

CR-COPEC: Causal Rationale of Corporate Performance Changes to Learn from Financial Reports

Ye Eun Chun^{1*}, Sunjae Kwon^{2*}, Kyunghwan Sohn³, Nakwon Sung⁴,
Junyoup Lee⁵, Byungki Seo¹, Kevin Compher⁶, Seung-won Hwang⁴, Jaesik Choi^{3,7†}

¹UNIST, ²UMass Amherst, ³KAIST, ⁴SNU

⁵ Ajou University, ⁶ Foresight Institute, ⁷ INEEJI

Abstract

In this paper, we introduce CR-COPEC called *Causal Rationale of Corporate Performance Changes* from financial reports. This is a comprehensive large-scale domain-adaptation causal sentence dataset to detect financial performance changes of corporate. CR-COPEC contributes to two major achievements. First, it detects causal rationale from 10-K annual reports of the U.S. companies, which contain experts' causal analysis following accounting standards in a formal manner. This dataset can be widely used by both individual investors and analysts as material information resources for investing and decision-making without tremendous effort to read through all the documents. Second, it carefully considers different characteristics which affect the financial performance of companies in twelve industries. As a result, CR-COPEC can distinguish causal sentences in various industries by taking unique narratives in each industry into consideration. We also provide an extensive analysis of how well CR-COPEC dataset is constructed and suited for classifying target sentences as causal ones with respect to industry characteristics. Our dataset and experimental codes are publicly available¹.

1 Introduction

Many critical decisions on events may require appropriate explanations of decisions based on accurate causal rationale. Justifying the root statements is directly related to identifying the causes of the events. When one could observe an event that is an apparent cause of a desired outcome, one can make a proper decision with confidence on events.

There has been extensive research in extracting causes of events in numerical data. As an example, Granger causality finds linear temporal

dependence between two (or more) temporal sequences (Bressler and Seth, 2011). Shapley values derive the contributions of individual input attributes when a decision is made by a complex function (Shapley, 1971). These methods can explain numerical causes of decisions made by automated systems, e.g. Robo-advisers in financial services (Hwang et al., 2016; Karuna, 2019; Chhatwani, 2022). However, human uses various types of information including numerical and textual inputs when making an important decision. As an example, analysts write summarized reports by extracting related causal information from multiple textual sources in conference calls, annual reports, earning statements and markets reports.

Our research goal is to extract causal rationales from financial reports. In general, when investing firms predict a certain financial performance, they provide analysts' reports to support their predictions. Thus, we want to generate appropriate explanations for performance changes of a corporate from official documents. Our algorithm classifies causal sentences from documents and provides binary classification results.

We consider fine-tuning of Pre-trained Neural Language Models (NLMs) for the causality modeling tasks. Pre-trained NLMs have been state-of-the-arts for many Natural Language Processing (NLP) tasks. For example, NLMs such as BERT (Devlin et al., 2018) and ALBERT (Lan et al., 2019) demonstrate outstanding performance in some tasks such as answering questions (Clark et al., 2020; Suissa et al., 2023) and computing conditional probabilities of masked words in a sentence (Kwon et al., 2022). Nonetheless, recent research indicates that the size of human-annotated data continues to be a significant factor influencing the performance of models (Gu et al., 2022; Mehrafarin et al., 2022).

There are previous works that introduce dataset for causality detection in the financial domain (El-Haj et al., 2016; Mariko et al., 2020). However,

*Both authors contributed equally to this work.

† Corresponding author: jaesik.choi@kaist.ac.kr

¹<https://github.com/CR-COPEC/CR-COPEC>

existing studies lacked consideration of industry-specific characteristics. Thus, it can be beneficial to consider it because the items in the financial statements that greatly affect financial performance are different for each company’s primary business.

Our main contribution is to collect sufficient annotations to achieve reasonable causality detection performance with NLMs. We achieved this by collecting over 283K sentences from 1,584 10-K annual reports, that give a detailed summary of the financial status and business operations of each company, along with audited financial statements. Then, we manually label individual sentences whether the sentences explain the cause of certain financial performance changes. We name the 283K pairs of a sentence and a corresponding label as *Causal Rationale of Corporate Performance Changes* (CR-COPEC). CR-COPEC is built on a large scale with guides of experts in the financial domain. Trained with our dataset, BERT can distinguish sentences containing main causes of its financial events, from annual reports written officially from most U.S. public companies. Thus, individual investors can save efforts to read a huge amount of reports by themselves.

However, we find collecting the dataset does not solve the problem as itself. One challenge we observe in the process of annotation is **diverse causality** in industries. That is, we find diverse causalities over different sectors. Another challenge is **imbalanced training**, where the number of data for each sector varies. Thus, we need to build a model carefully, as applying a common model to all sectors does not work in our problem setting. We provide extensive analysis on our CR-COPEC dataset to overcome these issues.

2 Related Work

The causal rationale is “the true sufficient rationales to fully predict and explain the outcome without superfluous information.” (Zhang et al., 2023) and extracting rationale is invaluable when a decision has to be made. Research on extracting rationale from text has been tried with various types of text documents (Blanco et al., 2008; Ittoo and Bouma, 2011; Lu et al., 2022). A model detecting and identifying rationale from chat messages was suggested in (Alkadhi et al., 2017). Bug reports from Chrome web browser were used as main sources to extract rationale (Rogers et al., 2012). Moreover, patent documents were utilized to discover design

rationale (Liang et al., 2012). To extract causal textual structures, one may consider a rule-based system where specific words such as ‘due to’, ‘owing to’ and ‘affects’ are listed to identify sentences including causal information for prediction (Girju and Moldovan, 2002; Chang and Choi, 2006; Sakai et al., 2015).

Diverse studies have explored causality extraction datasets (Xu et al., 2020; Ali et al., 2021; Yang et al., 2022). Some widely recognized datasets include SemEval-2007 task 4 (Girju et al., 2007), SemEval-2010 task 8 (Hendrickx et al., 2019), PDTB 2.0 (Prasad et al., 2008), BioInfer (Pyysalo et al., 2007), and ADE (Gurulingappa et al., 2012). Among these, three datasets stem from general domain, while two originate from the biomedical field (BioInfer and ADE). The maximum number of causality sizes of these datasets is less than 10K (PDTB 2.0, 9,190 causal examples). In contrast, CR-COPEC contains a larger substantial data size of 105,861 sentences, including 11,132 sentences with causal rationales. In addition, CO-COPEC targets the financial domain especially focused on corporate performance changes from financial reports. These characteristics contribute to the uniqueness of CR-COPEC and potentially helpful for the decision of financial experts.

