
Open4Business (O4B): An Open Access Dataset for Summarizing Business Documents

Amanpreet Singh*

Department of Computer Science
Stony Brook University
Stony Brook, NY 11790
amanpsingh@cs.stonybrook.edu

Niranjan Balasubramanian

Department of Computer Science
Stony Brook University
Stony Brook, NY 11790
niranjan@cs.stonybrook.edu

Abstract

A major challenge in fine-tuning deep learning models for automatic summarization is the need for large domain specific datasets. One of the barriers to curating such data from resources like online publications is navigating the license regulations applicable to their re-use, especially for commercial purposes. As a result, despite the availability of several business journals there are no large scale datasets for summarizing business documents. In this work, we introduce Open4Business (O4B), a dataset of 17,458 open access business articles and their reference summaries. The dataset introduces a new challenge for summarization in the business domain, requiring highly abstractive and more concise summaries as compared to other existing datasets. Additionally, we evaluate existing models on it and consequently show that models trained on O4B and a 7x larger non-open access dataset achieve comparable performance on summarization. We release the dataset, along with the code² which can be leveraged to similarly gather data for multiple domains.

1 Introduction

Generating effective summaries of documents is critical in business and corporate finance. Successful transactions in these domains often require timely processing of hundreds of long, complex and diverse documents. For instance, the due-diligence process in mergers and acquisitions requires the involved parties to go through tons of paperwork [1]. Concise summaries can save time and effort spent on reading entire documents.

However, it is challenging to obtain in-domain training data for fine-tuning automatic summarization systems. The documents are often proprietary and confidential or the access to the data itself could be ephemeral. Even with the necessary access, the effort needed for an expert to create reference summaries would be enormous. One way to circumvent these issues is to gather publicly accessible data such as articles from business journals. While most journals are published under licenses that prohibit their commercial use, in recent years there has been a steady increase in open access journals which have more lenient license terms with regards to copyrights and their re-use, even commercially.

This work makes the following contributions to further research and quick prototyping of summarization systems for business documents:

- i) O4B - a dataset consisting of 17,458 open access business articles with reference summaries.
- ii) Evaluation of current SOTA abstractive models on O4B. We also compare it extensively with other summarization datasets and another huge dataset which is not open access.

*Work done as an intern for SS&C Intralinks.

²<https://github.com/amanpreet692/0pen4Business>

iii) An open source implementation which can be used for dataset curation in other domains as well.

2 Related work

Many widely-used summarization datasets mainly cover news articles [2–4], with the exception of some that cover specific domains such as legislation [5] and intellectual property [6]. Cohan et al. [7] previously used ArXiv and PubMed research papers to generate abstracts of long documents. In the financial domain, El-Haj [8] released the Annual Reports of around 4000 companies listed on London Stock Exchange as part of a MultiLing 2019 shared task. Since Annual Reports are generally over 60 pages in length, the task was geared towards generating longer summaries of about 1000 words, which is computationally expensive for the current attention-based models.

Existing literature that align most closely with O4B, in terms of content and the curation process are Ammar et al. [9] and S2ORC [10]. Both consist of a huge volume of article records in multiple domains including business and finance which can be useful for several tasks. Even though the former allows commercial use, it doesn't have full text of the articles. Contrastingly, S2ORC has full text but prohibits commercial re-use as it was released under the CC BY-NC license (See Appendix A.1 for more details). Moreover S2ORC has not been used for benchmarking in summarization as of this writing. We attempt to overcome these shortcomings by providing the article abstract and full text without restrictions on re-use as part of O4B. We also benchmark O4B and compare it to a subset of S2ORC consisting of only business articles, referred to as S2ORCB in the evaluation section.

3 Open4Business dataset

We introduce a new summarization dataset, Open4Business (O4B), by parsing open access journal articles from the business domain to obtain their full text and abstract. Only open source tools are used for gathering the data to encourage re-use, and to maintain consistency with the open access theme of this work. Figure 1 shows a sample summary from O4B.

3.1 Dataset curation

O4B is created from a collection of business articles obtained from multiple publishers. We filtered this collection to only those articles which are open access, retrieved and converted their full-text PDF into a structured XML format, and extracted the abstracts as reference summaries. We detail each of these steps below and outline the number of data samples through these stages in Appendix A.2.

