dataset provides access to the full text of each paper.

4.2 Evaluation

We use the same evaluation method as Chan et al. (2019); Chen et al. (2020) and report the macro-averaged Precision (**P**), Recall (**R**) and F1-scores (**F₁**) in two different evaluation settings: **@5** (dummy keyphrases are appended to the predicted keyphrases to make the total count as 5) and **@M** (directly compare the generated keyphrases against the gold keyphrases). All the keyphrases are stemmed using PorterStemmer⁵ before comparison. As in prior works, we treat a keyphrase as *absent* if it is absent from the title and abstract, which helps us compare the performance of generating *absent* and *present* keyphrases for different models in a consistent way.

For the entire data including Title, Abstract and the additional data using different methods, we restrict the maximum sequence length to 800. This requires us to use an adaptive k for the additional number of sentences that can fit in the given maximum sequence length. The selection of the sequence length was based on the tradeoff between the performance and computational budget of the models. We also experiment with longer sequence lengths for the LED model in §5.3.

⁵https://www.nltk.org/_modules/nltk/stem/porter.html

5 Results

In this section, we first present the results for present and absent keyphrase generation performance in §5.1 for the different methods described in §3. We further explore our best performing method, i.e., T+A+summary into two sub-categories: extractive and abstractive summarization methods in §5.2. Finally, in §5.3 we explore the performance of the best performing model, i.e., Longformer Encoder-Decoder with sequence lengths much higher than 800. We run all the experiments three times and report the average.

5.1 Present & Absent Keyphrase Generation

Tables 3 and 4 present the results for the present and absent keyphrase performances, respectively, for our different experiments. The results demonstrate that adding extra information from the body of the given article to the baseline (T+A) is beneficial for both present and absent keyphrase generation. Augmenting the summary sentences provides the most substantial boost compared to any other method for both present and absent keyphrase performance. For instance, improvements in present keyphrase performance on F1@M metric are: catSeq (0.338 → 0.380), ExHiRD (0.325 → 0.366), LED (0.360 → 0.397). The absent keyphrase performance using the summary method improves up to four times the baseline performance; for example, catSeq (0.021 → 0.079), ExHiRD (0.030 → 0.080), LED (0.061 → 0.106). Interestingly, with One2Set, although F1@M for present keyphrase generation does not improve with Summary and Random methods, the trend for absent keyphrase performance is very similar to that of the other models. The discrepancy in the present keyphrase generation performance could be because One2Set might potentially be biased towards identifying the keyphrases from the earlier portions of the documents. For our setting of long documents, it may require better heuristics to initialize the control codes so as to attend to the later portions of the document. Still, in three out of four models, adding summary information gives us the best results. This seems to follow the natural intuition that the summary of the article contains the most topical information useful for the generation the keyphrases.

Figure 1 further validates the above intuition. Summary method is the richest source of present keyphrases. The total number of present keyphrases in the texts of all Summary samples

Model	Method	P@5	R@5	F1@5	P@M	R@M	F1@M
catSeq	Title+Abstract	0.167 ₂	0.343 ₃	0.224 ₂	0.333 ₃	0.343 ₄	0.338 ₁
	+ Ret-Aug	0.166 ₀	0.342 ₁	0.224 ₀	0.338 ₃	0.343 ₁	0.340 ₁
	+ Citations	0.170 ₁	0.351 ₃	0.229 ₂	0.360 ₂	0.351 ₃	0.355 ₁
	+ Non-Citations	0.178 ₁	0.365 ₃	0.240 ₂	0.365 ₃	0.366 ₃	0.366 ₂
	+ Random	0.175 ₂	0.359 ₅	0.236 ₃	0.370 ₄	0.359 ₅	0.364 ₄
	+ Summary	0.187₁	0.380₂	0.250₂	0.379₄	0.380₂	0.380₂
One2Set	Title+Abstract	0.203₃	0.418₄	0.273₃	0.300 ₆	0.425₆	0.351 ₂
	+ Ret-Aug	0.178 ₄	0.368 ₉	0.241 ₅	0.313 ₅	0.370 ₉	0.339 ₃
	+ Citations	0.192 ₁	0.394 ₄	0.258 ₂	0.316 ₅	0.396 ₃	0.351 ₁
	+ Non-Citations	0.192 ₆	0.395 ₁₃	0.259 ₉	0.316 ₆	0.397 ₁₄	0.352 ₂
	+ Random	0.191 ₄	0.391 ₇	0.257 ₄	0.314 ₉	0.394 ₈	0.349 ₃
	+ Summary	0.189 ₅	0.385 ₁₃	0.253 ₉	0.325₇	0.387 ₁₄	0.353₂
ExHiRD	Title+Abstract	0.184 ₂	0.381 ₃	0.236 ₂	0.320 ₃	0.382 ₃	0.325 ₁
	+ Ret-Aug	0.187 ₀	0.390 ₂	0.241 ₁	0.317 ₂	0.391 ₂	0.325 ₂
	+ Citations	0.194 ₀	0.402 ₁	0.249 ₀	0.336 ₁	0.403 ₁	0.340 ₀
	+ Non-Citations	0.197 ₁	0.410 ₆	0.257 ₈	0.353 ₂	0.407 ₁	0.354 ₁
	+ Random	0.195 ₀	0.403 ₁	0.249 ₀	0.352 ₂	0.404 ₁	0.351 ₁
	+ Summary	0.202₂	0.416₃	0.258₂	0.369₁	0.416₄	0.366₀
LED	Title+Abstract	0.193 ₁₁	0.389 ₂₁	0.258 ₁₅	0.333 ₁	0.392 ₂₂	0.360 ₉
	+ Ret-Aug	0.193 ₆	0.391 ₁₀	0.258 ₇	0.319 ₄	0.394 ₁₁	0.353 ₂
	+ Citations	0.201 ₅	0.403 ₉	0.268 ₆	0.338 ₄	0.407 ₁₀	0.369 ₂
	+ Non-Citations	0.209 ₃	0.421 ₆	0.280 ₄	0.348 ₉	0.425 ₇	0.382 ₃
	+ Random	0.213 ₀	0.428 ₁	0.284 ₀	0.354 ₆	0.431 ₀	0.389 ₃
	+ Summary	0.217₄	0.431₈	0.289₅	0.365₂	0.435₈	0.397₃