There are previous studies that introduce causal rationale corpus in the financial domain. To identify causal sentences from UK Preliminary Earning Announcements (PEAs), thirteen performance keywords including ‘sales’, ‘revenue’ and ‘turnover’ are used and selected sentences are annotated by human (El-Haj et al., 2016). Mariko et al. (2020)’s study is the most related our research. Herein, the authors built FinCausal corpus collected from financial news and websites which labeled with tags indicating the presence of causality and causal chunk as quantitative or non-quantitative. The corpus consists of two subtasks. The first subtask (FinCausal Task 1) is a binary classification task that targets to extract text having causality. The second subtask (FinCausal Task 2) is a relation extraction task identifying the substrings that indicate cause and effect. The main difference between CR-COPEC and FinCausal is that texts are written with a formal tone in 10-K reports since they should comply with regulatory rules. In addition, CR-COPEC solely concentrates on causal rationales of accounting items considering unique characteristics of various industries. Meanwhile, FinCausal dataset itself is written

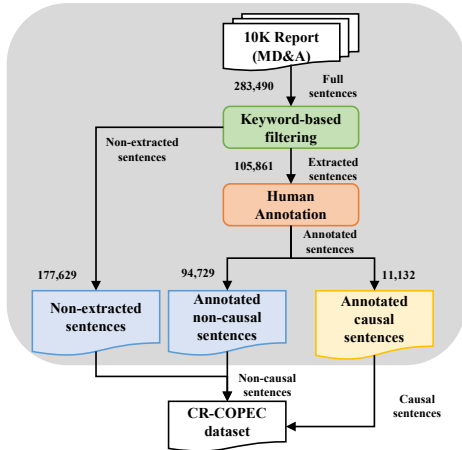


Figure 1: Overview of CR-COPEC dataset construction process.

with a casual tone in general because it is collected from news or web contents. The detailed comparison of two datasets is described in Appendix E

3 Causal Rationale of Corporate Performance Changes Dataset

Our goal is to collect causal rationale of sentences that contain causal rationales in predicting changes in corporate performance based on key accounting items within the target industrial sector. For this purpose, we target Management’s Discussion and Analysis (MD&A) section of 10-K reports since it provides the company’s perspective on its operations and financial results of the prior fiscal year. We gather MD&A reports filed in 1997 and 2017. MD&A of 1997 is gathered from the existing MD&A data repository (Kogan et al., 2009) and 2017’s is downloaded directly from the Securities and Exchange Commission (SEC) system. Then, we construct CR-COPEC dataset through keyword-based filtering (Section 3.1) and human annotation process (Section 3.2) illustrated in Figure 1. The validity of the dataset is verified in Section 3.3. In addition, we provide protocols of the dataset to train NLMs in Section 3.4.

3.1 Keyword-based Filtering

Since sentences we aim for detecting are rare in a single document, it is costly expensive to manually annotate every collected sentence. To mitigate the cost issue, previous studies extract text containing keywords such as domain specific terminologies (El-Haj et al., 2016; Fonseca et al., 2023) or causal phrases (Sakai et al., 2015; Dürlich et al., 2022). In our work, we utilize keyword-based filtering to extract candidate causal sentences. In particular, we

form a list of keywords² including causal trigger phrases of changes in financial performance. We filter out approximately 62.7% *non-extracted sentences* (177,629 out of 283,490) and simultaneously extract the remaining 37.3% *extracted sentences* (105,861 out of 283,490) from the total sentences of the MD&A section. Section 3.3 discusses on the coverage of the keyword list.

3.2 Human Annotation Process

3.2.1 Data Annotation following SIC

We describe the process of analyzing causal rationale of sentences by considering companies’ Standard Industrial Classification (SIC) that classifies the U.S. companies according to their primary business.³ For this, we divide 10-K reports into twelve industries which are predefined with respect to the SIC codes (French, 2019). These twelve categories include 1) consumer non-durables, 2) consumer durables, 3) manufacturing, 4) energy, 5) chemicals, 6) business equipment, 7) telephone, 8) utilities, 9) shops, 10) health, 11) finance and 12) others. Note that, throughout the paper, sectors are numbered in this order.

Since major factors that affect the financial performance of a company in each industry can be different, we build an annotation guideline for each sector by taking different accounting items among industries into consideration. Under the supervision of a finance faculty at the school business administration, we craft the annotation guideline by considering the main items of the balance sheets and income statements since those items are closely related to financial performance of companies. Then, we proceed annotation according to the annotation guideline. Each MD&A document is randomly assigned to an annotator. The number of documents and sentences from each sector is reported in Table 1. Examples of sentence containing main factors mostly observed in each sector are shown in Table 2. Additional examples of each sector’s main factors are demonstrated in Appendix B and causal/non-causal rationale of sentences can be found in Appendix C.

²1) result from, 2) result of, 3) due to, 4) cause, 5) impact, 6) because of, 7) increase, 8) decrease, 9) decline, 10) negative, 11) contract and 12) significant

³<https://www.sec.gov/info/edgar/siccodes.htm>

	1	2	3	4	5	6	7	8	9	10	11	12	Total
Causes	1,145 (5.70%)	346 (4.55%)	1,528 (4.36%)	497 (3.20%)	283 (5.62%)	2,505 (4.19%)	721 (3.24%)	490 (2.46%)	1,424 (4.30%)	563 (2.80%)	558 (2.18%)	1,072 (5.51%)	11,132 (3.93%)
Non-Causes	18,959	7,255	33,485	15,054	4,750	57,292	21,538	19,398	31,698	19,529	25,014	18,386	272,358
Total	20,104	7,601	35,013	15,551	5,033	59,797	22,259	19,888	33,122	20,092	25,572	19,458	283,490
# of Doc	140	61	248	79	36	379	76	55	138	127	119	126	1,584

Table 1: Dataset (CR-COPEC) Composition: number of sentences and documents in each sector (The ratio of causes to total sentences is in parenthesis).

Sector of Industry - Example [Document]
Consumer Durables - The sales increase for fiscal 1996 was principally due to improved sales of buses and ambulances. [Collins Industries, Inc., January, 1997]
Manufacturing - The increase in 1996 net sales was due primarily to increases in sales revenues recognized on the contracts to construct the first five Sealift ships, the Icebreaker and the forebodies for four double-hulled product tankers, which collectively accounted for 63% of the Company’s 1996 net sales revenue. [Avondale Industries, Inc., March, 1997]
Energy - Gas revenue increased \$32.9 million or 81% because of a 39% price increase combined with a 30% increase in production. [Cross Timbers Oil Co., March, 1997]
Chemicals - Loss of margin was principally due to sales price decreases and raw material price increases in the pyridine and related businesses, and higher manufacturing costs due to weather related problems in the first quarter 1994. [Cambrex Corp., March, 1997]
Finance - Mortgage investment income decreased for 1995 as compared to 1994 primarily due to the assignment to HUD of the mortgage on El Lago Apartments in June 1995. [American Insured Mortgage Investors Series 85 L P, March, 1997]

Table 2: Examples of causal rationale of sentences by each sector.

3.2.2 Annotator Sensitivity

We distinguish annotators into general and unskilled annotators based on the proficiency of the guideline. We regard annotators who 1) participated in the development of the annotation guideline and 2) labeled more than 30K sentences as general annotators, otherwise unskilled annotators. The general annotators label each extracted sentence as non-causal or causal one. 100,046 of the sentences were tagged by two general annotators.

We regard labels annotated by general annotators as a standard and train a model with these labels alone. We call this model as an initial teacher model. Then, we apply this teacher model to documents that are labeled from unskilled annotators. If all labels in a single MD&A document match with predictions of the teacher model, we add them to the previous training set. We train another teacher model with the new version of training set again and apply this model to the rest of other documents. We repeat this process until no matched document is found. As a result, we collect 1,584 10-K reports and 105,861 sentences containing 11,132 causal and 94,729 non-causal sentences.

Note that, all the sentences from the same document are assigned to the same annotator, and the annotators label the current sentence based on its previous context. Details on the annotation environment of annotators can be found in Appendix D.

3.3 The Validity of the Dataset

In this section, we discuss the validity of our dataset construction process; 1) Keyword-based filtering

	Causal	Non-causal	Overall
Extracted	38	347	385
Non-extracted	6	677	683
Overall	44	1,024	1,068

Table 3: The result of the coverage analysis of causal sentences extracted by the keywords.

and 2) Human Annotation.

