

## SCE-SUMMARY at the FNS 2020 shared task

<b>Marina Litvak</b> Shamoon College of Engineering (SCE) Beer-Sheva Israel marinal@ac.sce.ac.il	<b>Natalia Vanetik</b> Shamoon College of Engineering (SCE) Beer-Sheva Israel natalyav@sce.ac.il	<b>Tzvi Puchinsky</b> Shamoon College of Engineering (SCE) Beer-Sheva Israel tzvipu@ac.sce.ac.il
---	---	---

### Abstract

With the constantly growing amount of information, the need arises to automatically summarize this written information. One of the challenges in the summary is that it's difficult to generalize. For example, summarizing a news article is very different from summarizing a financial earnings report. This paper reports an approach for summarizing financial texts, which are different from the documents from other domains at least in three parameters: length, structure, and format. Our approach considers these parameters, it is adapted to hierarchical structure of sections, document length, and special "language". The approach builds an hierarchical summary, visualized as a tree with summaries under different discourse topics. The approach was evaluated using extrinsic and intrinsic automated evaluations, which are reported in this paper. As all participants of the Financial Narrative Summarisation (FNS 2020) shared task, we used FNS2020 dataset for evaluations.

### 1 Introduction

The area of text summarization exists for several decades, since the first work of Luhn (Luhn, 1958). The summarization approaches developed from extractive unsupervised statistical approaches to abstractive supervised methods, using deep learning models (Liu, 2019). However, the most advanced seq2seq models (transformers) are very limited in input size and, therefore, are inapplicable to long texts. Also, only few of state-of-the-art summarizers consider hierarchical structure of the input documents (Yang and Wang, 2008; Zhang et al., 2019), their key concepts (Ouyang et al., 2009; Plaza et al., 2011) or topics (Wang et al., 2013; Akhtar, 2017) and build a hierarchical summary (Christensen et al., 2014; Akhtar et al., 2019). Usually, hierarchical summary is built per document collection. The top level of hierarchy provides a general overview and users can navigate the hierarchy to drill down for more details on topics of interest.

There is a growing interest in the application of automatic and computer-aided approaches for extracting, summarising, and analysing both qualitative and quantitative financial data, as a series of FNP and related workshops (El-Haj, 2019; El-Haj et al., 2018) recently demonstrates. However, summarization of documents in financial domain is usually limited to summarization of financial news (Filippova et al., 2009; Yang and Wang, 2003; de Oliveira et al., 2002; Baralis et al., 2016; Zhang et al., 2018) which are not very different from the general news in length and format. Only few attempts were made to summarize financial reports (Isonuma et al., 2017), which are different from the news articles in at least four parameters: length, structure, format, and lexicon.

This paper reports an approach for hierarchical summarization of financial reports. Financial annual reports in the data of Financial Narrative Summarisation (FNS 2020) shared task<sup>1</sup> (El-Haj et al., 2020) are long, have many sections, and are written in "financial" language using many special terms, numerical data, and tables. Our system for summarization and hierarchical visualization financial reports (named by SCE-SUMM) considers discourse and topic hierarchical structure and builds an hierarchical view of the summarized report with interactive user interface.