Table 3: Results comparing present keyphrase performance with different methods of using additional data with the FULLTEXTKP dataset. Subscripts denote the standard deviation as multiples of ± 0.001 .

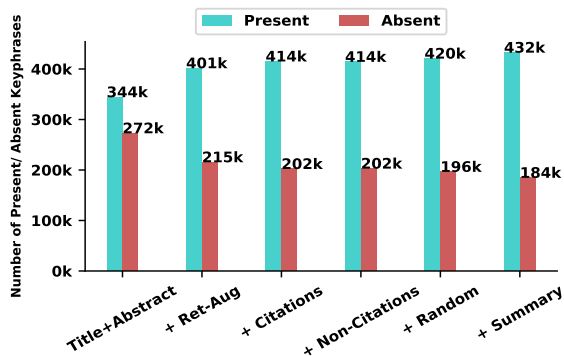


Figure 1: #Present and #Absent Keyphrases in terms of the whole input text (for various methods) in FULLTEXTKP dataset. + indicates concatenation of title and abstract with the specified information.

is about $\sim 88,000$ higher than just Title+Abstract samples. So, the model learns to generalize better when trained on such highly topical input texts of Summary method.

Interestingly, we find that using only the citation sentences results in worse performance compared to using non-citation sentences or just random sentences. This observation is in contrast to the observation of the use of citation sentences for the keyphrase extraction methods. We hypothesize the reason for such an observation is that keyphrase extraction methods focus mainly on identifying

keyphrases by capturing the statistical information such as tf-idf or word co-occurrences that appear in the target document and the cited ones (through the citation sentences), whereas the deep learning methods try to get a semantic or holistic understanding of the document itself (i.e., from all sentences).

Incorporating either non-citation sentences or random sentences gives very similar results because the entire text contains majority of non-citation sentences. Unfortunately, appending the sentences retrieved from the semantically similar documents (using SPECTER) performs worse than the baseline. We hypothesize that even though the sentences have semantic similarity on the embedding level, they could still be significantly different in terms of their meaning and may not contain exactly the same keyphrases. Rather, a lot of such sentences may confuse and deviate the model from predicting the gold keyphrases.

Among the four different models, Longformer Encoder-Decoder (LED) performs the best. We hypothesize that this happens because it is based on the more sophisticated Transformer model architecture and also because it is suited for reading long documents. However, LED training takes at least 5-10 times more than the catSeq model. We provide more details about the compute power and

