

PoinT-5: Pointer Network and T-5 based Financial Narrative Summarisation

Abhishek Singh

Samsung R&D Bangalore

abhishek.s.eee15@iitbhu.ac.in

Abstract

Companies provide annual reports to their shareholders at the end of the financial year that describes their operations and financial conditions. The average length of these reports is 80, and it may extend up to 250 pages long. In this paper, we propose our methodology PoinT-5 (the combination of Pointer Network and T-5 (Test-to-text transfer Transformer) algorithms) that we used in the Financial Narrative Summarisation (FNS) 2020 task. The proposed method uses Pointer networks to extract important narrative sentences from the report, and then T-5 is used to paraphrase extracted sentences into a concise yet informative sentence. We evaluate our method using ROUGE-N (1,2), L, and SU4. The proposed method achieves the highest precision scores in all the metrics and highest F1 scores in ROUGE 1, and LCS and only solution to cross MUSE solution baseline in ROUGE-LCS metrics.

1 Introduction

Annual Reports may extend up to 250 pages long as stated above, which contains different sections General Corporate Information, financial and operating cost, CEOs message, Narrative texts, accounting policies, Financial statement including balance sheet and summary of financial data documents. In the Financial narrative summarisation task, only the narrative section is summarised, which is not explicitly marked in the dataset, making it challenging and interesting. In recent years, previous manual small-scale research in the Accounting and Finance literature has been scaled up with the aid of NLP and ML methods, for example, to examine approaches to retrieving structured content from financial reports, and to study the causes and consequences of corporate disclosure and financial reporting outcomes (El-Haj et al., 2018).

Companies produce glossy brochures of annual reports with a much looser structure, and this makes automatic summarisation of narratives in UK annual reports a challenging task (El-Haj et al., 2020). Hence we summarize the narrative section of annual reports, particular narrative sentences that are spread loosely across the document need to be first identified and summarise those sentences. The summarisation limit is set to 1000 words, where the actual length of the report may go up to 250 pages long. Hence to summarize these long annual reports using a combination of extractive and abstractive summarisation.

The text summary method can be classified into two paradigms: extractive and abstractive. The extractive summarisation method extracts the meaningful sentences or a section of text from the original text and combines them (ranked or unranked) to form a summary (Cheng and Lapata, 2016; Narayan et al., 2018; Yasunaga et al., 2017; See et al., 2017). Whereas abstractive summarisation generates words and sentences that are similar in meaning to the given text to form a summary that may not be in actual text (Nallapati et al., 2016; Rush et al., 2015; Paulus et al., 2017; Li et al., 2018). When summarizing long documents such as in our case up to 250 pages long, extractive summarisation may not produce a coherent and readable summary, and abstractive summarisation cannot cover complete information using encoder-decoder architecture. One problem is that typical seq2seq frameworks often generate unnatural summaries consisting of repeated words or phrases (Li et al., 2018). Hence, we come up with a combination of extractive and abstractive summarisation to first select important narrative sentences and concisely convey them.

Pointer Networks (Vinyals et al., 2015) is used in various combinatorial optimization problems, such as Travelling Salesman Problem (TSP), Convex hull optimization. We used pointer networks in our task of financial narrative summarization to extract relevant narrative sentences in a particular order to have a logical flow in summary. These extracted sentences are paraphrased to summarise these sentences in an abstractive way using the T-5 sequence-to-sequence model. We train the complete model by optimizing the ROUGE-LCS evaluation metric through a reinforcement learning objective.

2 Related Works

In this section and Appendix B, we discuss related works in the fields of abstractive summarisation, extractive summarisation, combinations of these two methods, reinforcement learning applications, and summarisation of Financial Narratives and their methodology. Studies of human summarizers show that it is common to apply various operations while condensing, such as paraphrasing, generalization, sentence-level summarisation and reordering (Jing, 2002). We continue related work in Appendix B.

3 Data Description

The financial narrative summarisation dataset contains 3,863 annual reports for firms listed on LSE covering the period between 2002 and 2017 (El-Haj et al., 2019; El-Haj et al., 2014). Dataset is randomly split into training (75%), testing and validation (25%) Table:2. We use NLTK sentence tokenizer ¹ to tokenize sentence in the annual report and summary processing for all our experiments. Data is further described and analyzed in Appendix A.

4 Methodology

4.1 Model Description

Our model is composed of three parts that are first trained or executed individually and then finally brought together using policy gradient algorithm. As stated earlier the model is combination of extractive as well as abstractive methods. Initially dataset is provided in the form $\{x_i, y_i^j\}$, $1 \leq i \leq N$, $1 \leq j \leq 7$. Here i represents number of annual reports, j represents number of summaries for each annual reports and x, y represent report and summary respectively.

