FLUEnT: Financial Language Understandability Enhancement Toolkit

Sohom Ghosh and Sudip Kumar Naskar

November 9, 2022

FLUEnT: Financial Language Understandability Enhancement Toolkit

Sohom Ghosh
Jadavpur University
Kolkata, West Bengal, India
sohom1ghosh@gmail.com

Sudip Kumar Naskar
Jadavpur University
Kolkata, West Bengal, India
sudip.naskar@gmail.com

ABSTRACT

Over the years, promising returns have enticed the masses to invest in the stock markets. However, most people do not have the financial knowledge needed for making investment decisions. Even seasoned investors find it difficult to grasp all the available information. This is primarily due to the ever-changing market dynamics and information overload. Natural Language Processing based automated systems are the rescue to such problems. In this paper, we present the Financial Language Understandability Enhancement Toolkit (**FLUEnT**) for processing financial text. It consists of eight different tools for tasks like hypernym detection, numeral claim analysis, readability assessment, sustainability assessment, etc. The objective of the toolkit is to empower the masses and enable investors in making data-driven decisions. It is open-source under MIT license and is openly accessible from Colab and HuggingFace.[1],

CCS CONCEPTS

• **Applied computing** → *Economics*; • **Information systems** → **Information retrieval**; • **Computing methodologies** → **Information extraction**.

KEYWORDS

financial text processing, toolkit, natural language processing

1 INTRODUCTION

People who want to invest in stock markets often face various challenges due to the lack of financial knowledge. The financial domain is full of complicated concepts and jargons. Committing minor mistakes while investing can have adverse effect on the returns. Professional investors also get perplexed by the information overload which inhibits them from making decisions in real-time. To address these challenges, we developed the Financial Language Understandability Enhancement Toolkit (**FLUEnT**) which consists of eight different tools to cater to the needs of the general people and the investors. Figure 1 presents an overview of the toolkit and the functionalities of the constituent tools. We developed four of these tools. They are marked as **1, 2, 5** and **6** in Figure 1. For the remaining four tools (marked as **3, 4, 7** and **8** in Figure 1), we leverage existing open-source models and artefacts. The hypernym detection and readable assessment tools aim to enhance the financial literacy of the masses by providing them with suitable hypernyms (generic forms) of complex financial words (FW) and helping them to filter out the easy-to-understand (i.e., readable) content. The other tools of this toolkit help investors to summarize financial texts (FT) and understand the sentiment, sustainability, Forward-looking statements (FLS), Environmental, Social, and Governance (ESG) aspects of sentences present in FT. Furthermore, the claim detection tool (CD) looks to classify each numeral present in FT as in-claim or out-of-claim.

Our contributions

We have developed **FLUEnT** which can empower the investors in making data-driven decisions and aid in spreading financial literacy. Subsequently, we have deployed and open-sourced this toolkit[1] for non-commercial use. A live demonstration is available in YouTube[2]. The novelty of the system lies in the fact that, like a Swiss knife, it solves eight different use cases in real-time to empower seasoned as well as future investors. It intelligently picks up difficult words and numbers from financial texts and provides users with their hypernyms and 'claim' categories respectively. Moreover, it returns summary of the entered FT in addition to sentence wise sentiment, readability, sustainability, ESG and FLS classes.

We expect that the popularity of **FLUEnT** will grow over time among professional investors and common people who want to invest in the stock markets. Governments, policymakers and non-governmental organizations (NGOs) can use it readily for promoting financial literacy. Above all researchers working in this space can readily use the tools and libraries for their research.

2 RELATED WORK

Table 1 presents a list of related tools and their functionalities alongside **FLUEnT**. As of July 2022, only 9 out of 12 existing tools have a User Interface (UI) and only 6 of them are running live. Most of these tools deal with information extraction and present few analyses from financial reports. Unlike these tools, **FLUEnT** is relatively more comprehensive and it provides eight different functionalities. As we have two different variations of FinBERT, namely [3] and [16], we refer to them as FinBERT(a) and FinBERT(b) respectively. In addition to these tools, there are several proprietary tools and cloud services like SentiMine[3], Augmented Financial Analyst[4], etc. However, discussing them is beyond our scope as these tools require subscriptions.