Selecting open access articles: To identify open access articles, we used the ISSN of business and financial journals obtained from ISSN-GOLD-OA records [11], as well as a short list of related keywords (e.g. business, finance, and entrepreneur) and cross-referenced those with Crossref [12], a large scale metadata repository for published articles. Using an existing implementation of Crossref API³, we queried the Crossref REST service to retrieve Document Object Identifier (DOI), license details and other article metadata. We process the resulting list to only include articles with open access licenses, specifically those of type 'CC BY (3.0/4.0)'.

Retrieving article PDFs: Given the collection of relevant DOIs, we used Unpaywall [13], another repository that provides the best open access resource locations for a given article. We queried it using Unpywall⁴ and then retrieved the PDF versions of the articles. Out of the 21,397 DOIs, we were able to retrieve PDFs for 18,670. In some cases, there was no DOI record or there were connectivity issues, while for others the DOIs were pointing to non-business journals.

XML conversion and post-processing: We used GROBID [14] to convert PDFs to a structured XML format. We have considered only English-language articles that have abstracts. For each article we retained only the main text body of the article and removed section headings, tables, figures, citations, and other bibliographic references. The resulting collection contains the text of 17,458 articles along with a combination of their titles and abstracts as the gold summaries.

³<https://github.com/fabiobatalha/crossrefapi>

⁴<https://github.com/unpywall/unpywall>

Table 1: Comparison of Open4Business (O4B) with other summarization datasets

Dataset	Total Docs	Mean Words Doc	Mean Words Summary	Compression Ratio
BillSum	22,218	1,240.21	181.98	6.81
Newsroom	1,145,804	684.51	29.17	23.46
CNN/DM	311,971	707.65	49.66	14.25
ArXiv	215,500	5219.54	239.27	21.81
O4B	17,458	4764.89	173.58	27.45

Source
As a mechanism of great concern, internal control has attracted many scholars to study it, which has led the academic and practical circles to study the quality of internal control. Many scholars use a variety of different research methods to evaluate the internal control of enterprises. Doyle and Ashbaugh-Skaife found that the weaker the internal control of the company, the lower the accounting quality. Hoitash found that the severity and different nature of internal control deficiencies have different effects on audit fees. Kim found that companies with significant ...
Summary
Literature Review and Prospect of Enterprise Internal Control. This report systematically reviews the literature related to internal control from 1973 to 2017, and reviews the internal control literature from two aspects at home and abroad. Based on this, the future research of China's internal control is prospected.

Figure 1: Summary sample from O4B




Figure 2: Comparison of unique summary n-grams

3.2 Dataset characteristics

To get a measure of the conciseness and abtractiveness of the summaries in O4B, we compare it to a set of benchmark datasets [2, 4, 5, 7] from various domains. Following [6], we compare the reference summaries against their sources in terms of average length in words (to measure conciseness) and unique n-grams (to measure abtractiveness), respectively. As shown in Table 1, O4B has the highest compression ratio indicating highly compact summaries. The unique summary n-gram distributions shown in Figure 2 is comparable to ArXiv and CNN-DailyMail(CNN/DM) datasets, which are highly abtractive in nature. Also it has fewer extractive fragments than BillSum and Newsroom. As we show in the next section, our benchmarking experiments indicate that pre-training or fine-tuning on other datasets is not enough for obtaining high scores on O4B, highlighting the need for this dataset.

4 Evaluation

Since the focus of this work is solely on abtractive summarization, SOTA encoder-decoder transformer based models [15] are used for benchmarking O4B. PyTorch 1.5 implementations of Text-to-Text Transfer Transformer (T5) [16] and a distilled version of BART [17] (dBART) provided by [18] are used for evaluation. Off-the-shelf (base) variants of both the models are fine-tuned on a single 16 GB GPU setup for 2 epochs with a batch size of 2 for consistent comparisons⁵. For all benchmarking experiments, both the models are fine-tuned on 512 and 1024 tokens each from the source text.