In extraction process we assume that for every summary sentence there is matching sentence in the annual report. To train exaction model we need these corresponding sentences in the reports. Since, annual reports are not marked explicitly with sentences we followed ROUGE scores to extract these sentences as done in (Chen and Bansal, 2018; Nallapati et al., 2016). For every summary sentence we calculate ROUGE with every sentence in the report and then choose the sentence with maximum value.

$$j_t = \operatorname{argmax}_i (\text{ROUGE-L}_{\text{recall}}(d_i, s_t)) \quad (1)$$

In equation 1 j_t represent sentence with maximum ROUGE score for summary sentence s_t for every sentence in report d_i . For every report in training set there are multiple summaries. We calculate *ROUGE – Lsummary* for extracted sentences as mentioned in (Lin, 2004).

$$R_{lcs} = \frac{\sum_{i=1}^n LCS_U(r_i, c_i)}{m} \quad (2)$$

In equation 2 R_{lcs} represents ROUGE-L recall, r_i, c_i are report and summary sentences respectively and m is total number of words in extracted report sentences. Once summary level ROUGE-L recall is calculated we choose summary with maximum value for further processing. Once proxy sentences and a summary is selected applying above methods extraction model is trained. In extraction model sentence level representation of report sentences is calculated using hierarchical word to sentence level Bidirectional Long Short Term Memory (Bi-LSTM)(Hochreiter and Schmidhuber, 1997). First Bi-LSTM is

This work is licensed under a Creative Commons Attribution 4.0 International License. License details: <http://creativecommons.org/licenses/by/4.0/>.

¹<https://www.nltk.org/api/nltk.tokenize.html>

5 Results and Analysis

In this section we present results from our experiments and compare with different baselines MUSE (Litvak et al., 2010), Text-rank (Mihalcea and Tarau, 2004), Lex-Rank (Erkan and Radev, 2004), and Polynomial Summarisation (Litvak and Vanetik, 2013). We train three models in our experiments (PoinT-5) Pointer Network with T-5, Pointer Network alone, and Pre-trained Bert for text summarisation. In table

Metrics	Text-Rank	Lex-Rank	Polynomial	MUSE	Pre-trained Bert	Pointer Net	PoinT-5
Prec(R-L)	0.235	0.210	0.260	0.470	0.213	0.603*	0.605
Recall(R-L)	0.197	0.263	0.177	0.370	0.254	0.377	0.377
F-1(R-L)	0.206	0.218	0.205	0.407	0.225	0.455*	0.456
Prec(R-1)	0.414	0.337	0.324	0.483	0.241	0.611*	0.612
Recall (R-1)	0.118	0.269	0.253	0.413	0.378	0.392	0.393
F-1 (R-1)	0.172	0.264	0.274	0.433	0.283	0.465*	0.466
Prec (R-2)	0.229	0.193	0.147	0.311	0.114	0.448*	0.451
Recall (R-2)	0.044	0.107	0.088	0.198	0.138	0.220	0.222
F-1 (R-2)	0.070	0.120	0.105	0.234	0.118	0.289	0.289
Prec(R-SU4)	0.302	0.253	0.213	0.375	0.165	0.506*	0.508
Recall(R-SU4)	0.048	0.117	0.105	0.201	0.149	0.208	0.209
F-1(R-SU4)	0.079	0.140	0.135	0.253	0.149	0.286	0.288

Table 1: ROUGE Evaluation on Financial Narrative Summarisation data (Bold represents highest overall, Prec represents Precision and * represents second-highest overall)

1 bold is marked for highest value amongst all the solutions for the task including baselines. Pre-trained Bert for summarisation is not fine-tuned for this specific task hence its performance is not as good as Pointer Network and PoinT-5 (Pointer Network + T-5). PoinT-5 network gives highest results for precision on all the evaluation metrics ROUGE -1,2,L,SU4. From this it can be inferred that generated summaries are highly precise in extracting narrative sentences and matches with ground truth summaries. Recall is comparatively much lower than precision in most of the evaluated metrics which means that generated summaries do not cover all the information in ground-truth summaries. This large difference in precision and recall shows that generated summaries do not cover all the sentences possibly due to restriction imposed on the number of sentences during training to follow word limit in summaries. Less recall and high precision value shows that generated summaries provide highly relevant information but does not cover complete information in ground truth summaries. From results it is evident that there is not much difference in performance between pointer net performance and PoinT-5 which extraction played key role in the architecture. PoinT-5 and Pointer Networks are the only system to cross MUSE baseline in ROUGE-L by atleast 5 points. These models give highest F1 results in ROUGE-L and ROUGE_1 metrics and highest precision in ROUGE-L,1,2,SU4.

6 Conclusion and Future work

In this work we present our solution on Financial Narrative Summarisation(FNS2020) dataset using PoinT-5 method explained in 4. It is combination of both extractive and abstractive methods using Pointer Network and T-5. With these methods we are able to achieve highest precision score in every evaluation metric and achieve highest F-1 scores in ROUGE-LCS and ROUGE-1.