[1]https://colab.research.google.com/drive/1-KBBKByCU2bkyAUDwW-h6QCSqWI8z127?usp=sharing
https://huggingface.co/spaces/sohomghosh/FLUEnT

This manuscript is the authors' pre-print version. It has been accepted at the 6th Joint International Conference on Data Science & Management of Data (10th ACM IKDD CODS and 28th COMAD).
https://doi.org/10.1145/3570991.3571067

[2]https://youtu.be/Bp8Ij5GQ59I
[3]https://www.lseg.com/about-lseg/labs/sentimine
[4]https://yseop.com/solutions/augmented-financial-analyst

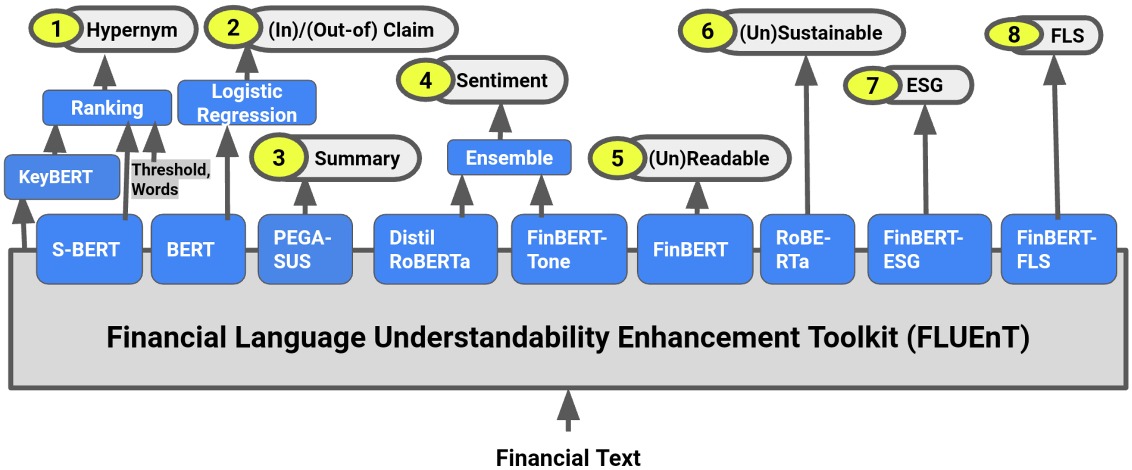

Figure 1: Overview of Financial Language Understandability Enhancement Toolkit

Tools	UI	Live	Functionalities
Financial Term Visualization [29]	No	No	Risk assessment, FT Identification & Visualization from financial reports
FINCHAN [2]	Yes	No	Syntactic & semantic information extraction, & Summarization, Text-to-speech conversion of financial instant messages
FIN10K [21]	Yes	Yes	Extracts relevant portions from 10-K reports & visualizes risk levels & sentiments of keywords
Financial Chatbot [7]	Yes	No	Document search, Topic extraction & Clustering
RegMiner [30]	Yes	Yes	Extraction & Visualization of restrictions present in regulatory documents
ClimateQA [23]	Yes	No	Extraction of climate related sections from financial reports using question answering
FinBERT(b) [16]	No	No	Sentiment Analysis, FLS Assessment & ESG Assessment
FedNLP [19]	Yes	Yes	Summarization, Sentiment Analysis, Topic Models, Federal Funds Rate Movement Rate Prediction
EDGAR-CRAWLER [22]	No	No	Extraction of texts from financial reports
FinRead [13]	Yes	Yes	Readability Assessment
FiNCAT-2 [9]	Yes	Yes	Claim Detection
Financial_Analyst_AI [demo link]	Yes	Yes	Voice-to-Text, Summarization, Sentiment Analysis, FLS Assessment, Company Names & Location Identification
FLUEnT [Demo] [Video] [Colab]	Yes	Yes	**Keyword & Hypernym Detection, Claim Detection, Summarization, Sentiment Analysis, Readability Assessment, Sustainability Assessment, ESG Assessment & FLS assessment**