4.1 ROUGE-N benchmarking

The O4B benchmark ROUGE-(1/2/L) scores are shown in Table 2. Benchmark scores from both the models without any fine-tuning are considered as baselines. During fine-tuning, the models with the best ROUGE-2 score on the dev set are saved as checkpoints. Both the models perform better when they witness a longer portion of the input. But, this comes at cost of longer training times, which increases proportionally with the input length. Overall, the low performance numbers for T5 and dBART - specially for ROUGE-2, indicates that the task of long document summarization is

⁵Fine-tuning larger models for longer will likely yield better scores but here we only present basic benchmarking results.

Table 2: ROUGE-N Metrics, RG: ROUGE, O4B: Open4Business

Architecture	Tokens	RG-1	RG-2	RG-L
T5-base	512	24.30	7.21	16.33
T5-O4B	512	42.50	14.88	24.62
	1024	45.34	17.39	26.31
dBART-base	512	28.46	8.10	17.89
dBART-O4B	512	43.55	15.95	25.25
	1024	45.98	17.90	26.59

Table 3: Transfer learning experiment results, RG: ROUGE

(a) Performance on S2ORCB test set				(b) Performance on O4B test set			
Model	RG-1	RG-2	RG-L	Model	RG-1	RG-2	RG-L
T5-O4B	45.52	17.85	26.83	T5-O4B	45.34	17.39	26.31
T5-S2ORCB	45.59	18.17	27.32	T5-S2ORCB	44.04	16.58	25.67
dBART-O4B	46.11	18.26	27.04	dBART-O4B	45.98	17.90	26.59
dBART-S2ORCB	47.36	19.47	28.26	dBART-S2ORCB	45.72	17.79	26.35
Mean Difference	0.66	0.60	0.855	Mean Difference	0.78	0.46	0.44

challenging by itself, which is consistent with the findings in [7], [19] on the ArXiv and PubMed datasets. Sample summaries generated by both the models are presented in Appendix A.4.

4.1.1 Error analysis

We observe two main types of errors: word-level repetition, and hallucinating content that is not present in or supported by the input. Consider the sample: *The results show that a mixture of global and local smoothing improves non-parametric estimation, while the local parts allow for the desired flexibility to detect local features.* The source article is about financial data-modeling with sparse mentions of non-parametric estimation but the model combines this information by lifting some text from the source and ignoring other essential parts. Here, the original text is - *Those mixtures are interesting because the global parts can borrow the strength from a larger sample and have a smoothing effect, while the local parts allow for the desired flexibility to detect local features.*

4.2 Transfer-learning tests with S2ORCB

To validate the effectiveness of both O4B and open access articles, we conduct transfer-learning experiments. First, we fine-tune T5 and dBART models on O4B and test on the S2ORCB dataset, a large collection of over 119k summary/article pairs. S2ORCB is described in detail in Appendix A.3. Here, we consider 1024 source text tokens since that resulted in better benchmark scores. Given the datasets, a sizeable difference is observed in the training times as well. Results in Table 3a show that the models fine-tuned on O4B (T5-O4B and dBART-O4B) achieve comparable performance to corresponding models fine-tuned on the S2ORCB dataset. Second, we reverse the setup and fine-tune models on the training portions of S2ORCB and then test on O4B. Table 3b indicates a similar gap in performance between the transfer models versus the ones trained on the target dataset (O4B in this case). The gaps hint at the challenges in the summarization task on the O4B dataset that cannot be solved by simply training on a larger dataset. In spite of the huge disparity in dataset sizes, similar model performances indicate that O4B can be effectively utilized for summarization systems.

5 Conclusions

Automatic summarization in business domain is hampered by the lack of access to large scale datasets. Thus, we introduce Open4Business(O4B), a new dataset that consists of open access articles and their abstracts. This presents a new challenge for summarization systems which we hope spurs further research and development of new methods for this domain.

Acknowledgments

We would like to thank Prakash Kanchinadam and also the legal team at SS&C Intralinks for the fruitful discussions and their administrative support in this endeavor. We thank all anonymous reviewers for their feedback.