Model	Method	P@5	R@5	F1@5	P@M	R@M	F1@M
catSeq	Title+Abstract	0.007 ₀	0.017 ₁	0.010 ₁	0.030 ₂	0.017 ₁	0.021 ₁
	+ Ret-Aug	0.009 ₀	0.021 ₁	0.012 ₁	0.037 ₁	0.021 ₁	0.027 ₁
	+ Citations	0.013 ₁	0.031 ₂	0.019 ₂	0.052 ₃	0.031 ₂	0.039 ₃
	+ Non-Citations	0.021 ₀	0.046 ₁	0.029 ₀	0.068 ₁	0.046 ₁	0.055 ₀
	+ Random	0.022 ₁	0.049 ₂	0.030 ₁	0.073 ₃	0.049 ₂	0.059 ₃
	+ Summary	0.033₁	0.069₂	0.045₁	0.093₂	0.069₂	0.079₂
One2Set	Title+Abstract	0.018 ₁	0.042 ₂	0.026 ₁	0.043 ₁	0.042 ₂	0.043 ₀
	+ Ret-Aug	0.017 ₀	0.039 ₂	0.023 ₁	0.044 ₂	0.039 ₂	0.041 ₁
	+ Citations	0.021 ₁	0.047 ₂	0.029 ₁	0.053 ₁	0.047 ₂	0.049 ₁
	+ Non-Citations	0.023 ₁	0.053₂	0.032 ₁	0.054 ₁	0.053₂	0.053 ₂
	+ Random	0.024₀	0.053₂	0.033₁	0.056₁	0.053₂	0.054₂
	+ Summary	0.023 ₃	0.050 ₆	0.032 ₄	0.054 ₄	0.051 ₆	0.052 ₅
ExHiRD	Title+Abstract	0.012 ₀	0.027 ₂	0.016 ₁	0.042 ₁	0.026 ₀	0.030 ₀
	+ Ret-Aug	0.010 ₀	0.022 ₀	0.013 ₀	0.035 ₁	0.022 ₀	0.025 ₀
	+ Citations	0.019 ₀	0.041 ₀	0.024 ₀	0.060 ₁	0.041 ₀	0.045 ₁
	+ Non-Citations	0.024 ₀	0.051 ₀	0.031 ₀	0.071 ₁	0.051 ₀	0.055 ₀
	+ Random	0.026 ₀	0.054 ₁	0.032 ₀	0.076 ₁	0.054 ₁	0.059 ₁
	+ Summary	0.037₁	0.074₃	0.046₁	0.100₂	0.074₃	0.080₂
LED	Title+Abstract	0.027 ₀	0.063 ₁	0.038 ₀	0.061 ₁₁	0.063 ₁	0.061 ₅
	+ Ret-Aug	0.014 ₁	0.034 ₂	0.020 ₁	0.053 ₁	0.034 ₂	0.041 ₁
	+ Citations	0.025 ₂	0.058 ₄	0.035 ₃	0.073 ₂	0.058 ₄	0.065 ₃
	+ Non-Citations	0.033 ₁	0.071 ₂	0.045 ₁	0.088 ₁	0.071 ₂	0.079 ₁
	+ Random	0.036 ₂	0.078 ₄	0.049 ₂	0.093 ₁	0.078 ₄	0.085 ₂
	+ Summary	0.047₁	0.100₁	0.064₁	0.112₄	0.100₁	0.106₂

Table 4: Results comparing absent keyphrase performance with different methods of using additional data with the FULLTEXTKP dataset. Subscripts denote the standard deviation as multiples of ± 0.001 .

Method	P@5	R@5	F1@5	P@M	R@M	F1@M
Abs_Sum	0.201 ₁	0.405 ₂	0.268 ₁	0.326 ₄	0.409 ₂	0.363 ₃
Ext_Sum	0.217₄	0.431₈	0.289₅	0.365₂	0.435₈	0.397₃
Abs_Sum	0.020 ₁	0.048 ₂	0.029 ₁	0.062 ₁	0.048 ₂	0.054 ₁
Ext_Sum	0.047₁	0.100₁	0.064₁	0.112₄	0.100₁	0.106₂

Table 5: Results comparing the performance of Abstractive Summarization vs. Extractive Summarization methods. Subscripts denote the standard deviation as multiples of ± 0.001 .

time in Appendix B.

We performed a 2-tailed statistical significance test with alpha⁶ value as 0.05 and the p-values⁷ less than 0.05 were considered statistically different. We observed that the results were statistically different for all models except One2Set. With One2Set, the additional information to the T+A baseline does not yield statistically different results.

5.2 Analysis on Summarization Methods

In this section, we further investigate the capability of our best performing method, i.e., extractive summarization, for keyphrase generation in comparison with an abstractive summarization ap-

⁶Probability of rejecting the null hypothesis when it was actually true

⁷Probability of getting a result that is as extreme or more extreme when the null hypothesis is true

proach, abbreviated as $T+A+Abs_Sum$ or simply Abs_Sum . For the abstractive summarization, we use BigBirdPegasus (Zaheer et al., 2020) model, pretrained on ArXiv dataset (Cohan et al., 2018) to first generate the abstractive summaries. Next, we augment the sentences from abstractive summaries to the title and abstract of the articles. Further, we train our best performing model (i.e., LED) using $T+A+Abs_Sum$ method. In Table 5, we compare the results of $T+A+Abs_Sum$ with the $T+A+Summary$ (extractive) method. Note that, by default, we use Summary for the Extractive Summarization method (Zheng and Lapata, 2019), also abbreviated as Ext_Sum .