Table 1: Comparison of FLUEnT with existing non-proprietary tools

3 CONSTITUENT TOOLS

FLUEnT consists of eight different tools. Inputs, outputs, development process, and performance for each of these tools is summarized in Table 2. In this section, we present a detailed explanation for all of them. We chose the underlying models based on their performance and availability.

3.1 Hypernym Detection (HD)

Complex terms can be explained easily using their generic forms or hypernyms. For example, we can explain the FT "*alternative debentures*" by mentioning it's hypernym i.e. "*it is a kind of bond*". A tool to detect hypernyms is useful to learn financial jargons effortlessly. Given an FT, we extract the top three keywords from it using KeyBERT [15]. Users have an option to look for hypernyms of these keywords or other FW they manually enter. Chopra and Ghosh [5] fine-tuned a FinBERT(a) [3] model on the FinSim-3 dataset [18] using the sentence BERT architecture [27] for financial hypernym detection. For all the keywords or FW, we use the fine-tuned sentence BERT embeddings to calculate its cosine similarity with a set of seventeen pre-defined hypernyms. We provide users with the hypernyms corresponding to the entered financial words only when their similarity is more than the threshold set by the user using the slider present in the tool.

3.2 Claim Detection (CD)

Executives try to lure investors by making claims which may not always be true. The sentence, "*In the year 2021, the markets were bullish. We expect to boost our sales by 80% this quarter.*" has two numerals 2021 and 80%. Among these two, "*2020*" is 'out-of-claim' and "*80%*" is 'in-claim'. The CD tool can alert investors by detecting numerals in FT which are 'in-claim'. For each of the numbers present in an FT, we extract its BERT-base [6] embedding given a context window of 6 words before and after it. Subsequently, we use Logistic Regression to classify it as either in-claim or out-of-claim. The methodology of the CD tool is described in [11] and [8]. The

Tool	Input	Output	Base Models	Developer	Development Dataset (Size)	Performance
HD	FW	Generic form of each terms	SBERT+FinBERT(a)	Ours (Ghosh et al.)	FinSim-3 (1,050 FW)	Accuracy: 0.9170
CD	FT	Each numeral in FT: in-claim or out-of-claim	BERT-base	Ours (Ghosh et al.)	FinNum-3 English (10,720 FT)	Macro-F1: 0.8238
SM	FT	Summary of the entire FT	PEGASUS	Passali et al.	Bloomberg articles (2,000 FT)	Rouge-L: 18.14
SA	FT	Each sentences present in FT: positive, negative or neutral	1) BERT-base 2) DistillRoBerta-base	1) Huang et al. (finbert.ai) 2) Romero M.	1) Analyst reports of S&P 500 firms (10,000 FT) 2) Financial PhraseBank (4,840 FT)	1) Accuracy: 0.882 2) Accuracy: 0.9823
RA	FT	Each sentences present in FT: readable or non-readable	FinBERT(a)	Ours (Ghosh et al.)	FinRAD (13,112 FW definitions)	AUROC: 0.9927
SN	FT	Each sentences present in FT: sustainable, non-sustainable or none	RoBERTa-base	Ours (Ghosh et al.)	FinSim-4-ESG (2,265 FT)	Accuracy: 0.9317
ESG	FT	Each sentences present in FT: Environmental, Social, Governance or None	FinBERT(b)	Huang et al. (finbert.ai)	Annual & ESG reports of firms (2,000 FT)	Accuracy: 0.895
FLS	FT	Each sentences present in FT: Specific-FLS, Non-specific FLS or Not-FLS	FinBERT(b)	Huang et al. (finbert.ai)	MD&A sections of annual reports of Russell 3000 firms (3,500 FT)	Accuracy: 0.853

Table 2: Different constituent tools and their characteristics. FT & FW means financial texts & words respectively.