3.3.1 Verifying Keyword-based Filtering

We randomly sample a small number of sentences from full sentences and classify causal sentences manually. Table 3 is the result of analyzing the coverage of keywords. From 1,068 sampled sentences, 385 sentences are extracted with the keyword list, and the rest of 683 sentences are filtered out during the keyword-based filtering process. Out of extracted sentences, 38 sentences are annotated as causal one. Meanwhile, six sentences are turned out to be causal rationale of sentences from non-extracted sentences. From these results, we can see that the keywords can be an effective indicator of causal rationale of sentences showing high recall performance of 86.4% (38 out of 44). The ratio of causal rationale of sentences and the non-extracted sentences are 9.8% (38 out of 385) and 0.9% (6 out of 683) respectively.

3.3.2 Verifying Human Annotation Process

Human annotation task is sensitive to the reliability of the general annotators. Thus, to verify the reliability between annotators, we conduct the following steps. To be specific, this process includes three phases: 1) Random sampling, 2) Additional annotation and 3) Reliability calculation. Figure 2

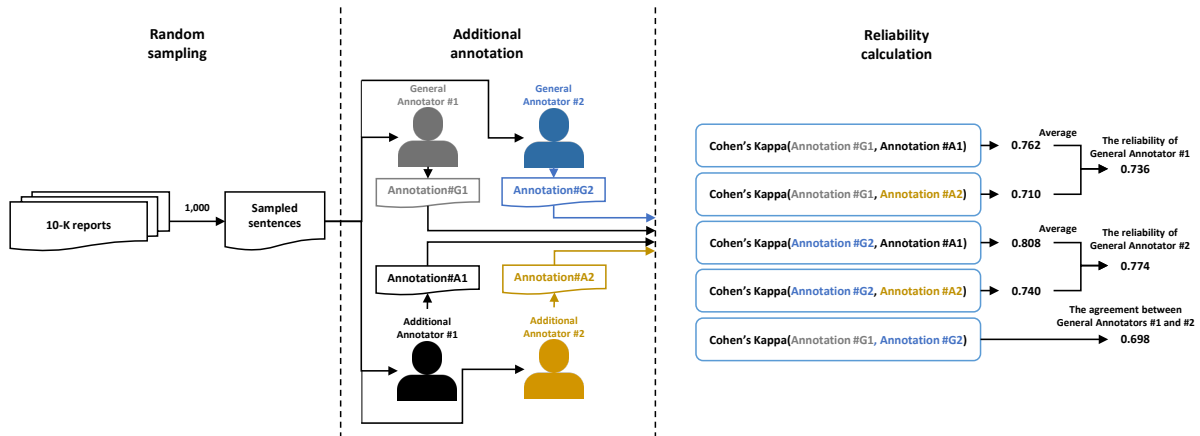


Figure 2: Overview of the verifying the human annotations of general annotators.

	1	2	3	4	5	6	7	8	9	10	11	12
Train/Valid	5.71	4.58	4.38	3.18	5.68	4.18	3.24	2.45	4.29	2.80	2.19	5.53
Test	5.56	4.37	4.25	3.34	5.23	4.24	3.28	2.58	4.33	2.85	2.08	5.38

Table 4: Total Dataset Causal Composition: the ratio of causal sentences in each sector (%).

is a conceptual diagram for the process.

To begin with, we randomly sample 1,000 sentences from 10-K reports. Then, we ask for the general annotators and two additional annotators⁴ to annotate the sampled sentences. In this case, those additional annotators were trained on the annotation guideline and corpus labeling. As a result, we can get annotation #G1, and annotation #G2 from the general annotators and annotation #A1, and annotation #A2 from the additional annotators.

Finally, we calculate the reliability of the sampled sentences of the general annotators #1 and #2. For this purpose, we use Cohen’s Kappa (Cohen, 1960), which evaluates the inter reliability of two annotators. As a result, we observe that the agreements of the general annotator #1 are 0.762, and 0.710 from the additional annotators #1 and #2, respectively, being able to be interpreted as a substantial agreement ($\kappa > 0.6$) (Viera et al., 2005). Furthermore, the agreements of the general annotator #2 are 0.808, and 0.740 from the additional annotators #1 and #2, respectively. Finally, the agreement between general annotators is 0.698.

3.4 Dataset to Validate Training Performance

This section describes how we recompose CR-COPEC dataset as suitable to train NLM. In ad-

⁴Herein, additional annotator #1 is working in the computational linguistic domain for more than 6 years and additional annotator #2 is a financial expert who worked at Financial Supervisory Service for more than 8 years. Note that, both general annotators are $\kappa > 0.7$, when compared to additional annotator #2, indicating that the annotations of the general annotators are reliable.

dition, we provide protocols for different version of datasets through Section 3.4.1 and Section 3.4.2.

3.4.1 Protocols for Total Dataset

In CR-COPEC dataset, the ratio of causal to non-causal is imbalanced, the robustness of learning causality can be affected by the ratio of causal in each train/valid/test dataset. Thus, we carefully select train, valid and test dataset so that each sector composes similar causal ratio between different datasets. In this process, we iterate random selection until the ratio difference is less than 0.5%. The proportion of train/valid/test follows a 81%/9%/10% of total CR-COPEC dataset. The causal composition of train/valid/test dataset for each sector is reported in Table 4.

3.4.2 Protocols for Fraction of Dataset

In Section 4.3, we compare cross-sector performances to figure out transferable information between sectors. As we observe in Table 1, the size of each sector varies from 5K to 60K. Thus, in order to control the effect of size on performance, we randomly sample 3,500 sentences from each sector of train dataset. Since the smallest number of train datasets among all sectors is less than 4,000.

4 Analysis on CR-COPEC Dataset

4.1 Experimental Settings

We conduct experiments with baseline models on CR-COPEC consisting of sentences and corresponding labels in supervised learning. Since the dataset is highly imbalanced, we use area under

		Sector Test Dataset												Avg.
		1	2	3	4	5	6	7	8	9	10	11	12	
Sector-only Same-size Model	1	58.62	60.75	60.08	42.05	75.72	72.88	60.70	52.66	65.80	62.03	48.43	64.01	60.31
	2	59.95	69.05	62.69	46.76	71.59	77.87	67.26	46.41	69.17	62.59	45.51	61.91	61.73
	3	59.13	55.69	61.96	44.80	55.34	72.98	55.22	58.77	61.57	63.90	47.76	58.98	58.01
	4	56.43	41.58	52.77	54.36	49.77	63.19	45.31	54.00	51.35	63.11	46.22	53.72	52.65
	5	66.03	65.98	69.52	56.98	70.62	76.30	63.08	68.32	67.37	71.14	53.12	58.43	65.57
	6	64.19	58.75	63.27	48.93	62.96	75.88	61.79	56.18	60.33	66.47	50.45	54.97	60.35
	7	53.29	59.14	53.89	45.32	64.58	71.26	65.47	45.70	60.99	65.09	49.98	58.49	57.77
	8	52.23	45.31	52.86	52.30	56.82	60.65	54.31	61.59	58.43	60.64	51.43	49.09	54.64
	9	60.64	72.23	62.85	49.89	58.22	76.83	61.03	57.43	77.62	71.42	59.46	60.94	64.05
	10	59.67	62.55	64.04	57.62	58.70	75.54	70.99	50.39	66.17	74.15	54.27	58.55	62.72
	11	52.03	30.06	47.69	30.53	41.49	60.51	49.30	39.93	46.64	46.57	42.41	47.26	44.54
	12	54.90	56.75	55.42	46.45	55.56	70.37	50.05	42.89	55.03	60.06	42.19	64.25	54.49
Avg.		58.09	56.49	58.92	48.00	60.11	71.19	58.71	52.86	61.70	63.93	49.27	57.55	

Table 5: Cross-sector test with models trained on the same size of each sector (AUPRC, %).