References

- [1] Peter Howson. *Due diligence: The critical stage in mergers and acquisitions*, pages 151–160. Gower Publishing, Ltd., 2003.
- [2] Abigail See, Peter J Liu, and Christopher D Manning. Get to the point: Summarization with pointer-generator networks. *arXiv preprint arXiv:1704.04368*, 2017.
- [3] Shashi Narayan, Shay B Cohen, and Mirella Lapata. Don’t give me the details, just the summary! *Topic-aware Convolutional Neural Networks for Extreme Summarization*. In, 2018.
- [4] Max Grusky, Mor Naaman, and Yoav Artzi. Newsroom: A dataset of 1.3 million summaries with diverse extractive strategies. *arXiv preprint arXiv:1804.11283*, 2018.
- [5] Vladimir Eidelman. Billsun: A corpus for automatic summarization of us legislation. In *Proceedings of the 2nd Workshop on New Frontiers in Summarization*, pages 48–56, 2019.
- [6] Eva Sharma, Chen Li, and Lu Wang. Bigpatent: A large-scale dataset for abstractive and coherent summarization. *arXiv preprint arXiv:1906.03741*, 2019.
- [7] Arman Cohan, Franck Dernoncourt, Doo Soon Kim, Trung Bui, Seokhwan Kim, Walter Chang, and Nazli Goharian. A discourse-aware attention model for abstractive summarization of long documents. *arXiv preprint arXiv:1804.05685*, 2018.
- [8] Mahmoud El-Haj. Multiling 2019: Financial narrative summarisation. In *Proceedings of the Workshop MultiLing 2019: Summarization Across Languages, Genres and Sources*, pages 6–10, 2019.
- [9] Waleed Ammar, Dirk Groeneveld, Chandra Bhagavatula, Iz Beltagy, Miles Crawford, Doug Downey, Jason Dunkelberger, Ahmed Elgohary, Sergey Feldman, Vu A. Ha, Rodney Michael Kinney, Sebastian Kohlmeier, Kyle Lo, Tyler C. Murray, Hsu-Han Ooi, Matthew E. Peters, Joanna L. Power, Sam Skjonsberg, Lucy Lu Wang, Christopher Wilhelm, Zheng Yuan, Madeleine van Zuylen, and O. Etzioni. Construction of the literature graph in semantic scholar. In *NAACL-HLT*, 2018.
- [10] Kyle Lo, Lucy Lu Wang, Mark Neumann, Rodney Kinney, and Daniel S. Weld. S2ORC: The Semantic Scholar Open Research Corpus. In *Proceedings of ACL, 2020*. URL <https://www.aclweb.org/anthology/2020.acl-main.447>.
- [11] A Bruns, C Lenke, C Schmidt, and NC Taubert. Issn-matching of gold oa journals (issn-gold-oa) 3.0. bielefeld university, 2019.
- [12] Crossref. <https://www.crossref.org/>, 2020.
- [13] Heather Piwowar, Jason Priem, Vincent Larivière, Juan Pablo Alperin, Lisa Matthias, Bree Norlander, Ashley Farley, Jevin West, and Stefanie Haustein. The state of oa: a large-scale analysis of the prevalence and impact of open access articles. *PeerJ*, 6:e4375, 2018.
- [14] Grobid. <https://github.com/kermitt2/grobid>, 2008–2020.
- [15] Alexander R Fabbri, Wojciech Kryściński, Bryan McCann, Caiming Xiong, Richard Socher, and Dragomir Radev. Summeval: Re-evaluating summarization evaluation. *arXiv preprint arXiv:2007.12626*, 2020.
- [16] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140):1–67, 2020.

- [17] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*, 2019.
- [18] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. Huggingface’s transformers: State-of-the-art natural language processing. *ArXiv*, abs/1910.03771, 2019.
- [19] Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter J Liu. Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. *arXiv preprint arXiv:1912.08777*, 2019.

A Appendices

A.1 Background

Traditionally, most of the journals have published under licenses which restrict their re-use in some manner eg. CC BY-NC⁶, or prohibit mining them. The Open Access movement, which has gained popularity in recent years, allows peer-reviewed articles to be made available immediately for public without any copyright restrictions. Several levels of open access publishing exist with each having its own set of regulations. The readers are encouraged to read about these as it is essential to be aware of any potential restrictions on re-use of what seems to be public data. Consequently, only the the Gold Open Access level articles licensed as CC BY⁷ have been used to create this Open4Business(O4B).