CD model was trained on FinNum-3 (English) dataset [4]. We have further released two tools FiNCAT [10] and FiNCAT-2 [9] to help investors in detecting claims present in numerals within FT.

3.3 Summarization (SM)

In today's fast-moving world, time and money are almost equivalent. With the advent of Big Data, investors are overloaded with information; they do not have the time to assimilate all the information. Thus, the SM tool aims to help them by removing irrelevant and less relevant facts and providing them with only the necessary information. We integrated the financial summarizer built by Passali et al. [25] in our toolkit. The SM tool provides a summary of the entered FT using the PEGASUS [32] model.

3.4 Sentiment Analysis (SA)

Lately, financial opinion mining has gained huge interest. Some of the open-sourced models include FinBERT-tone[5] (a derivative of FinBERT(b) [16]) developed by fine-tuning BERT-base [6] on analyst reports of S&P 500 firms, and distilRoberta-financial-sentiment[6] developed by fine-tuning DistillRoBERTa-base [28] on the Financial PhraseBank dataset [24]. For each sentence in an FT, we evaluate both these models and produce the label with the greater probability. The output labels are: 'positive', 'negative' and 'neutral'.

3.5 Readability Assessment (RA)

To ensure that the non-investors who want to invest in the stock market do not get overwhelmed, it is essential to present them with information which is easy to understand ('readable'). Since the formula-based readability scores (like Automated Readability Index, Coleman Liau index, etc.) do not hold good for the financial domain, we proposed a new financial readability assessment dataset, FinRAD [14], and a FinBERT(a) [3] based neural model to classify definition of financial terms. We use this model to assess whether each sentence in the entered FT is 'readable' or not. Subsequently, we have developed a tool FinRead [13] to address this.

3.6 Sustainability Assessment (SN)

Socially conscious investors look for sustainable avenues for investments. We used the FinSim-4-ESG (shared task 2) [17] dataset to fine-tune a RoBERTa-base model [20] for classification of each sentences present in an FT into three classes 'sustainable', 'non-sustainable' or none (represented by '-') [12].

3.7 ESG Assessment (ESG)

Investors look for ESG ratings of companies they want to invest in. It is very tedious to read ESG reports of every organization. For each sentence in an FT, this tool detects whether it is related to 'Environment', 'Social', 'Governance' or none. Huang et al.[7] developed the underlying model by fine-tuning the FinBERT(b) model [16].

3.8 FLS Assessment (FLS)

FLS help investors to understand the future conditions of the financial market. Huang et al. proposed FinBERT-FLS[8] for classifying financial texts as 'Specific-FLS', 'Non-specific FLS' or 'Not-FLS'. It was developed by fine-tuning FinBERT(b) [16] on a set of 3,500 manually annotated financial sentences. We use the FinBERT-FLS model for classifying each sentence present in the entered FT into the above mentioned classes.

4 SYSTEM OVERVIEW

In this section, we discuss the underlying technologies and elaborate the user interface of **FLUEnT** in details.

4.1 Implementation Details

These constituent models have been trained in PyTorch [26] using HuggingFace Transformers [31]. We carried out the experiments on

[5] https://huggingface.co/yiyanghkust/finbert-tone
[6] https://huggingface.co/mrm8488/distilroberta-finetuned-financial-news-sentiment-analysis
[7] https://huggingface.co/yiyanghkust/finbert-esg
[8] https://huggingface.co/yiyanghkust/finbert-fls

Google Colab[9] (runtime: GPU). The user interface has been created using Gradio [1] and hosted on Colab and HuggingFace Spaces[1].