Model	AUPRC (%)
LSTM	83.65
Bi-LSTM	83.38
ELECTRA Base	84.92
BERT Base	85.13

Table 6: Classification performance of baseline models; LSTM, Bi-LSTM, ELECTRA Base, and BERT Base.

the precision-recall curve (AUPRC) of causal sentences as the evaluation metric (Davis and Goadrich, 2006). Models are optimized for the validation dataset. A trained model gives us the probability of causality on each sentence. Experiments are conducted on Google Colab Pro with one Tesla V100-SXM2-16GB GPU and four Intel(R) Xeon(R) CPU @ 2.00GHz CPUs.

4.2 Experimental Results on Baseline Models

We use LSTM, Bidirectional LSTM (Bi-LSTM) (Graves et al., 2005), ELECTRA Base (Clark et al., 2020) and BERT Base as baselines to compare the performance of extracting rationale of financial performance changes. As an input, each sentence is tokenized by ELECTRA and BERT (Devlin et al., 2018) tokenizers respectively for ELECTRA Base and BERT Base. For the LSTM models, we use Glove (Pennington et al., 2014) word embedding.

For the baseline experiments, we use total CR-COPEC dataset described in Section 3.4.1. As shown in Table 6, models based on transformers (ELECTRA Base and BERT Base) performed better than RNN based models. Because BERT Base achieved the highest AUPRC score (85.13%), we utilize BERT Base for followed experiments.

4.3 Examination of Sectors’ Characteristics

CR-COPEC consists of twelve different industries. We observe that each industry possesses its own field of interests with respect to reasons of financial performance changes. Therefore, training each sector individually is required for a better classification. However, the number of data for each indus-

try varies around [5k-60k], in which some sectors do not contain enough data to detect causal rationale precisely. Thus, this section tests a hypothesis if training other sectors together helps detecting causal sentences in a target sector. For this purpose, we conduct cross-sector test and compare the performance of models trained on various versions of the CR-COPEC dataset.

First, we fine-tune BERT Base on the equal size of fraction dataset described in Section 3.4.2 and test across sectors. We call this model as Sector-only Same-size. Then, we select neighbor sectors based on the results from the cross-sector test with valid dataset (Table 5). If the models trained on a neighbor sector detect target sector sentences better than the model trained on target sector, we assume that this sector can support to train a model. For example, in the case of Sector 8, we select Sector 5 as a neighbor which shows 68.32% in AUPRC (> 61.59%). After we select neighbors for each sector, we fine-tune BERT Base with the target and all selected neighbor sector datasets. We define this model as Neighbor model. Besides, for equal size comparison, we randomly select data from Total dataset to the same size of each Neighbor model trained on. The model trained on this dataset is called Total-random. In addition, we also train model on the whole size of each sector dataset and call it as Sector-only. Finally, we name BERT Base trained on Total dataset from the Section 4.1 as Total-all. We compare the performances of these four different types of models on in details below.

From the results in Table 5, we also find that the average performance of models on Sector 6 is the highest (71.19%) and Sector 11 (49.27%) is the lowest. This result demonstrates a gap in the level of difficulty of detecting causality on target sectors. Detecting causal sentences in Sector 6 is relatively easier and on Sector 11 is more difficult than other sectors. Furthermore, Sector-only Same-size model trained on fraction dataset of Sector 5

Test performances of classification models on test data of each sector




Figure 3: Test performances of classification models on valid dataset of each sector (AUPRC). The size of dataset for Neighbor and Total-random is presented in parenthesis.

shows the highest average performances across all sectors. Meanwhile, the model trained on fraction dataset of Sector 11 shows the lowest average score. We interpret these results as the sector generality and specificity. It can be considered as Sector 5 consists of more general causality patterns that can support to other sectors easier, while Sector 11 contains sector specific information so harder to transferable to other sectors. Hence, we find that Sector 11 is difficult to find causal sentences and also hard to be transferable to other sectors. We assume that since Sector 11 is the sector of finance industry, its unique characteristics present the sector specificity in our dataset. However, this transferable interaction between sectors is asymmetric, so other sectors can easily be a help to Sector 11.

Figure 3 compares the performances of Sector-only, Total-random, Neighbor and Total-all models on each sector. Overall, Neighbor models show significant higher performances compared to Sector-only and Total-random. We conduct t-tests of Neighbor and Total-random to determine the statistical significance of the difference in performances of AUPRC ($p < .01$). In Sector 1 and 4, Neighbor is higher than Total-all, though the training dataset’s size is much smaller than Total dataset.

4.4 Effects of Dataset Sizes

We check the effect of the number of training data with BERT Base trained on CR-COPEC. We increase the amount of training data and verify the causal extracting performance on each step. At each step, we randomly select a set of training data by increasing 10,000 sentences, then report AUPRC on CR-COPEC test dataset. AUPRC of




Figure 4: AUPRC of BERT Base by the number of training data from CR-COPEC.

BERT Base increased with a huge gap (6.64%) from 10K to 20K. Then they increase gradually until 200K then slightly drop afterward. We provide the detailed results in Figure 4. Figure 3 shows that the gap between Total-all and Neighbor is bigger when the number of training data is relatively smaller than Total dataset, in case of Sector 9, 10 and 12.

4.5 Explaining Detected Causal Rationales

We analyze input features used for causal rationale classification. Local Interpretable Model-agnostic Explanations (LIME) (Ribeiro et al., 2016) is an explanation method that makes the predictions of a classifier interpretable by learning a linear interpretable model locally around the prediction. With this explanation technique, we can visualize the most important features affecting the prediction. To see the difference on the predicted sentences from various models, we applied LIME on our trained models (Figure 5).

The probability of a causal sentence from Sector 1 is predicted as highest with Total-all, compared to Sector-only, Total-random and Total-all models. The Neighbor model also accurately predicts the correct answer (> 0.5). Furthermore, the

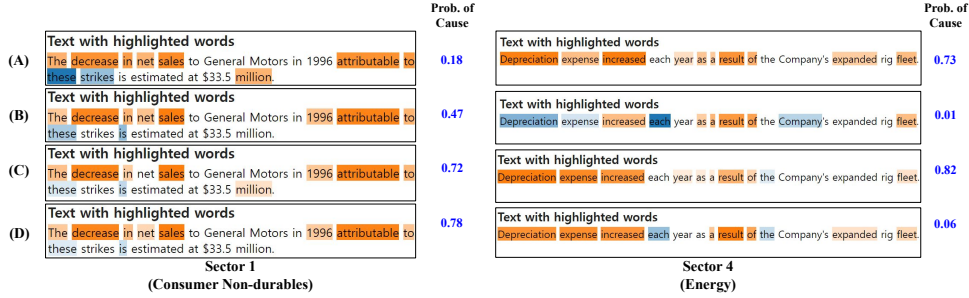


Figure 5: Interpretation of predictions from (A) ‘Sector-only’, (B) ‘Total-random’, (C) ‘Neighbor’ and (D) ‘Total-all’ models with LIME. Features contributing on causal rationale are highlighted in orange, the opposites are in blue.

word ‘strike’ works as negative to causal rationale in Sector-only model. However, it becomes neutral in models trained on more data.