In our future work we would like to address several limitation of our method such as factual correctness in summaries which is very important in financial domain as done in (Zhang et al., 2019) in summarizing radiology reports. To improve precision of our generated summaries under 1000 words we would formulate a penalty if system generates more than 1000 words during training of RL algorithm rather than restricting algorithm to fixed number of sentences.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Eddy Cardinaels, Stephan Hollander, and Brian J White. 2018. Automatic summaries of earnings releases: Attributes and effects on investors' judgments. *Available at SSRN 2904384*.
- Yen-Chun Chen and Mohit Bansal. 2018. Fast abstractive summarization with reinforce-selected sentence rewriting. *arXiv preprint arXiv:1805.11080*.
- Jianpeng Cheng and Mirella Lapata. 2016. Neural summarization by extracting sentences and words. *arXiv preprint arXiv:1603.07252*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Mahmoud El-Haj, Paul Rayson, Steven Young, and Martin Walker. 2014. Detecting document structure in a very large corpus of uk financial reports.
- Mahmoud El-Haj, Paul Rayson, and Andrew Moore. 2018. The first financial narrative processing workshop (fnp 2018). In *Proceedings of the LREC 2018 Workshop*.
- Mahmoud El-Haj, Paul Rayson, Paulo Alves, Carlos Herrero-Zorita, and Steven Young. 2019. Multilingual financial narrative processing: Analysing annual reports in english, spanish and portuguese. *Multilingual Text Analysis: Challenges, Models, And Approaches*, page 441.
- Mahmoud El-Haj, Ahmed AbuRa'ed, Nikiforos Pittaras, and George Giannakopoulos. 2020. The Financial Narrative Summarisation Shared Task (FNS 2020). In *The 1st Joint Workshop on Financial Narrative Processing and MultiLing Financial Summarisation (FNP-FNS 2020, Barcelona, Spain)*.
- Günes Erkan and Dragomir R Radev. 2004. Lexrank: Graph-based lexical centrality as salience in text summarization. *Journal of artificial intelligence research*, 22:457–479.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Wan-Ting Hsu, Chieh-Kai Lin, Ming-Ying Lee, Kerui Min, Jing Tang, and Min Sun. 2018. A unified model for extractive and abstractive summarization using inconsistency loss. *arXiv preprint arXiv:1805.06266*.
- Hongyan Jing. 2002. Using hidden markov modeling to decompose human-written summaries. *Computational linguistics*, 28(4):527–543.
- Piji Li, Lidong Bing, and Wai Lam. 2018. Actor-critic based training framework for abstractive summarization. *arXiv preprint arXiv:1803.11070*.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Marina Litvak and Natalia Vanetik. 2013. Mining the gaps: Towards polynomial summarization. In *Proceedings of the Sixth International Joint Conference on Natural Language Processing*, pages 655–660.
- Marina Litvak, Mark Last, and Menahem Friedman. 2010. A new approach to improving multilingual summarization using a genetic algorithm. In *Proceedings of the 48th annual meeting of the association for computational linguistics*, pages 927–936.
- Rada Mihalcea and Paul Tarau. 2004. Textrank: Bringing order into text. In *Proceedings of the 2004 conference on empirical methods in natural language processing*, pages 404–411.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Ramesh Nallapati, Bowen Zhou, Caglar Gulcehre, Bing Xiang, et al. 2016. Abstractive text summarization using sequence-to-sequence rnns and beyond. *arXiv preprint arXiv:1602.06023*.
- Shashi Narayan, Shay B Cohen, and Mirella Lapata. 2018. Ranking sentences for extractive summarization with reinforcement learning. *arXiv preprint arXiv:1802.08636*.
- Romain Paulus, Caiming Xiong, and Richard Socher. 2017. A deep reinforced model for abstractive summarization. *arXiv preprint arXiv:1705.04304*.

Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2019. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*.

Alexander M Rush, Sumit Chopra, and Jason Weston. 2015. A neural attention model for abstractive sentence summarization. *arXiv preprint arXiv:1509.00685*.

Abigail See, Peter J Liu, and Christopher D Manning. 2017. Get to the point: Summarization with pointer-generator networks. *arXiv preprint arXiv:1704.04368*.

Zhaopeng Tu, Zhengdong Lu, Yang Liu, Xiaohua Liu, and Hang Li. 2016. Modeling coverage for neural machine translation. *arXiv preprint arXiv:1601.04811*.

Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. Pointer networks. In *Advances in neural information processing systems*, pages 2692–2700.

Qicai Wang, Peiyu Liu, Zhenfang Zhu, Hongxia Yin, Qiuyue Zhang, and Lindong Zhang. 2019. A text abstraction summary model based on bert word embedding and reinforcement learning. *Applied Sciences*, 9(21):4701.

Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256.

Michihiro Yasunaga, Rui Zhang, Kshitijh Meelu, Ayush Pareek, Krishnan Srinivasan, and Dragomir Radev. 2017. Graph-based neural multi-document summarization. *arXiv preprint arXiv:1706.06681*.

Yuhao Zhang, Derek Merck, Emily Bao Tsai, Christopher D Manning, and Curtis P Langlotz. 2019. Optimizing the factual correctness of a summary: A study of summarizing radiology reports. *arXiv preprint arXiv:1911.02541*.

Appendices

A Data Analysis

Table 2 presented total summaries and annual reports in train, validation, and test set. During our analysis, we found that most of the annual reports contain 100-200 sentences. There are 279 summaries with more than 500 sentences, which is large. Whereas in summaries, the average number of sentences is 50. Therefore summaries are one fourth on average of annual reports and up to one-tenth in many cases. We present these analysis in Figure 2.

Data Type	Training	Validation	Testing
Report full text	3,000	363	500
Gold Summaries	9,873	1,250	1673

Table 2: FNS 2020 Shared Task Dataset

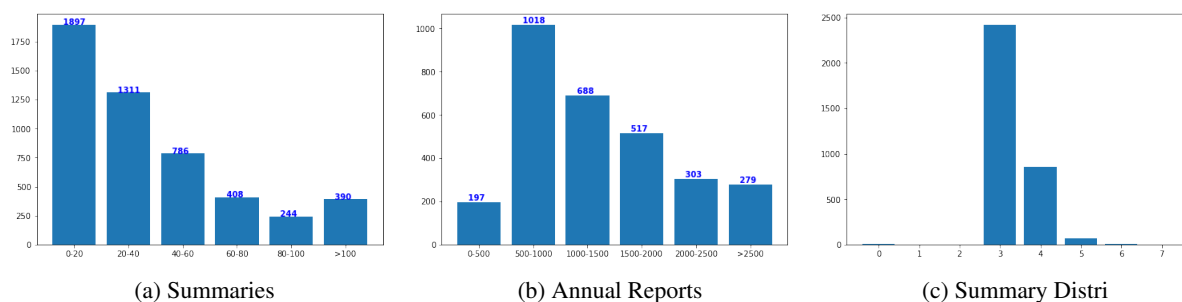


Figure 2: Sentence wise Distribution (a,b), Number of Summary Distribution (c)

B Extended Related Work

(Li et al., 2018) proposes a training framework based on the actor-critic model. They apply the attention-based sequence-to-sequence model as the actor to conduct summary generation. For the critic, they combine the maximum likelihood estimator with a well designed global summary quality estimator. (Nallapati et al., 2016) propose RNN based encoder-decoder model for abstractive summarisation. They apply bi-directional GRU-RNN at the encoder side and uni-directional GRU in decoder with attention. In their approach, each mini-batch’s decoder-vocabulary is restricted to words in the source documents of that batch. (Rush et al., 2015) also propose attention based sequence to sequence model for abstractive summarisation. (Paulus et al., 2017) states that attentional, RNN-based encoder-decoder models for abstractive summarisation have achieved good performance on short input and output sequences. For longer documents and summaries, however, these models often include repetitive and incoherent phrases. Hence they combine standard word prediction with the global sequence prediction training of RL which makes resulting summaries become more readable.

(Narayan et al., 2018) propose a reinforcement learning-based sentence ranking approach in extractive summarisation. (Cheng and Lapata, 2016) use attention architecture for words and sentence level extraction in extractive summarisation method. (Yasunaga et al., 2017) proposes a multi-document summarization system that exploits the representational power of deep neural networks and the sentence relation information encoded in graph representations of document clusters. Specifically, they apply Graph Convolutional Networks on sentence relation graphs.

(See et al., 2017) propose a novel “Pointer Generator” networks for abstractive summarisation. In this model a word is chosen with probability P_{gen} from overall vocabulary and $1 - P_{gen}$ from current sentence using attention weights. They apply coverage mechanism by (Tu et al., 2016) to avoid word repetitions in the summary which is common in large documents.

(Chen and Bansal, 2018) apply combination abstractive and extractive summarisation by join training using reinforcement learning. They apply pointer network for extraction (different from Pointer Generator) and RNN based encoder-decoder for abstraction. (Hsu et al., 2018) propose unified extractive and abstractive with a hierarchical sentence and word level attention model using novel inconsistency loss. (Wang et al., 2019) used T-5 (Raffel et al., 2019) based sentence representation for combined extractive and abstractive training.

Financial Narrative Summarisation has been explored in the past by (Cardinaels et al., 2018). They provide multiple evidence that algorithm-based summaries are less positively biased than management summaries.