4.2 Demonstration Interface

The Graphical User Interface (GUI) of **FLUEnT** primarily consists of two sections: the inputs (Ref: Figure 2) and the outputs (Ref: Figure 3. In the input section the user enters an FT in the text-box above (`TB-1`) and sets a confidence `threshold` using the slider. The GUI also provides a number of examples. On clicking the "*Get Keywords for Hypernym Detection*" button, the GUI shows the top three keywords extracted from the entered text based on KeyBERT [15]. The keywords are shown in the text-box (`TB-2`) below. The user can look for hypernyms corresponding to these keywords. The user can also alter the contents of this text-box (`TB-2`) by manually entering keywords of his/her choice.

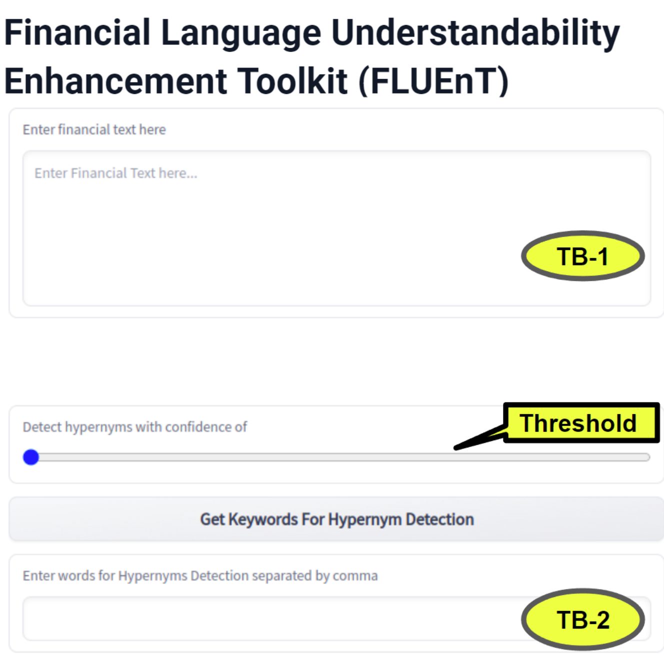

Figure 2: Inputs with text-boxes (`TB-1`, `TB-2`) and `threshold` field marked

The output section consists of eight different tools presented in four tabs. Each of these tools can be used independently. This saves time as well as computing resources. At any given time, users can select any of the four tabs and press a `Get` button corresponding to the tool they want to use. The HD tool extracts generic forms of the FW entered in the `TB-2` (or, the extracted keywords). These hypernyms are presented only when their similarity score is above the threshold set by using the slider. The CD tool first extracts numerals from the FT entered in `TB-1`. Subsequently, it classifies each of the numerals as 'in-claim' or 'out-of-claim'. The 'in-claim' and 'out-of-claim' numerals are presented in red and green colour respectively. The remaining six tools use the FT entered in `TB-1`