In the example of Sector 4, sector specific models including Sector-only and Neighbor correctly predict the sentence as causal. Meanwhile, Total-random model regard the term ‘depreciation expense’ as negative one at prediction. This is because sentences containing causal rationale of ‘depreciation expense’ are annotated differently across sectors, so training all sectors together brings confusion in this case.

4.6 Comparisons with an Related Dataset

In order to verify the novelty of CR-COPEC, we compare our data with the existing financial causality detection corpus. For this, we conduct cross-dataset experiments between CR-COPEC and the FinCausal task 1 (Mariko et al., 2020). In specific, we compare the test performances in a cross-dataset setting where we fine-tune BERT Base on 1) FinCausal practice dataset, then test on both our test dataset and FinCausal trial dataset, and 2) CR-COPEC train dataset, then test on both of test dataset and FinCausal trial dataset.

Table 7 show that the performance of each model trained on a specific dataset is obviously decreased in its counterpart. This indicates that there are significant differences between the two datasets and the existing dataset cannot fully address the causality detection in 10-K reports and vice versa. We believe this is owing to the purpose and the target text of the datasets being different. Specifically, the CR-COPEC dataset aims to extract causal rationale sentences based on accounting items that are able to cause companies’ performance changes. Besides, CR-COPEC targets formal public text. On the other hand, FinCausal targets relatively casual text since FinCausal dataset tries to elicit comprehensive causality texts in financial documents.

Train \ Test	CR-COPEC	FinCausal
CR-COPEC	85.13	12.81
FinCausal	9.75	69.94

Table 7: The experimental results of the cross-dataset between CR-COPEC and FinCausal dataset (AUPRC).

	1	2	3
Longformer (AUPRC)	84.50	84.44	84.85

Table 8: The performances (AUPRC) on the difference size of input sentence n-grams ($n = 1, 2, 3$).

4.7 Multiple Sentence Modeling

Since the evidence for a causal statement can be scattered across more than one sentence, we conducted additional experiments on sentence n-gram settings to provide context for causal sentences. Each sentence n-gram consists of one target sentence and n-1 previous sentences, which serve as the context. Moreover, by following Mariko et al. (2020)’s setting, we conducted experiments for cases where $n = 1, 2, 3$, that the length of the context can have up to two sentences. Finally, since the number of tokens in sentence n-grams often exceeds the maximum token input size of BERT (512 tokens), all experiments were performed with Longformer Base (Beltagy et al., 2020), which accepts longer inputs (4,096 tokens). Appendix F provides detailed descriptions of the experimental settings. Experimental results of Table 8 shows that the performance of the multi-sentence (bi- and tri-gram) modelings achieves slightly higher performance at tri-gram modeling setting than single the uni-gram modeling but it is not significant ($p > 0.2$). The result implies that, in many cases, the evidences for a causal statement can be found within one sentence.

4.8 Experimental Results on a Large Language Model

This section reports the experimental results on a state-of-the-art large language model (ChatGPT; GPT-3.5-Turbo) (Ouyang et al., 2022) with our baseline model (BERT Base). For this, 1,000 instances were randomly sampled from the test set.

Model	Voting	F1 Score
BERT Base	-	0.870
GPT-3.5-Turbo Zero Shot	1	0.698
GPT-3.5-Turbo Zero Shot	5	0.714
GPT-3.5-Turbo 5-Shot	1	0.734
GPT-3.5-Turbo 5-Shot	5	0.747

Table 9: Classification performance of 1,000 sampled test set on our baseline and GPT-3.5-Turbo. Voting indicates the number of voters participating in inference based on majority voting. When voting is 1, it means single inference. Otherwise, when voting is 5, it means the majority voting of 5 times inferences.

The experiments were performed in the prompting-based zero-shot inference (Zero Shot) and 5-shot inference (Brown et al., 2020) where 5 positive and negative examples from the same sector as the input text were randomly selected from the training set. In addition, we conducted experiments on both of single inference and majority voting five inferences by following Arora et al. (2022)’s setting. Detailed experimental settings and prompts are described in Appendix G.

Table 9 shows that ChatGPT can successfully carry out a specific part of causal rationale text analysis without additional fine-tuning or zero-shot in-context learning. Additionally, 5-shot settings outperform zero-shot settings approximately 3%p. Furthermore, it has been demonstrated that employing majority voting surpasses the single inference-based performance. However, those performances are still significantly lower ($p < 10^{-4}$) than that of the fine-tuned smaller model; BERT Base. This is because ChatGPT’s results are sensitivity to minor fluctuations in financial condition or change in minor accounting items within (Interest income in the Sector 1) the company’s industrial sector. These results imply the need for an annotated corpus based on the domain knowledge on the specialized field even in the large language model’s era.

4.9 Potential Uses of Dataset

Financial analysts are key information intermediaries in capital markets, with their research, focused on uncovering private information and interpreting public data, being highly valued by investors. The value of analyst research can stem from two main sources: uncovering private information and interpreting public information, as exemplified by studies such as (Ivković and Jegadeesh, 2004; Asquith et al., 2005). To uncover this private information and reinterpret public information, analysts often analyze various linguistic

patterns in annual reports. Given the sheer volume of corporate filings today, not only are they paying significant time and money to analyze these reports, but their relying on manual annotation also reduces the accuracy of the analysis. Otherwise, NLP models trained with our dataset can automate annotations enhancing both speed and scalability. Moreover, analysts typically specialize in specific areas, requiring domain knowledge in industry specifics and market patterns. This includes understanding industry dynamics, competitive landscape, and regulatory environment, along with learning unique financial metrics and varying accounting practices across industries. Analysts may need to grasp new metrics and varying accounting practices across industries to understand how financial statements are prepared and analyzed, allowing them to make meaningful comparisons. Our dataset is expected to significantly reduce the learning curve and associated costs for analysts transitioning between industries.

5 Conclusion

We introduce a novel large scale dataset to extract causal rationales from financial reports in various sector. The dataset was annotated by considering accounting factors according to the industry. We further validate the process of building our CR-COPEC dataset. Finally, through the qualitative and quantitative analysis, we observed that the model trained with CR-COPEC recognize the clue of causal sentences. We hope that our work would promote the study of causality detection in the financial text domain.

Acknowledgement

Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No.2022-0-00984, Development of Artificial Intelligence Technology for Personalized Plug-and-Play Explanation and Verification of Explanation, No.2019-0-00075, Artificial Intelligence Graduate School Program (KAIST), NO.2022-0-00184, Development and Study of AI Technologies to Inexpensively Conform to Evolving Policy on Ethics).

Ethical Consideration

The main ethical concern in this work is a data license and releasing issue of 10-K database. First of all, 10-K database is publicly available, so a

license is not required to access the database. In addition, according to information presented on www.sec.gov is considered public information and may be copied or further distributed by users of the web site without the SEC's permission. Secondly, to all annotators, we paid them by obeying the minimum wage standard of national law. Besides, the general annotators were full-time employees.

Limitation

CR-COPEC may include two types of bias: 1) protocol bias, 2) annotators' subjectivity bias. First of all, we clearly defined in this paper that we annotate sentences based on annotation rules regarding financial statement items. Thus, our dataset may have different criteria from other causality detection problems (e.g, FinCausal dataset). Secondly, we are aware that the subjective judgment of annotators may exist in the annotation process. Therefore, inconsistency of annotations between annotators may occur for some sentences, which may affect data quality degradation and model performance. To control this issue, we verified the human annotation process (Section 3.3.2) and obtained substantial agreement results by measuring the kappa score between general annotators and additional annotators for 1000 sample sentences. The reliability of each general annotator is 0.7 or higher, and the agreement between general annotators is 0.698.