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

<sup>1</sup><http://wp.lancs.ac.uk/cfie/fns2020/>

## 2 The SCE-SUMM System

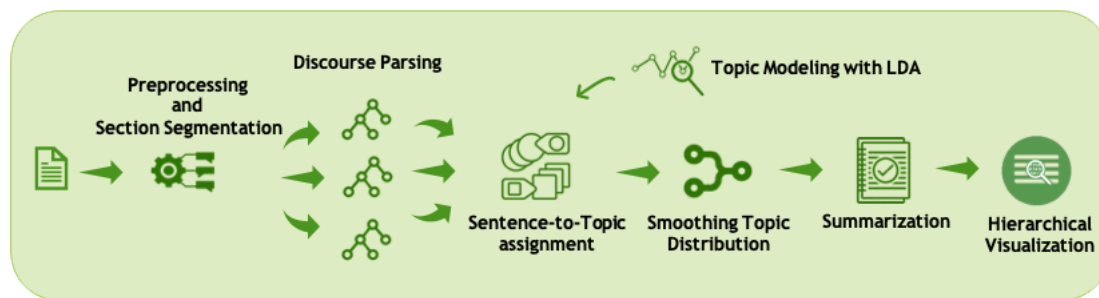


Figure 1: SCE-SUMM pipeline

SCE-SUMM utilizes two main methods: topic modeling (TM) and discourse parsing (DP). The pipeline of the proposed methodology is depicted in Figure 1 and includes the following steps:

**Text preprocessing**, that includes text cleaning, sentence splitting and tokenization. We developed our own tool that cleaned text before segmenting it to sentences and tokens. Financial reports usually contain a lot of sections, figures, and tables. Because the text files in the FNS-2020 dataset were obtained by converting pdf files to plain texts, these texts contain a lot of “noise” left from broken tables and meta-data such as section and page numbers. We cleaned the noise by measuring the ratio between text and numbers and ratio between number of words and whitespaces. Lines with ratio less than 0.4 were removed. Then, regular expressions were applied to find and mark such entities as URL, phone number, date, time, email. Finally, non-Unicode characters were filtered out.

**Section segmentation**, where section headers are identified and a document is segmented into sections. The section titles were extracted following the heuristic rules saying that (1) each title appears in a separate line, (2) does not end with period mark, and (3) contains only few (up to 5) words with (4) each word either starting with capital case letter or containing only upper case letters. The extracted candidates were then compared against the list of 13 manually edited titles<sup>2</sup>. The candidate that obtained Jaccard similarity above 0.4 to one of the titles from the list was extracted as a title. The text body between two consequent titles was marked as a section.

**Discourse parsing** of each section. For discourse parsing we used the CODRA parser (Joty et al., 2015). CODRA parser performs two-part process: (1) a discourse segmenter creates a segmentation analysis on the sentence level and EDU’s for the discourse parsing process and (2) a discourse parser parses the text on sentence level and document level to identify relations between parts of sentences and sentences in the document. The rhetorical analysis of the parser starts from a breaking a text into Elementary Discourse Units (EDUs). Because EDUs do not span across multiple sentences, this segmentation task finds EDUs inside the sentence boundaries. As a result, some sentences (actually, most, according to our observations) are split into EDUs. Every EDU is marked as a *nucleus* (an essence part) or a *satellite* (a complementary part of the related nucleus), based on the relation that they are connected to. Internal (relation) nodes represent different inter-sentence relations: elaboration, same-unit, etc.

**Topic modeling**. For topic modeling we applied Latent Dirichlet Allocation (LDA) model (Blei et al., 2003). It was applied on all files in the FNS-2020 dataset with predefined number of topics<sup>3</sup>.

**Topic-to-text assignment**, where each sentence (or sentence part) represented by a leaf node of the discourse tree, is assigned to one of the topics obtained by LDA. We refer topic probabilities  $p(t|w)$  for all sentence  $S$  words  $w \in S$  as their topic-related importance scores. Therefore, we extract a dominant

<sup>2</sup>Titles that appear in almost every report in FNS-2020 dataset, such as: ‘chairman statement’, ‘chief executive officer CEO review’, ‘chief executive officer CEO report’, ‘governance statement’, ‘remuneration report’, ‘business review’, ‘financial review’, ‘operating review’, ‘highlights’, ‘auditors report’, ‘risk management’, ‘chairman governance introduction’, ‘corporate social responsibility CSR disclosures’.

<sup>3</sup>We experimented with 4, 6, and 10 topics, and finally decided to keep 10 topics as best performing value.

topic ( $t \in T$ ) for each sentence  $S$ , as a topic with the maximal normalized sum of topic probabilities for all sentence words  $w \in S$ :  $\max_{t \in T} \frac{\sum_{w \in S} p(t|w)}{\sum_{w \in S} 1}$ .