[9]https://research.google.com/colaboratory/

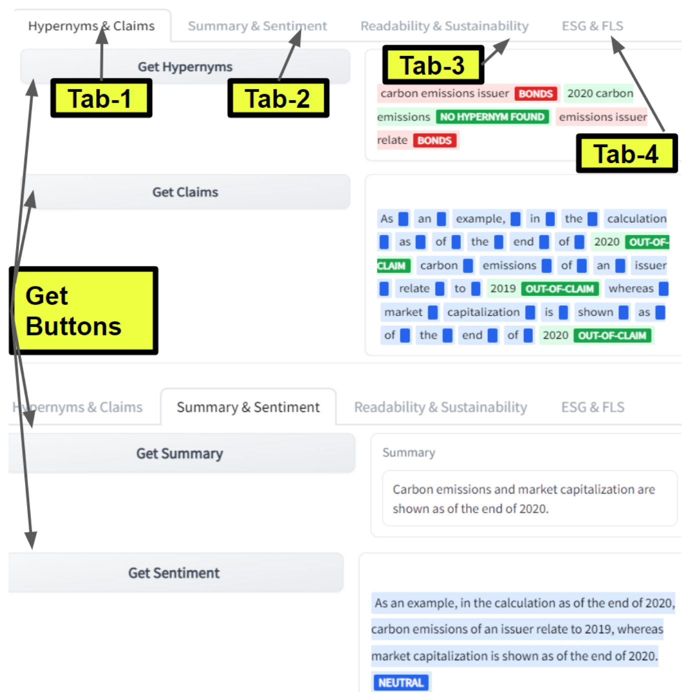

Figure 3: Outputs (HD, CD, SM, and SA). Similar outputs are generated for RA, SN, ESG and FLS.

for making predictions. We highlight each of the sentences present in the FT and mention the predicted categories next to them. This enhances the usability of **FLUEnT**.

FLUEnT is available on Google Colab, HuggingFace Spaces[1] and a tutorial video is available[2] for the convenience of the users.

5 CONCLUSION

In this paper, we presented **FLUEnT**, a toolkit that helps in improving the comprehensibility of complex FT. It performs several tasks on financial texts, like HD, CD, SM, SA, RA, etc. In future, we want to add various other features like uploading documents (PDFs) as input, extracting relevant portions from these documents which relates to finance and then performing various tasks on these portions. We also want to work on collecting feedback from the users and develop a browser-based extension that will scan content from the financial web pages and help investors in understanding it. Another direction for future work is to develop a multi-task model which will reduce the overall size of the tool and improve its throughput.

REFERENCES

[1] Abubakar Abid, Ali Abdalla, Ali Abid, Dawood Khan, Abdulrahman Alfozan, and James Zou. 2019. Gradio: Hassle-Free Sharing and Testing of ML Models in the Wild. *arXiv preprint arXiv:1906.02569* (2019).

[2] Abejide Ade-Ibijola. 2016. FINCHAN: A Grammar-Based Tool for Automatic Comprehension of Financial Instant Messages. In *Proceedings of the Annual Conference of the South African Institute of Computer Scientists and Information Technologists* (Johannesburg, South Africa) *(SAICSIT '16)*. Association for Computing Machinery, New York, NY, USA, Article 1, 10 pages. https://doi.org/10.1145/2987491.2987518

[3] Dogu Araci. 2019. FinBERT: Financial Sentiment Analysis with Pre-trained Language Models. arXiv:1908.10063 [cs.CL] https://arxiv.org/abs/1908.10063

[4] Chung-Chi Chen, Hen-Hsen Huang, Yu-Lieh Huang, Hiroya Takamura, and Hsin-Hsi Chen. 2022. Overview of the ntcir-16 finnum-3 task: investor's and manager's fine-grained claim detection. In *Proceedings of the 16th NTCIR Conference on Evaluation of Information Access Technologies*. NII, Tokyo, Japan, 87–91. http://research.nii.ac.jp/ntcir/workshop/OnlineProceedings16/pdf/ntcir/01-NTCIR16-OV-FINNUM-ChenC.pdf

[5] Ankush Chopra and Sohom Ghosh. 2021. Term Expansion and FinBERT fine-tuning for Hypernym and Synonym Ranking of Financial Terms. In *Proceedings of the Third Workshop on Financial Technology and Natural Language Processing (FinNLP@IJCAI 2021)*. -, Online, 46–51. https://aclanthology.org/2021.finnlp-1.8

[6] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Association for Computational Linguistics, Minneapolis, Minnesota, 4171–4186. https://doi.org/10.18653/v1/N19-1423

[7] Boris Galitsky and Dmitry Ilvovsky. 2019. On a Chatbot Conducting a Virtual Dialogue in Financial Domain. In *Proceedings of the First Workshop on Financial Technology and Natural Language Processing*. Macao, China, 99–101. https://aclanthology.org/W19-5517