CR-COPEC is designed to detect causal statements that can limit further analyses. In the future, those without considering the positive or negative impact of performance changes. Additional analysis based on subcategories such as positive/negative/neutral on the impact of causal statements on a corporate performance changes is required in the future.

References

- Wajid Ali, Wanli Zuo, Rahman Ali, Xianglin Zuo, and Gohar Rahman. 2021. Causality mining in natural languages using machine and deep learning techniques: A survey. *Applied Sciences*, 11(21):10064.
- R. Alkadhi, T. Lata, E. Guzman, and B. Bruegge. 2017. Rationale in development chat messages: An exploratory study. In *2017 IEEE/ACM International Conference on Mining Software Repositories*, pages 436–446.
- Simran Arora, Avanika Narayan, Mayee F Chen, Laurel Orr, Neel Guha, Kush Bhatia, Ines Chami, and Christopher Re. 2022. Ask me anything: A simple strategy for prompting language models. In *The Eleventh International Conference on Learning Representations*.
- Paul Asquith, Michael B Mikhail, and Andrea S Au. 2005. Information content of equity analyst reports. *Journal of financial economics*, 75(2):245–282.
- Iz Beltagy, Matthew E Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*.
- Eduardo Blanco, Núria Castell, and Dan Moldovan. 2008. Causal relation extraction. In *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC'08)*, pages 310–313.
- Steven L. Bressler and Anil K. Seth. 2011. Wiener-granger causality: A well established methodology. *NeuroImage*, 58(2):323 – 329.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Du-Seong Chang and Key-Sun Choi. 2006. Incremental cue phrase learning and bootstrapping method for causality extraction using cue phrase and word pair probabilities. *Information Processing & Management*, 42(3):662 – 678.
- Malvika Chhatwani. 2022. Does robo-advisory increase retirement worry? a causal explanation. *Managerial Finance*.
- Kevin Clark, Minh-Thang Luong, Quoc V Le, and Christopher D Manning. 2020. Electra: Pre-training text encoders as discriminators rather than generators. *arXiv preprint arXiv:2003.10555*.
- Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and psychological measurement*, 20(1):37–46.
- Jesse Davis and Mark Goadrich. 2006. The relationship between precision-recall and roc curves. In *Proceedings of the 23rd international conference on Machine learning*, pages 233–240.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: pre-training of deep bidirectional transformers for language understanding. *CoRR*, abs/1810.04805.
- Luisse Dürlich, Sebastian Reimann, Gustav Finnveden, Joakim Nivre, and Sara Stymne. 2022. Cause and effect in governmental reports: Two data sets for causality detection in swedish. In *Proceedings of the LREC 2022 workshop on Natural Language Processing for Political Sciences*, pages 46–55.

- Mahmoud El-Haj, Paul Rayson, Steve Young, Andrew Moore, Martin Walker, Thomas Schleicher, and Vasiliki Athanasakou. 2016. Learning tone and attribution for financial text mining. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 1820–1825.
- Mauro Fonseca, Fernando Delbianco, Ana Maguitman, and Axel J Soto. 2023. Assessing causality among topics and sentiments: The case of the g20 discussion on twitter. *Journal of Information Science*, page 01655515231160034.
- Kenneth R French. 2019. Kenneth r. french - detail for 12 industry portfolios. https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/Data_Library/det_12_ind_port.html.
- Roxana Girju and Dan I Moldovan. 2002. Text mining for causal relations. In *Proceedings of the Fifteenth International Florida Artificial Intelligence Research Society Conference*, pages 360–364.
- Roxana Girju et al. 2007. Semeval-2007 task 04: Classification of semantic relations between nominals. In *SemEval 2007*, pages 13–18.
- Alex Graves, Santiago Fernández, and Jürgen Schmidhuber. 2005. Bidirectional lstm networks for improved phoneme classification and recognition. In *International Conference on Artificial Neural Networks*, pages 799–804.
- Yuxian Gu, Xu Han, Zhiyuan Liu, and Minlie Huang. 2022. Ppt: Pre-trained prompt tuning for few-shot learning. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8410–8423.
- Harsha Gurulingappa et al. 2012. Development of a benchmark corpus to support the automatic extraction of drug-related adverse effects from medical case reports. *Journal of biomedical informatics*, 45(5):885–892.
- Iris Hendrickx et al. 2019. Semeval-2010 task 8: Multi-way classification of semantic relations between pairs of nominals. *arXiv preprint arXiv:1911.10422*.
- Yunseong Hwang, Anh Tong, and Jaesik Choi. 2016. Automatic construction of nonparametric relational regression models for multiple time series. In *Proceedings of the International Conference on Machine Learning*, volume 48, pages 3030–3039.
- Ashwin Ittoo and Gosse Bouma. 2011. Extracting explicit and implicit causal relations from sparse, domain-specific texts. In *International Conference on Application of Natural Language to Information Systems*, pages 52–63.
- Zoran Ivković and Narasimhan Jegadeesh. 2004. The timing and value of forecast and recommendation revisions. *Journal of Financial Economics*, 73(3):433–463.
- Ammireddy Karuna. 2019. [Here is what’s behind the ai revolution](#). *Forbes Technology Council*.
- Shimon Kogan, Dmitry Levin, Bryan R. Routledge, Jacob S. Sagi, and Noah A. Smith. 2009. Predicting risk from financial reports with regression. In *Proceedings of The Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 272–280.
- Sunjae Kwon, Zonghai Yao, Harmon Jordan, David Levy, Brian Corner, and Hong Yu. 2022. [MedJEx: A medical jargon extraction model with Wiki’s hyperlink span and contextualized masked language model score](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. Albert: A lite bert for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942*.
- Yan Liang, Ying Liu, Chun Kit Kwong, and Wing Bun Lee. 2012. Learning the “whys”: Discovering design rationale using text mining — an algorithm perspective. *Computer-Aided Design*, 44(10):916–930.
- Jinghui Lu, Linyi Yang, Brian Namee, and Yue Zhang. 2022. A rationale-centric framework for human-in-the-loop machine learning. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6986–6996.
- Dominique Mariko, Hanna Abi Akl, Estelle Labidurie, Stephane Durfort, Hugues De Mazancourt, and Mahmoud El-Haj. 2020. The financial document causality detection shared task (fincausal 2020). In *Proceedings of the 1st Joint Workshop on Financial Narrative Processing and MultiLing Financial Summarization*, pages 23–32.
- Houman Mehrafarin, Sara Rajaei, and Mohammad Taher Pilehvar. 2022. On the importance of data size in probing fine-tuned models. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 228–238.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. [Glove: Global vectors for word representation](#). In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543.