From Table 5, we observe that the extractive summary added to T+A performs better than the abstractive summary (added to T+A) by 0.034 (F1@M) and 0.052 (F1@M) for present and absent keyphrase generation, respectively. The behavior could be explained as follows: (1) Extractive summaries contain sentences that are more central to the paper (Zheng and Lapata, 2019). Thus, they tend to bring complementary information from all parts of the paper that improves the performance of keyphrase generation models; and (2) We manually inspected a small subset of the dataset and observed that about two-thirds of the sentences in the abstrac-

Type	Method	Seq-Len	P@5	R@5	F@5	P@M	R@M	F1@M
Present Keyphrases	Summary	800	0.217 ₄	0.431 ₈	0.289 ₅	0.365 ₂	0.435 ₈	0.397 ₃
	Random	800	0.213 ₀	0.428 ₁	0.284 ₀	0.354 ₆	0.431 ₀	0.389 ₃
		1500	0.211 ₂	0.424 ₃	0.282 ₂	0.370 ₃	0.427 ₃	0.397 ₁
		2000	0.215 ₂	0.428 ₃	0.286 ₂	0.372 ₄	0.431 ₃	0.399 ₄
		2500	0.212 ₄	0.425 ₆	0.283 ₅	0.389 ₁	0.427 ₇	0.407 ₃
Absent Keyphrases	Summary	800	0.047 ₁	0.100 ₁	0.064 ₁	0.112 ₄	0.100 ₁	0.106 ₂
	Random	800	0.036 ₂	0.078 ₄	0.049 ₂	0.093 ₁	0.078 ₄	0.085 ₂
		1500	0.044 ₀	0.097 ₁	0.061 ₁	0.110 ₂	0.097 ₁	0.103 ₁
		2000	0.047 ₂	0.100 ₃	0.063 ₂	0.114 ₂	0.100 ₃	0.107 ₃
		2500	0.049 ₂	0.108 ₄	0.067 ₂	0.126 ₂	0.108 ₄	0.116 ₃

Table 6: Results comparing the performance of Longformer Encoder-Decoder Model with source texts of different sequence lengths. Subscripts denote the standard deviation as multiples of ± 0.001 .

tive summaries were from the Introduction section of the papers, which although contain extended but potentially overlapping or paraphrased description of the Abstract section. Further, we observed that there were about 10-15% repetitive statements in the abstractive summaries. We conclude that augmenting repetitive statements or statements with redundant information as the abstract do not aid the keyphrase generation performance.

5.3 Exploring longer sequence lengths with Longformer Encoder-Decoder

Scientific documents are often very long sequences of text, ranging from a few thousands to tens of thousands of tokens. Longformer Encoder-Decoder, proposed by Beltagy et al. (2020), has the capability to deal with long sequences of text. We therefore explore the effect of sequence lengths longer than 800 on the present and absent keyphrase performance. We prepare different versions of the *T+A+Random* with the maximum sequence lengths of 800, 1500, 2000 and 2500. We chose the *Random* version instead of directly using the entire source document since *Random* would be a better representative of the long document if we need to truncate the sequence length to a fixed budget (i.e., rather than truncating and completely removing a part of the document, we ensure coverage of sentences from the entire document). Note that these sequence length values are based on the word-level tokenization based on NLTK’s word-tokenizer. The sequence length value based on subword tokenization (used by Transformers) is actually much longer, but fortunately, LED is able to handle sequence length up to 16000 tokens.

The results in Table 6 show that both present and

absent keyphrase performance generally increases as we increase the source text length. In the table, we also compare the performance of these different versions of *T+A+Random* with the *T+A+Summary* with a much smaller sequence length of 800, so as to justify the purpose of smaller intended summary. We observe that the performance of *Summary-800* is quite comparable to *Random-2000* for both present and absent keyphrase performance, particularly for the most important metric F1@M. This further validates our hypothesis that summary indeed gives us a richer source of information. Just the sequence length of 800 can condense the information that could otherwise be extracted using sequence length of 2000. Furthermore, the summary has the benefit of fitting a smaller memory and computational resources budget. We observed that with the increasing sequence lengths, the training time increased at least at a linear rate and used considerable amount of GPU and RAM memory. In the case we add more sentences to the summary, we expect the quality to degrade and converge to that of the *Random* versions and furthermore, it will no more stand out as a summary of the document.

In Table 7, we show sample predictions for the best performing method *T+A+ (Extractive) Summary* using LED model. The gold keyphrases and the predicted keyphrases are tokenized and stemmed before comparison. We make the following observations from the table. First, the model can predict accurately both the present and absent keyphrases. Second, some predicted keyphrases are near-matches of the gold keyphrases, and despite their quality as keyphrases, due to the limitation of the metric (exact match), they do not contribute to the model performance.