[8] Sohom Ghosh and Sudip Kumar Naskar. 2022. Detecting context-based in-claim numerals in Financial Earnings Conference Calls. *International Journal of Information Technology* – (2022). https://doi.org/10.1007/s41870-022-00952-7

[9] Sohom Ghosh and Sudip Kumar Naskar. 2022. FiNCAT-2: An enhanced Financial Numeral Claim Analysis Tool. *Software Impacts* 12 (2022), 100288. https://doi.org/10.1016/j.simpa.2022.100288

[10] Sohom Ghosh and Sudip Kumar Naskar. 2022. FiNCAT: Financial Numeral Claim Analysis Tool. In *Companion Proceedings of the Web Conference 2022 (WWW '22 Companion)* (Virtual Event, Lyon, France). Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3487553.3524635

[11] Sohom Ghosh and Sudip Kumar Naskar. 2022. Lipi at the ntcir-16 finnum-3 task: ensembling transformer based models to detect in-claim numerals in financial conversations. In *Proceedings of the 16th NTCIR Conference on Evaluation of Information Access Technologies*. NII, Tokyo, Japan, 92–94. http://research.nii.ac.jp/ntcir/workshop/OnlineProceedings16/pdf/ntcir/02-NTCIR16-FINNUM-GhoshS.pdf

[12] Sohom Ghosh and Sudip Kumar Naskar. 2022. Ranking Environment, Social And Governance Related Concepts And Assessing Sustainability Aspect Of Financial Texts. In *Proceedings of the Fourth Workshop on Financial Technology and Natural Language Processing (FinNLP@IJCAI-ECAI 2022)*. -, Vienna, Austria, 87–92. https://mx.nthu.edu.tw/~chungchichen/FinNLP2022_IJCAI/14.pdf

[13] Sohom Ghosh, Shovon Sengupta, Sudip Naskar, and Sunny Kumar Singh. 2021. FinRead: A Transfer Learning Based Tool to Assess Readability of Definitions of Financial Terms. In *Proceedings of the 18th International Conference on Natural Language Processing (ICON)*. NLP Association of India (NLPAI), National Institute of Technology Silchar, Silchar, India, 658–659. https://aclanthology.org/2021.icon-main.81

[14] Sohom Ghosh, Shovon Sengupta, Sudip Kumar Naskar, and Sunny Kumar Singh. 2022. FinRAD: Financial Readability Assessment Dataset - 13,000+ Definitions of Financial Terms for Measuring Readability. In *Proceedings of the The 4th Financial Narrative Processing Workshop (FNP@LREC2022)*. European Language Resources Association, Marseille, France, 1–9. http://lrec-conf.org/proceedings/lrec2022/workshops/FNP/pdf/2022.fnp-1.1.pdf

[15] Maarten Grootendorst. 2020. KeyBERT: Minimal keyword extraction with BERT. https://doi.org/10.5281/zenodo.4461265

[16] Allen Huang, Hui Wang, and Yi Yang. 2020. FinBERT—A Large Language Model Approach to Extracting Information from Financial Text. http://dx.doi.org/10.2139/ssrn.3910214

[17] Juyeon Kang, Mehdi Kchouk, Sandra Bellato, Mei Gan, and Ismail El Maarouf. 2022. FinSim4-ESG Shared Task: Learning Semantic Similarities for the Financial Domain. Extended edition to ESG insights. In *Proceedings of the Fourth Workshop on Financial Technology and Natural Language Processing (FinNLP@IJCAI-ECAI 2022)*. -, Vienna, Austria, 57–63. https://mx.nthu.edu.tw/~chungchichen/FinNLP2022_IJCAI/9.pdf

[18] Juyeon Kang, Ismail El Maarouf, Sandra Bellato, and Mei Gan. 2021. FinSim-3: The 3rd Shared Task on Learning Semantic Similarities for the Financial Domain. In *Proceedings of the Third Workshop on Financial Technology and Natural Language Processing*. -, Online, 31–35. https://aclanthology.org/2021.finnlp-1.5