- Rashmi Prasad et al. 2008. The penn discourse treebank 2.0. In *LREC*.
- Sampo Pyysalo et al. 2007. Bioinfer: a corpus for information extraction in the biomedical domain. *BMC bioinformatics*, 8:1–24.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. ” why should i trust you?” explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 1135–1144.
- B Rogers, James Gung, Yechen Qiao, and Janet E Burge. 2012. Exploring techniques for rationale extraction from existing documents. In *Proceedings of the International Conference on Software Engineering*, pages 1313–1316.
- Hiroyuki Sakai, Hiroko Nishizawa, Shogo Matsunami, and Hiroki Sakaji. 2015. Extraction of causal information from pdf files of the summary of financial statements of companies. *Transactions of the Japanese Society for Artificial Intelligence*, 30:172–182.
- Lloyd Shapley. 1971. Cores of convex games. *International Journal of Game Theory*, 1(1):11–26.
- Omri Suissa, Maayan Zhitomirsky-Geffet, and Avshalom Elmalech. 2023. Question answering with deep neural networks for semi-structured heterogeneous genealogical knowledge graphs. *Semantic Web*, 14(2):209–237.
- Anthony J Viera, Joanne M Garrett, et al. 2005. Understanding interobserver agreement: the kappa statistic. *Family medicine*, 37(5):360–363.
- Jinghang Xu, Wanli Zuo, Shining Liang, and Xianglin Zuo. 2020. A review of dataset and labeling methods for causality extraction. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 1519–1531.
- Jie Yang, Soyeon Caren Han, and Josiah Poon. 2022. A survey on extraction of causal relations from natural language text. *Knowledge and Information Systems*, 64(5):1161–1186.
- Wenbo Zhang, Tong Wu, Yunlong Wang, Yong Cai, and Hengrui Cai. 2023. Towards trustworthy explanation: On causal rationalization. In *Proceedings of ICML 2023*.

A Annotation Guideline

Main Category	Sub-category	Items	Sector											
			1	2	3	4	5	6	7	8	9	10	11	12
Income Statements	Sales	Operating revenue	1	1	1	1	1	1	1	1	1	1	1	1
		Service revenue	0	0	0	0	0	0	0	0	1	0	0	1
		Engineering service revenue	0	1	1	0	0	1	0	0	0	0	0	0
		Oil and gas oil and gas revenue	0	0	0	1	0	0	0	0	0	0	0	0
		Lease revenue	1	0	0	0	0	0	0	0	0	0	0	0
		Contract logistics revenue	0	0	0	1	0	0	0	0	0	0	0	0
	Cost	Sale cost	1	1	1	1	1	1	1	1	1	1	1	1
		Net sales	1	1	1	1	1	1	1	1	1	1	1	1
	Gross Profit	Gross profit	1	1	1	1	1	1	1	1	1	1	1	1
		Gross profit margin	1	1	1	1	1	1	1	1	1	1	1	1
		Research development expense	0	0	0	0	1	1	0	0	0	1	0	0
	Operating Expenses	Amortization expense	0	0	0	0	1	1	0	0	0	1	0	0
		Depreciation expense	0	0	1	1	0	0	0	0	0	0	0	1
		Legal and entitlement cost	0	0	0	0	0	0	0	0	0	0	1	0
		Expenditure associated with Environmental liabilities	0	0	0	0	0	0	0	0	0	0	0	1
		Operating Incomes	Operating income	1	1	1	1	1	1	1	1	1	1	1
	Non-operating Items	Non operating income (net) Interest expense	1	1	1	1	1	1	1	1	1	1	1	1
		Interest income	0	0	0	0	0	0	0	0	0	0	1	0
		Purchased services and other expenses	0	0	0	0	0	0	0	0	0	0	0	1
		Loan payment	0	1	1	0	0	0	0	0	0	0	0	0
		Net Income	Net income	1	1	1	1	1	1	1	1	1	1	1
		-	Manufacturer price	1	1	1	1	1	1	1	1	1	1	1
Balance Sheets	Liabilities	Debt	0	0	0	0	1	1	0	0	0	1	0	0
	Assets	Mortgage loan	0	0	0	0	0	0	0	0	0	0	1	0

Table 10: The annotation guideline for each sector. Herein, 1 and 0 indicate signals for *causal rationale of sentences* and *non-causal rationale of sentences* respectively. Note that **the guideline is not an absolute standard**, and it is **possible to flexibly annotate** if it is judged as causal sentences according to the annotator’s point of view. For items not included in this guideline, the causality of a target sentence is annotated by depending on annotators’ subjection.

B Main Factors

Main factors	Sector												
	Consumer non-durables	Consumer durables	Manufac.	Energy	Chemicals	Business equipment	Telephone	Utilities	Shops	Health	Finance	Others	
Material cost	■				■								
Exchange rate		■											
Sales volume		■	■										
Contract			■										
Production				■									
R&D Expense						■				■			
# of Customers							■	■					
# of Shops									■				
Interest income											■		
Service revenue												■	

Figure 6: Examples of each sector’s main factors.

C Additional Examples of Causal and Non-causal rationale of sentences

Sector of Industry - Example [Document]
Consumer Non-Durables - The gross profit margin as a percentage of sales improved from 44.3% in Fiscal Year 1995 to 46.8% in Fiscal Year 1996, principally due to lower green coffee and material costs, and lower plant overhead costs. [Brothers Gourmet Coffees, Inc., July, 1997]
Consumer Durables - Other income (expense), net increased primarily as a result of foreign currency exchange gains. [Breed Technologies, Inc., September, 1997]
Manufacturing - The increase in 1996 net sales was due primarily to increases in sales revenues recognized on the contracts to construct the first five Sealift ships, the Icebreaker and the forebodies for four double-hulled product tankers, which collectively accounted for 63% of the Company's 1996 net sales revenue. [Avondale Industries, Inc., March, 1997]
Energy - Oil and gas sales decreased due to a decrease in production volumes. [Castle Energy Corp., December, 1997]
Chemicals - Revenues decreased eight percent to \$3,955 million in 1996 from \$4,282 million in 1995 primarily reflecting lower SM prices and, to a lesser extent, lower volumes for PO and derivatives. [Arco Chemical Co., February, 1997]
Business Equipment - Research and development expenses increased from \$10.1 million (2% of net revenues) in fiscal 1995 to \$34.6 million (21% of net revenues) in fiscal 1996 due to the increase in Software development resulting from the acquisition of the three Software studios in calendar 1995. [Acclaim Entertainment, Inc., November, 1997]
Telephone - Revenue from cable television operations increased by \$90,713 or 24.6%, over the corresponding year ended May 31, 1996 as a result of regulated price increases, increases in the number of cable television subscribers and acquisitions. [Century Communications Corp., August, 1997]
Utilities - Gas operating revenues increased \$36.7 million, or 21.0%, due to increased volumes as a result of customer growth and higher gas costs. [WPS Resources Corp., March, 1997]
Shops - Aggregate sales generated by franchised stores increased by \$10,001,000, or 13.1%, to \$86,485,000 for calendar year ended December 31, 1995, as compared to \$76,484,000 for the same period in 1994, due to an increase of the number of franchised stores, as well as higher sales volume per store. [Sterling Vision, Inc., April, 1997]
Health - The increase in research and development expenses in 1996 and 1995 was due primarily to higher expenditures for the Actiq Cancer Pain Program, new product development and other expenditures for product development, including clinical trials. [Anesta Corp., March, 1997]
Finance - The investment income decrease resulted from a declining asset base, in large part resulting from loan repayments.
Others - The Company's largest revenue source is from the marketing and administration of extended vehicle service contracts ("VSCs") under the EasyCare(R) name, which provided 99% of revenues for 1996. [Automobile Protection Corp-AUPRCCO, March, 1997]

Table 11: Additional examples of causal rationale of sentences by each sector.