Snippets from title and abstracts and summary with model predictions	Gold Keyphrases
Segmentation of brain MR images using intuitionistic fuzzy clustering algorithm <sep> A new algorithm ... using intuitionistic fuzzy clustering (IFCM), is proposed in this paper.	segmentation, brain mr image, intuitionistic fuzzy set
GPU -based simulation of wireless body area network <sep> ... in a local workstation to perform model simulations of a ... the goal of this project is to gain an understanding of the gpu -based performance of wban model simulation ...	gpu, wireless body area network, model, matlab
Spectrum sharing for directional systems <sep> Dynamic spectrum access systems will need to ... An example involving a satellite downlink antenna and a broadband wireless access system using directional antennas is presented ...	dynamic spectrum access, directionality
Automatic classification of anuran sounds using convolutional neural networks <sep> ... networks with mel-frequency cepstral coefficients (mfccs) as input for the task of ... <sep> Different machine learning approaches for anuran classification ...	anuran, convolutional neural networks, wireless sensor networks, mfcc, machine learning

Table 7: Predicted keyphrases by Longformer Encoder-Decoder model are highlighted in cyan . The source text is in the format: Title <sep> Abstract <sep> Summary-Sent_1 ... <sep> Summary-Sent_k

6 Conclusion

In this paper, we explore numerous ways of incorporating additional data from the body of the scholarly documents for improving the performance of the keyphrase generation task. We conclude from the results that the extractive summary sentences, that are more central to the paper, provide the most topical information for boosting both present and absent keyphrase generation performance. The citation, non-citation, and random sentences also bring complementary information to improve the performance modestly. Some other ways such as augmenting semantically similar sentences from other papers or augmenting abstractive summary sentences, that bring repetitive or redundant information to titles and abstracts, are not effective. We present a comprehensive analysis with four models including LED, yet unexplored for the task. Our work aims at breaking apart the barriers of using only titles and abstracts, and presents a large-scale dataset with full texts for keyphrase generation.

7 Limitations

One of the limitations of the proposed methods is the increased (up to 2-3 times) compute time and memory for the training of the models compared to the conventional training of the models using only T+A of the articles. We provide more details in Appendix B. Further, our best performing method $T+A+Summary$ requires an additional modest overhead of pre-computing summaries for all the articles.

Another potential limitation of our work is that we can not directly compare our models on the widely used datasets, e.g., KP20k, Inspec, Krapivin,

NUS since these datasets do not have full texts of the papers. To be comprehensive, we considered the performance of four models on the new dataset.

We encourage future work in the direction of better ways of integrating external information including more sophisticated approaches for the summarization of scientific documents.

8 Ethical Considerations

Since keyphrase generation finds direct application in many downstream tasks such as recommender systems, reviewer matching, articles clustering for fast retrieval, etc., it is important to consider the ethical implications of these models. The models are not perfect and might make misleading predictions at times. The users must make their own discrimination in deploying the current models for decision-making systems.

Acknowledgements

This research is supported in part by NSF CAREER award #1802358, NSF CRI award #1823292, NSF IIS award #2107518, and UIC Discovery Partners Institute (DPI) award. Any opinions, findings, and conclusions expressed here are those of the authors and do not necessarily reflect the views of NSF or DPI. We thank AWS for computational resources used in this study. We also thank our anonymous reviewers for their constructive feedback and suggestions.

References

Rabah Al-Zaidy, Cornelia Caragea, and C. Lee Giles. 2019. Bi-lstm-crf sequence labeling for keyphrase