[19] Jean Lee, Hoyoul Luis Youn, Nicholas Stevens, Josiah Poon, and Soyeon Caren Han. 2021. FedNLP: An Interpretable NLP System to Decode Federal Reserve Communications. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Virtual Event, Canada) *(SIGIR '21)*. Association for Computing Machinery, New York, NY, USA, 2560–2564. https://doi.org/10.1145/3404835.3462785

[20] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. https://arxiv.org/abs/1907.11692

[21] Yu-Wen Liu, Liang-Chih Liu, Chuan-Ju Wang, and Ming-Feng Tsai. 2016. FIN10K: A Web-Based Information System for Financial Report Analysis and Visualization. In *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management* (Indianapolis, Indiana, USA) *(CIKM '16)*. Association for Computing Machinery, New York, NY, USA, 2441–2444. https://doi.org/10.1145/2983323.2983328

[22] Lefteris Loukas, Manos Fergadiotis, Ion Androutsopoulos, and Prodromos Malakasiotis. 2021. EDGAR-CORPUS: Billions of Tokens Make The World Go Round. In *Proceedings of the Third Workshop on Economics and Natural Language Processing*. Association for Computational Linguistics, Punta Cana, Dominican Republic, 13–18. https://doi.org/10.18653/v1/2021.econlp-1.2

[23] Alexandra Luccioni, Emily Baylor, and Nicolas Duchene. 2020. Analyzing sustainability reports using natural language processing, In Tackling Climate Change with Machine Learning workshop at NeurIPS 2020. *arXiv preprint arXiv:2011.08073*. https://arxiv.org/abs/2011.08073

[24] Pekka Malo, Ankur Sinha, Pyry Takala, Pekka Korhonen, and Jyrki Wallenius. 2013. FinancialPhraseBank-v1.0. https://www.researchgate.net/publication/251231364_FinancialPhraseBank-v10

[25] Tatiana Passali, Alexios Gidiotis, Efstathios Chatzikyriakidis, and Grigorios Tsoumakas. 2021. Towards Human-Centered Summarization: A Case Study on Financial News. In *Proceedings of the First Workshop on Bridging Human–Computer Interaction and Natural Language Processing*. Association for Computational Linguistics, Online, 21–27. https://www.aclweb.org/anthology/2021.hcinlp-1.4

[26] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems 32*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Eds.). Curran Associates, Inc., 8024–8035. http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf

[27] Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. 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)*. Association for Computational Linguistics, Hong Kong, China, 3982–3992. https://doi.org/10.18653/v1/D19-1410

[28] Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. *ArXiv* abs/1910.01108 (2019).

[29] Ming-Feng Tsai and Chuan-Ju Wang. 2012. Visualization on Financial Terms via Risk Ranking from Financial Reports. In *Proceedings of COLING 2012: Demonstration Papers*. The COLING 2012 Organizing Committee, Mumbai, India, 447–452. https://aclanthology.org/C12-3056

[30] Karolin Winter, Manuel Gall, and Stefanie Rinderle-Ma. 2020. RegMiner: Taming the Complexity of Regulatory Documents for Digitalized Compliance Management. In *Proceedings of the Best Dissertation Award, Doctoral Consortium, and Demonstration & Resources Track at BPM 2020*. 112–116. http://ceur-ws.org/Vol-2673/paperDR10.pdf

[31] 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. 2019. HuggingFace's Transformers: State-of-the-art Natural Language Processing. https://doi.org/10.48550/ARXIV.1910.03771

[32] Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter Liu. 2020. PEGASUS: Pre-training with Extracted Gap-sentences for Abstractive Summarization. In *Proceedings of the 37th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 119)*, Hal Daumé III and Aarti Singh (Eds.). PMLR, 11328–11339. https://proceedings.mlr.press/v119/zhang20ae.html