**Topic distribution smoothing.** We noticed that after single text nodes (that stand for sentences or sentence parts) are assigned to topics, we can get unexpected topic distribution where two parts of the same sentence or two adjacent sentences inside the same paragraph and/or belonging to the same discourse relation are assigned to different topics, and transition from one topic to another is not coherent.<sup>4</sup> We decided to smooth topic distribution by extrapolating one dominant topic on entire block of adjacent sentences and sentence parts, connected by a direct discourse relation. We denote nodes with at least one leaf node as “simple” and all leaves in its sub-tree are finally assigned to one dominant topic, so that a “random” noise is left out. The implement this approach as follows. We know that all leaf nodes are arranged in the natural sequential order of their texts from right-to-left (top-down) in a discourse tree. We assume that the important information usually comes first (important part of a sentence usually precedes its complementary part, and a sentence stating some fact usually precedes a sentence that elaborates more about this fact) and, therefore, upper right nodes and nucleuses should propagate their topics on their siblings. According to this assumption and our empirical observations on each parameter’s influence, the final impact factor  $NI$  of node  $n$  is calculated as follows.  $NI(n) = \sum_{i=1}^3 w_i \times f_i(n)$ , where:

$f_1$  is a relative depth feature  $rd(n) = \frac{h(t)+1-d(n)}{h(t)+1}$ ,  $h(t)$  is a tree height,  $d(n)$  is  $n$ ’s depth

$f_2$  is a position feature  $pos(n) = \begin{cases} 1, & \text{if } n \text{ is on right} \\ 0, & \text{else} \end{cases}$

$f_3$  is a discourse label feature  $l(n) = \begin{cases} 1, & \text{if } n \text{ is nucleus} \\ 0, & \text{else} \end{cases}$

$w_1 = 0.5$ ,  $w_2 = 0.3$ , and  $w_3 = 0.2$ .

Then, the final dominant topic for a “simple” sub-tree is calculated as follows:  $\max_{t \in T} \{\sum_{n \in leaves} NI(n) * score_{t,n}\}$ . After topic-to-sentence assignment (at previous stage), every leaf node has non-zero value for only one dominant topic, other topics have  $score_{t,n} = 0$ .

**Summarization** of entire report (regardless visualization) and of each section (for visualization needs) was performed as follows. All topics  $t$  are ranked by their importance  $TI(t)$  (normalized sum of their probabilities for all document/section words). Then, summaries are created by extraction of nucleuses from each topic, in the topics’ importance order, until the maximum length limit is reached. As for entire report a summary should not exceed 1000 words according to the shared task instructions, we limit a section summary to 100 words.

**Hierarchical visualization.** At this stage SCE-SUMM creates an interactive html file with the data from all the stages for a user to browse. The file contains the following sections: (1) original text; (2) processed XML text after cleaning and section segmentation; (3) discourse trees for all the sections; (4) sentences (nodes) with assigned topics after smoothing; (5) the final hierarchical tree with the section summaries, and (6) a general report summary. For visualization and interactive user’s navigation, the following tree structure of a document is built and present to a user: root represents an entire document and points to its sections, each section is split to major topics inside this section after smoothing, and each topic points to a summary of this particular section focused on the chosen topic. Visualization is performed in interactive manner, upon a user’s request. Demo video<sup>5</sup> demonstrates all interactive options provided by the system.

### 3 Experiments

#### 3.1 Dataset

The Financial Narrative Summarisation (FNS 2020) shared task aims to demonstrate the value and challenges of applying automatic text summarisation to financial text written in English, usually referred to

<sup>4</sup>We assume that in a natural topic distribution, that is usually observed in general domains, topics must flow from one paragraph (or sections or cluster of sentences) to another, without mix of topics inside clusters.