A.2 O4B through curation stages

O4B is built from open access business journals published by publishers such as Scientific Research Publishing (SCIRP)⁸ and Multidisciplinary Digital Publishing Institute (MDPI)⁹. Through every step in the dataset curation process as explained in the main paper, several potential data samples are filtered out due to processing issues as shown in Table 4. After final post-processing, remaining articles were split 80/10/10 to have a training, dev and test set of 13,966, 1,746 and 1,746 articles respectively.

Table 4: Data samples throughout pipeline stages

Step	Number of samples
License Metadata Filtering (Crossref)	21,397
Open Access URL Retrieval (Unpaywall)	20,432
PDF Mining	18,670
XML Conversion (GROBID)	18,544
Final Post-processing	17,458

A.3 Deriving S2ORCB from S2ORC

S2ORCB is derived from S2ORC by including only those articles that belong to the business domain. S2ORC has the metadata for about 1.8 million business articles but full text is not available for all of them. Thus, the records were then filtered to include only those which had both the article abstracts and full body text. After parsing the text, it is subject to similar post-processing as O4B and any bibliographical references are subsequently removed. Finally, S2ORCB is obtained with article full text as source, and a combination of title and abstract as the reference summary similar to O4B.

⁶Creative Commons-NonCommercial.

⁷<https://creativecommons.org/licenses/by/4.0/>

⁸<https://www.scirp.org/>

⁹<https://www.mdpi.com/>

S2ORCB is split with 95926, 11991 and 11991 articles in the train, dev and test set respectively. Table 5 shows the mean word count in source and summary for each of the splits in S2ORCB.

Table 5: S2ORC-Business(S2ORCB) Statistics

Split	Description	Mean word count
Train	Source	4611
	Summary	186
Dev	Source	4557
	Summary	187
Test	Source	4606
	Summary	186

A.4 Sample generated summary

We present sample generated summaries from both T5 and dBART along with the gold reference summary in Figure 3. As highlighted in the figure, following instances of grammatical inconsistencies are observed in the generated summaries starting from the top:

- i) inconsistent spelling of the term *polarization*
- ii) global *bipolarization* is explained by the *polarization*
- iii) *We conclude by the main conclusions* seems incoherent
- iv) *evolution* is **not** followed by *patterns* as per the source text

<p>Gold Summary Empirical Evidence of Economic Bipolarization in Africa. This report examines the degree of polarization in African countries' per capita GDP distribution between 1966 to 2004. We first use a nonparametric analysis and find that the countries tend to cluster in two classes of per capita GDP. Secondly, by using the Wolfson's bipolarization measure, the results reveal that bipolarization has been accelerating during the two first decades and is still growing. We relate the evolution of polarization during the period to the business sectors. We find that the specialization of the countries is the main factor explaining its evolution, namely, in agriculture and industry sectors.</p>
<p>T5 The degree of economic polarization in Africa: a nonparametric analysis of the per capita GDP distribution in 34 African countries. We show that there exists a bipolarization of the economies by using three methodologies. We primarily estimate the distribution of the GDP per capita in 34 countries, aiming to identify multimodality of the distributions, during the period 1966-2004. But the methodology does not give precise measure of the observed changes in the degree of polarisation over time. For this purpose, we find how the global bipolarized is explained by the polarized into four sectors: agriculture, mining, industries, and services. We conclude by the main conclusions and some policy implications.</p>
<p>dBART Bipolarization of African Economies: A Nonparametric Analysis. We examine the degree of economic polarization in Africa, using the per capita GDP distribution during the period 1966-2004. We show that there exists a bipolarization of the economies by using three methodologies. The paper is based on data drawn from the World Bank Africa Database. Estimates are based on Epanechnikov kernel functions, and in each case, the smoothing parameter function is determined following Silverman [37, 48]. The results are robust to the kernel function. There appear different patterns in the evolution patterns of the African countries over time, and the evolution of Wolfson's bipolarization index is explained by the polarization into four sectors: agriculture, mining, industries, and services.</p>

Figure 3: Generated summary sample from T5 and dBART with the corresponding Gold (reference) summary. The portions in generated summary indicating grammatical incoherence are highlighted.