- extraction from scholarly documents. In *Proceedings of The Web Conference (WWW 2019)*, San Francisco, California, USA.
- Isabelle Augenstein, Mrinal Das, Sebastian Riedel, Lakshmi Vikraman, and Andrew McCallum. 2017. [SemEval 2017 task 10: ScienceIE - extracting keyphrases and relations from scientific publications](#). In *SemEval (2017)*, pages 546–555. ACL.
- Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. *arXiv:2004.05150*.
- Erion Cano and Ondřej Bojar. 2018. Keyphrase generation: A text summarization struggle. *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Long Papers)*, Melbourne, Australia, page 1383–1392.
- Cornelia Caragea, Florin Adrian Bulgarov, Andreea Godea, and Sujatha Das Gollapalli. 2014. [Citation-enhanced keyphrase extraction from research papers: A supervised approach](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1435–1446, Doha, Qatar. Association for Computational Linguistics.
- Hou Pong Chan, Wang Chen, Lu Wang, and Irwin King. 2019. [Neural keyphrase generation via reinforcement learning with adaptive rewards](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2163–2174, Florence, Italy. Association for Computational Linguistics.
- Jun Chen, Xiaoming Zhang, Yu Wu, Zhao Yan, and Zhoujun Li. 2018. [Keyphrase generation with correlation constraints](#). In *EMNLP*, pages 4057–4066.
- Wang Chen, Hou Pong Chan, Piji Li, Lidong Bing, and Irwin King. 2019a. [An integrated approach for keyphrase generation via exploring the power of retrieval and extraction](#). In *Proceedings of the 2019 Conference of the NAACL, Volume 1*, pages 2846–2856. ACL.
- Wang Chen, Hou Pong Chan, Piji Li, and Irwin King. 2020. [Exclusive hierarchical decoding for deep keyphrase generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1095–1105, Online. Association for Computational Linguistics.
- Wang Chen, Yifan Gao, Jiani Zhang, Irwin King, and Michael R Lyu. 2019b. [Title-guided encoding for keyphrase generation](#). In *AAAI*, pages 6268–6275.
- Arman Cohan, Franck Dernoncourt, Doo Soon Kim, Trung Bui, Seokhwan Kim, Walter Chang, and Nazli Goharian. 2018. A discourse-aware attention model for abstractive summarization of long documents. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2:615–621.
- Arman Cohan, Sergey Feldman, Iz Beltagy, Doug Downey, and Daniel S Weld. 2020. [Specter: Document-level representation learning using citation-informed transformers](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2270–2282.
- Sujatha Das Gollapalli and Cornelia Caragea. 2014. [Extracting keyphrases from research papers using citation networks](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 28(1).
- Shizhe Diao, Yan Song, and Tong Zhang. 2020. [Keyphrase generation with cross-document attention](#). *arXiv preprint arXiv:2004.09800*.
- Corina Florescu and Cornelia Caragea. 2017. [PositionRank: An unsupervised approach to keyphrase extraction from scholarly documents](#). In *Proceedings of the 55th ACL (Vol 1)*, pages 1105–1115. ACL.
- Ygor Gallina, Florian Boudin, and Beatrice Daille. 2019. [KPTimes: A large-scale dataset for keyphrase generation on news documents](#). In *Proceedings of the 12th International Conference on Natural Language Generation*, pages 130–135, Tokyo, Japan. Association for Computational Linguistics.
- Sujatha Das Gollapalli and Cornelia Caragea. 2014. [Extracting keyphrases from research papers using citation networks](#). In *AAAI*, pages 1629–1635.
- David Hall, Daniel Jurafsky, and Christopher D. Manning. 2008. [Studying the history of ideas using topic models](#). In *EMNLP*, pages 363–371.
- Khaled M. Hammouda, Diego N. Matute, and Mohamed S. Kamel. 2005. [Corephrase: Keyphrase extraction for document clustering](#). In *Proceedings of the 4th International Conference on Machine Learning and Data Mining in Pattern Recognition*, MLDM’05, pages 265–274, Berlin, Heidelberg. Springer-Verlag.
- Anette Hulth. 2003. [Improved automatic keyword extraction given more linguistic knowledge](#). In *Proceedings of the 2003 conference on Empirical methods in natural language processing*, pages 216–223.
- Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2019. [Billion-scale similarity search with gpus](#). *IEEE Transactions on Big Data*.
- Su Nam Kim, Olena Medelyan, Min-Yen Kan, and Timothy Baldwin. 2010. [Semeval-2010 task 5: Automatic keyphrase extraction from scientific articles](#). In *SemEval ’10*, pages 21–26. ACL.
- Diederik P Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). *International Conference on Learning Representations (ICLR)*.
- Mikalai Krapivin, Aliaksandr Autayeu, and Maurizio Marchese. 2008. [Large dataset for keyphrases extraction](#). Technical Report DISI-09-055.