<sup>5</sup>[https://drive.google.com/file/d/14qMRUHZIwaVoSltaLPSiH6NZx13M\\_9ue/view](https://drive.google.com/file/d/14qMRUHZIwaVoSltaLPSiH6NZx13M_9ue/view)

as financial narrative disclosures. The task dataset has been extracted from UK annual reports published in PDF file format. UK annual reports are lengthy documents with around 80 pages on average, some annual reports could span over more than 250 pages, while the summary length should not exceed 1000 words. The training set includes 3,000 annual reports, with 3-4 human-generated summaries as gold standard. For the evaluation process the test set of 500 files were provided. To address the time limitations and processing long files<sup>6</sup> the project reduced the length of the original files (to 15000 characters) to be able to process in feasible time limit (20 minutes per file at most).

### 3.2 Tools and runtime environment

For LDA, we used the Python gensim4 package. Corpus tf-idf vectorization and K-means clustering were performed by the Python sklearn package. For running Rouge, we used ROUGE 2.05 java package (Ganesan, 2018). Our approach was implemented in Python and run on Intel Pentium Gold G5400 with 16GB memory server with 40GB swap file configured.

### 3.3 Results

Automatic evaluation was performed using ROUGE metrics (Lin, 2004) which work by comparing an automatically produced summary against a set of reference summaries (typically human-produced). We applied three ROUGE metrics—ROUGE-1, ROUGE-2, and ROUGE-L. We compared our approach with two baseline methods—MUSE (Litvak et al., 2010) and POLY (Litvak and Vanetik, 2013). MUSE is a supervised approach based on a genetic algorithm, it was trained on 30 randomly selected gold standard summaries provided with FNS-2020 dataset. POLY is unsupervised approach based on linear programming, it was applied with Maximal Weighted Term Sum (OBJ1 in (Litvak and Vanetik, 2013)) objective function. Table 1 show the results, with recall, precision, and F-measure for each metric. The best scores are marked in bold and the second best are marked by grey background. It can be seen that SCE-SUMM performs better than POLY (both are unsupervised), and even outperforms MUSE (which is supervised) in one metric (ROUGE-L, Precision), meaning that its summaries are less “scattered” and more coherent (and therefore probably more readable) than other summaries. The comparative results with other systems participating in the FNS 2020 shared task can be seen in (El-Haj et al., 2020).<sup>7</sup>

system	R-1 R	R-1 P	R-1 F	R-2 R	R-2 P	R-2 F	R-L R	R-L P	R-L F
MUSE	<b>0.483</b>	<b>0.413</b>	<b>0.433</b>	<b>0.311</b>	<b>0.198</b>	<b>0.234</b>	<b>0.486</b>	0.381	<b>0.419</b>
POLY	0.324	0.253	0.274	0.147	0.088	0.105	0.270	0.182	0.212
SCE-SUMM	0.290	0.396	0.324	0.150	0.153	0.144	0.290	<b>0.396</b>	0.324

Table 1: Rouge results.

## 4 Conclusions and Future Work

This paper describes a new method for hierarchical summarization of financial reports, based on integrating the discourse structure and topic modeling. Our approach is mainly based on heuristic assumptions. In future, we intend to develop a new supervised method utilizing the discourse and topics knowledge. Also, we would like to apply this method and its extension to educational materials, which also have highly hierarchical structure and an evolving flow of topics in a discourse. Hierarchical summarization can help to organize those materials in a hierarchical structure and provide users with interactive navigation to the topics of interest. SCE-SUMM’s source code is available<sup>8</sup> and can be run using the provided instructions<sup>9</sup>.

<sup>6</sup>mostly, due to a very time-consuming discourse parsing

<sup>7</sup>Due to significant updates in the SCE-SUMM’s code since the task submission the scores might differ.