- Zhiyuan Liu, Wenyi Huang, Yabin Zheng, and Maosong Sun. 2010. [Automatic keyphrase extraction via topic decomposition](#). In *EMNLP '10*, pages 366–376. ACL.
- Zhiyuan Liu, Peng Li, Yabin Zheng, and Maosong Sun. 2009. [Clustering to find exemplar terms for keyphrase extraction](#). In *2009 EMNLP*.
- Lúís Marujo, Anatole Gershman, Jaime Carbonell, Robert Frederking, and João P Neto. 2013. Supervised topical key phrase extraction of news stories using crowdsourcing, light filtering and co-reference normalization. *arXiv preprint arXiv:1306.4886*.
- Olena Medelyan, Eibe Frank, and Ian H. Witten. 2009. [Human-competitive tagging using automatic keyphrase extraction](#). In *EMNLP '09*, pages 1318–1327. ACL.
- Rui Meng, Sanqiang Zhao, Shuguang Han, Daqing He, Peter Brusilovsky, and Yu Chi. 2017. [Deep keyphrase generation](#). In *ACL*, pages 582–592. ACL.
- Rada Mihalcea and Paul Tarau. 2004. Texttrank: Bringing order into text. In *2004 EMNLP*.
- Thuy Dung Nguyen and Min-Yen Kan. 2007. Keyphrase extraction in scientific publications. In *International conference on Asian digital libraries*, pages 317–326. Springer.
- Krutarth Patel and Cornelia Caragea. 2021. [Exploiting position and contextual word embeddings for keyphrase extraction from scientific papers](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1585–1591, Online. Association for Computational Linguistics.
- Vahed Qazvinian, Dragomir R. Radev, and Arzucan Özgür. 2010. [Citation summarization through keyphrase extraction](#). In *Proceedings of the 23rd International Conference on Computational Linguistics, COLING '10*, pages 895–903, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Anna Ritchie, Simone Teufel, and Stephen Robertson. 2006. [How to find better index terms through citations](#). In *Proceedings of the Workshop on How Can Computational Linguistics Improve Information Retrieval?*, CLIR '06, pages 25–32, Stroudsburg, PA, USA. Association for Computational Linguistics.
- TYSS Santosh, Debarshi Kumar Sanyal, Plaban Kumar Bhowmick, and Partha Pratim Das. 2021. Gazetteer-guided keyphrase generation from research papers. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pages 655–667. Springer.
- Xianjie Shen, Yinghan Wang, Rui Meng, and Jingbo Shang. 2021. Unsupervised deep keyphrase generation. *arXiv preprint arXiv:2104.08729*.
- I. Y. Song, R. B. Allen, Z. Obradovic, and M. Song. 2006. [Keyphrase extraction-based query expansion in digital libraries](#). In *Proceedings of the 6th ACM/IEEE-CS Joint Conference on Digital Libraries (JCDL '06)*, pages 202–209.
- Lucas Sterckx, Cornelia Caragea, Thomas Demeester, and Chris Develder. 2016. [Supervised keyphrase extraction as positive unlabeled learning](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1924–1929, Austin, Texas. Association for Computational Linguistics.
- Lucas Sterckx, Thomas Demeester, Johannes Deleu, and Chris Develder. 2015. [Topical word importance for fast keyphrase extraction](#). In *WWW'15*, pages 121–122. ACM.
- Avinash Swaminathan, Haimin Zhang, Debanjan Mahata, Rakesh Gosangi, Rajiv Ratn Shah, and Amanda Stent. 2020. [A preliminary exploration of GANs for keyphrase generation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8021–8030, Online. Association for Computational Linguistics.
- Vijay Viswanathan, Graham Neubig, and Pengfei Liu. 2021. Citationie: Leveraging the citation graph for scientific information extraction. In *Joint Conference of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (ACL-IJCNLP)*, Virtual.
- Xiaojun Wan and Jianguo Xiao. 2008a. [Single document keyphrase extraction using neighborhood knowledge](#). In *23rd National Conference on AI, AAAI'08*, pages 855–860. AAAI Press.
- Xiaojun Wan and Jianguo Xiao. 2008b. Single document keyphrase extraction using neighborhood knowledge. In *AAAI*, volume 8, pages 855–860.
- Lee Xiong, Chuan Hu, Chenyan Xiong, Daniel Campos, and Arnold Overwijk. 2019. [Open domain web keyphrase extraction beyond language modeling](#). 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 5174–5183, Hong Kong, China. Association for Computational Linguistics.
- Jiacheng Ye, Ruijian Cai, Tao Gui, and Qi Zhang. 2021a. Heterogeneous graph neural networks for keyphrase generation. *arXiv preprint arXiv:2109.04703*.
- Jiacheng Ye, Tao Gui, Yichao Luo, Yige Xu, and Qi Zhang. 2021b. [One2set: Generating diverse keyphrases as a set](#). In *Proceedings of ACL*.
- Xingdi Yuan, Tong Wang, Rui Meng, Khushboo Thaker, Peter Brusilovsky, Daqing He, and Adam Trischler. 2020. [One size does not fit all: Generating and evaluating variable number of keyphrases](#). In *Proceedings*

of the 58th Annual Meeting of the Association for Computational Linguistics, pages 7961–7975, Online. Association for Computational Linguistics.

Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, et al. 2020. Big bird: Transformers for longer sequences. *Advances in Neural Information Processing Systems*, 33:17283–17297.

Hao Zheng and Mirella Lapata. 2019. Sentence centrality revisited for unsupervised summarization. *In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6236–6247.

Methods	Time	CUDA Memory
Title+Abstract	3.6 hrs/ epoch	10 GB
+Body	5.5 hrs/ epoch	30 GB

Table 8: Time and Memory consumption using the conventional method (T+A) and the new methods (T+A+Body) where Body could be sentences retrieved using any of the methods like Summary, Citation, Non-Citation, Random, Ret-Aug, with total sequence length 800.

A Hyperparamters

For the catSeq model, we use the same hyperparameters as Chan et al. (2019). For ExHiRD and One2Set, we use the respective repositories with their respective hyperparameters. For the LED model, we did a hyperparameter sweep over the different learning rates $\in \{1e-3, 1e-4, 1e-5, 2e-5, 5e-5, 8e-5\}$ and finally chose the learning rate as $5e-5$, optimizer as Adam (Kingma and Ba, 2015), train batch size as 12. For preprocessing the articles to generate summary using PACSUM (TF-IDF), we use the default hyperparameters $\beta = 0$, $\lambda_1 = 0$, $\lambda_2 = 1$. We ran each experiment thrice with different random seeds using V100 (16GB) and A6000 (48GB) GPU machines.

B Time and Memory Consumption

In Table 8, we provide the comparison of time and memory utilization between the conventional (i.e., Title+Abstract) and new methods (i.e., Title+Abstract+Body) in reference to our best performing model (i.e., LED). Body refers to sentences retrieved using any of the methods described in §3, i.e., Random, Citation, Non-Citation, Summary, Retrieval-Augmentation.

The conventional method takes about 3.6 hours/epoch and 10 GB of CUDA memory with A6000 GPU for training. Whereas the newer methods take 5.5 hours/epoch and about 30 GB of CUDA memory for training. The inference time is less than 1 GPU hour for both types of methods.

C PacSum Algorithm

D Retrieval-Augmentation Algorithm

Algorithm 1 Summarization: PacSum (TF-IDF)

Input Φ : Document split into a list of sentences $(s_i)_{i \in [1, N]}$ (tokenized into list of words); hyperparameters: β ; λ_1 ; λ_2

Output \mathbb{S} : Summary of document Φ

- 1: $N \leftarrow$ Number of sentences in document
- 2: IDF $\leftarrow \left[\log(N - \text{count}(w, \Phi) + 0.5) - \log(\text{count}(w, \Phi) + 0.5) | w \in s_i; s_i \in \Phi \right]$
- 3: TF $\leftarrow \left[\text{count}(w, s_i) | w \in s_i; s_i \in \Phi \right]$
- 4: $\forall i, j \in [1, N], \text{sim}(i, j) \leftarrow \begin{cases} \mathbf{1} & i == j \\ \sum_{w \in s_i \& s_j} (\text{TF}[s_i][w] * \text{TF}[s_j][w] * \text{IDF}[w])^2 & i != j \end{cases}$
- 5: $\tau = \min(\text{sim}) + \beta * (\max(\text{sim}) - \min(\text{sim}))$
- 6: $\forall i, j \in [1, N], \text{sim}(i, j) = \text{sim}(i, j) - \tau$
- 7: $\text{FS}(i) \leftarrow \begin{cases} - \sum_{\forall j, j < i} \text{sim}(i, j) & \text{sim}(i, j) > \tau \\ 0 & \text{otherwise} \end{cases}$
- 8: $\text{BS}(i) \leftarrow \begin{cases} \sum_{\forall j, j > i} \text{sim}(i, j) & \text{sim}(i, j) > \tau \\ 0 & \text{otherwise} \end{cases}$
- 9: $\text{centrality}(i) \leftarrow \lambda_1 * \text{FS}(i) + \lambda_2 * \text{BS}(i), \forall i \in [1, N]$
- 10: $\text{indices} \leftarrow \text{argsort}(\text{centrality}, \text{desc}=\text{True})$
- 11: $\mathbb{S} \leftarrow \Phi[\text{indices}]$
- 12: Return \mathbb{S}

Algorithm 2 Retrieval-Augmentation

Input Collection of scientific articles \mathbb{C}

Output Collection of scientific articles \mathbb{D} with Title \oplus Abstract \oplus Semantically similar sentences

- 1: $\Phi \leftarrow \bigcup_{a \in \mathbb{C}_{\text{train}}} \left\{ \bigcup_{s_j \in (T_a \oplus A_a)} s_j \right\} \triangleright T_a$: Title, A_a : Abstract, Φ : Set of sentences from titles and abstracts of all articles in training dataset
- 2: $\mathbb{D} \leftarrow \emptyset$
- 3: **for all** $a \in \mathbb{C}$ **do**
- 4: $\Phi_q \leftarrow \bigcup_{s_j \in (T_a \oplus A_a)} s_j$
- 5: $\Phi_{\text{key}} \leftarrow \Phi - \Phi_q$
- 6: $E_{\text{key}} \leftarrow (\text{embed}(s_i))_{s_i \in \Phi_{\text{key}}}$
- 7: $\vec{e}_q \leftarrow \text{embed}(T_a \oplus A_a)$
- 8: $\text{similarity} \leftarrow (\vec{e}_{\text{key}} \cdot \vec{e}_q)_{\vec{e}_{\text{key}} \in E_{\text{key}}}$
- 9: $\text{indices} \leftarrow \text{argmax}_k(\text{similarity})$
- 10: $\text{extra_info} \leftarrow (\Phi_{\text{key}}[i])_{i \in \text{indices}}$
- 11: $a \leftarrow T_a \oplus A_a \oplus \text{extra_info}$
- 12: $\mathbb{D} \leftarrow \mathbb{D} \cup \{a\}$
- 13: **end for**
- 14: Return \mathbb{D}