<sup>8</sup><https://github.com/Tzvi23/Hierarchical-Summarization-Part1>

<sup>9</sup><https://drive.google.com/drive/folders/1YxnNQ-9ebPX1Gtd6Dmr0to6UIBgudf1C>

## References

- Nadeem Akhtar, Hira Javed, and Tameem Ahmad. 2019. Hierarchical summarization of text documents using topic modeling and formal concept analysis. In *Data Management, Analytics and Innovation*, pages 21–33. Springer.
- Nadeem Akhtar. 2017. Hierarchical summarization of news tweets with twitter-lda. In *Applications of Soft Computing for the Web*, pages 83–98. Springer.
- Elena Baralis, Luca Cagliero, and Tania Cerquitelli. 2016. Supporting stock trading in multiple foreign markets: a multilingual news summarization approach. In *Proceedings of the Second International Workshop on Data Science for Macro-Modeling*, pages 1–6.
- David M Blei, Andrew Y Ng, and Michael I Jordan. 2003. Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan):993–1022.
- Janara Christensen, Stephen Soderland, Gagan Bansal, et al. 2014. Hierarchical summarization: Scaling up multi-document summarization. In *Proceedings of the 52nd annual meeting of the association for computational linguistics (volume 1: Long papers)*, pages 902–912.
- Paulo Cesar Fernandes de Oliveira, Khurshid Ahmad, and Lee Gillam. 2002. A financial news summarization system based on lexical cohesion. In *Proceedings of the International Conference on Terminology and Knowledge Engineering, Nancy, France*.
- 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, 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)*.
- Mahmoud El-Haj. 2019. Multiling 2019: Financial narrative summarisation. In *Proceedings of the Workshop MultiLing 2019: Summarization Across Languages, Genres and Sources*, pages 6–10.
- Katja Filippova, Mihai Surdeanu, Massimiliano Ciaramita, and Hugo Zaragoza. 2009. Company-oriented extractive summarization of financial news. In *Proceedings of the 12th Conference of the European Chapter of the ACL (EACL 2009)*, pages 246–254.
- Kavita Ganesan. 2018. Rouge 2.0: Updated and improved measures for evaluation of summarization tasks. *arXiv preprint arXiv:1803.01937*.
- Masaru Isonuma, Toru Fujino, Junichiro Mori, Yutaka Matsuo, and Ichiro Sakata. 2017. Extractive summarization using multi-task learning with document classification. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2101–2110.
- Shafiq Joty, Giuseppe Carenini, and Raymond T Ng. 2015. Codra: A novel discriminative framework for rhetorical analysis. *Computational Linguistics*, 41(3):385–435.
- 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.
- Yang Liu. 2019. Fine-tune bert for extractive summarization. *arXiv preprint arXiv:1903.10318*.
- Hans Peter Luhn. 1958. The automatic creation of literature abstracts. *IBM Journal of research and development*, 2(2):159–165.
- You Ouyang, Wenjie Li, and Qin Lu. 2009. An integrated multi-document summarization approach based on word hierarchical representation. In *Proceedings of the ACL-IJCNLP 2009 Conference Short Papers*, pages 113–116.
- Laura Plaza, Alberto Díaz, and Pablo Gervás. 2011. A semantic graph-based approach to biomedical summarization. *Artificial intelligence in medicine*, 53(1):1–14.

- Chi Wang, Xiao Yu, Yanen Li, Chengxiang Zhai, and Jiawei Han. 2013. Content coverage maximization on word networks for hierarchical topic summarization. In *Proceedings of the 22nd ACM international conference on Information & Knowledge Management*, pages 249–258.
- Christopher C Yang and Fu Lee Wang. 2003. Automatic summarization for financial news delivery on mobile devices. In *WWW (Posters)*.
- Christopher C Yang and Fu Lee Wang. 2008. Hierarchical summarization of large documents. *Journal of the American Society for Information Science and Technology*, 59(6):887–902.
- Yong Zhang, Erdan Chen, and Weidong Xiao. 2018. Extractive-abstractive summarization with pointer and coverage mechanism. In *Proceedings of 2018 International Conference on Big Data Technologies*, pages 69–74.
- Xingxing Zhang, Furu Wei, and Ming Zhou. 2019. Hibert: Document level pre-training of hierarchical bidirectional transformers for document summarization. *arXiv preprint arXiv:1905.06566*.