Section of Industry - Example [Document]
Consumer Non-Durables - Total general and administrative costs decreased by \$79,000 in 1995 due primarily to the absence of a management fee for 1995. [Highwater Ethanol, LLC, January, 2017]
Consumer Durables - Bank borrowings during 1995 were attributable to the Silver Furniture acquisition and the refinancing of Silver Furniture's bank indebtedness. [Chromcraft Revington, Inc., March, 1997]
Manufacturing - Because components are sold directly to the Company's manufacturing sources, the Company is not aware of the precise quantities sourced from particular suppliers. [Fossil, Inc., March, 1997]
Energy - Due to the apparent age of the material, no fine or enforcement action is expected. [Arabian Shield Development Co, March, 1997]
Chemicals - Due to personnel additions to the department, employee wages increased approximately \$56,700 in 1996. [American Vanguard Corp., March, 1997]
Business Equipment - Government contracts are subject to negotiated overhead rates, and work performed under government contracts is subject to audit and adjustments of amounts paid to the Company. [Ibis Technology Corp., 1997]
Telephone - Cost of services related to the wireless telephone operations during the year ended May 31, 1996 was \$26,129, an increase of \$3,977 or 18.0% as compared to the year ended May 31, 1995. [Century Communications Corp., August, 1997]
Utilities - The remainder of the increase was attributable to increases in ad valorem taxes, repair and maintenance expense mainly related to the WCLSF and the employee incentive plan which rewards certain of Tejas' employees with bonuses when the company achieves certain annual financial growth targets. [Tejas Gas Corp., March, 1997]
Shops - The Board may increase or decrease the number of shares under the Program or terminate the Program in its discretion at any time. [Boise Cascade Co., February, 2017]
Health - The 1995 results were also negatively impacted by a reduction of the Company's income tax benefit resulting from reserves established related to the expiration of certain state operating losses. [American White Cross Inc., April, 1997]
Finance - In the last three years, inflation has not had a significant impact on the Company because of the relatively low inflation rate. [Weeks Corp., March, 1997]
Others - In addition, the timing of revenue is difficult to forecast because the Company's sales cycle is relatively long. [Claremont Technology Group Inc., September, 1997]

Table 12: Examples of non-causal rationale of sentences by each section.

D Details on Annotation Environment

A	B	C	D	E	F
sentence	section	document	date	loc	true_label
ITEM 7.	1	129598108	19970331	0	no
MANAGEMENT'S DISCUSSION AND ANALYSIS OF FINANCIAL CONDITION AND RESULTS OF OPERATIONS Results of Operations Overview On March 31, 1996, Calgene and Monsanto Company ("Monsanto") entered into a transaction under which Monsanto contributed Gargiulo Inc. ("Gargiulo"), \$30 million cash and certain oils and produce related technology in exchange for a 49.9% equity interest in Calgene.	1	129598108	19970331	1	no
Gargiulo is a grower, packer, marketer and distributor of tomatoes, strawberries and other produce.	1	129598108	19970331	2	no
Gargiulo tomato producing operations are conducted principally in Florida, California, Puerto Rico and Mexico.	1	129598108	19970331	3	no
Gargiulo berry production operations are conducted principally in northern California.	1	129598108	19970331	4	no
On February 29, 1996, Gargiulo and Collier Enterprises consummated an asset purchase agreement whereby Gargiulo acquired substantially all the assets subject to the assumption of certain specified liabilities of the produce business conducted by certain affiliates of Collier Enterprises under the trade name Collier Farms ("Collier").	1	129598108	19970331	5	no
Collier is an agricultural producer of tomatoes and other vegetables in Florida, and engages in the packaging, marketing and distribution of those products in the commodity markets.	1	129598108	19970331	6	no
The Company's tomato operations consist of the combined business of Calgene Fresh, which was organized in 1992 to develop and produce genetically engineered premium tomatoes, and Gargiulo.	1	129598108	19970331	7	no
Effective January 1, 1997, Calgene changed its fiscal year from June 30 to December 31.	1	129598108	19970331	8	no
Accordingly, the financial presentation in this report for 1996 is for the six month period from July 1, 1996, to December 31, 1996.	1	129598108	19970331	9	no
Prior fiscal year's operations are as previously reported and cover twelve month periods ended June 30, 1996, 1995 and 1994.	1	129598108	19970331	10	no
Revenues Calgene's product sales in the six month period ended December 31, 1996 increased 250% to \$69.8 million from \$19.9 million in the comparable period of the prior year.	1	129598108	19970331	11	no
The increase is due to the inclusion of Gargiulo's operations which resulted in higher fresh market produce sales of \$50.4 million.	1	129598108	19970331	12	yes
Product sales in fiscal 1996 increased 95.5% to \$95.7 million from \$49.0 million in fiscal 1995.	1	129598108	19970331	13	no
The increase resulted from the inclusion of Gargiulo's product sales of \$48.3 million in the fourth quarter of fiscal 1996.	1	129598108	19970331	14	no
Product sales in fiscal 1995 increased 38.3% to \$49.0 million from \$35.4 million in fiscal 1994.	1	129598108	19970331	15	no
The increase reflects \$6.8 million higher tomato sales, \$4.2 million higher cotton seed sales, and \$2.3 million higher specialty oleochemical sales.	1	129598108	19970331	16	no
Product development revenues in the six month period ended December 31, 1996 decreased by 14.2% to \$729,000 from \$850,000 in the corresponding period of the prior year due to the conclusion of several research contracts and a \$125,000 benchmark milestone payment which was recognized in the prior year.	1	129598108	19970331	17	yes
This decrease was partly offset by the addition of several new research contracts.	1	129598108	19970331	18	no
Product development revenues in fiscal 1996 increased by 43.6% to \$9.3 million from \$6.5 million in the prior fiscal year.	1	129598108	19970331	19	no
The increase was due to a \$3.3 million increase in technology license sales reflecting the net impact of non-recurring sales of \$7.0 million and \$3.8 million in fiscal 1996 and 1995, respectively.	1	129598108	19970331	20	no
Product development revenues in fiscal 1995 increased by 113.5% to \$6.5 million from \$3.0 million in the prior fiscal year.	1	129598108	19970331	21	no
The increase was primarily due to a \$3.4 million increase in technology license sales.	1	129598108	19970331	22	yes
Gross Profit Calgene's gross profit on net product sales was negative \$2.3 million in the six month period ended December 31, 1996 as compared to a negative gross profit of \$2.2 million in the comparable period of the prior year primarily due to the Company's produce operations.	1	129598108	19970331	23	yes
Tomato production yields were reduced due to adverse weather conditions in several growing regions, and gross profit on strawberry production was negatively impacted by exceptionally low prices experienced by the industry.	1	129598108	19970331	24	yes
Calgene's gross profit on net product sales was \$5.3 million in fiscal 1996 as compared to a negative gross profit of \$4.7 million in fiscal 1995.	1	129598108	19970331	25	no
The \$10.0 million improvement is attributable to a \$10.4 million reduction of negative gross profit at Calgene Fresh, and the inclusion of Gargiulo fourth quarter gross profit of \$2.1 million.	1	129598108	19970331	26	no
This improvement was partly offset by a \$2.1 million decrease in gross profit from cotton operations primarily due to surplus inventory write downs and higher product costs.	1	129598108	19970331	27	no

Figure 7: Illustrative example of the annotation form.

Annotation files are based on the Comma-Separated Values (CSV) format. In this case, all text in each report of a corporation is gathered in a single CSV file. The file consists of 6 columns: sentence, section, document, date, location (loc), and true_label. The sentence column contains the text information of the report that is split by each sentence. The section column indicates the SIC code of the corporation. The document column represents the unique document ID of the report. Date is the date the report was published. In addition, the location (loc) column signifies the sentence number of the text. Finally, true_label is the human annotated label for a sentence.

Each file was assigned to an annotator. The annotator coded sentences based on the annotation guidelines and the previous context. Specifically, for consecutive sentences related to the same accounting item, only the sentences that provide a reason for the performance change are labeled as causal rationale. For example, in the sentences at location 11, although product sales (which could be a causal rationale for performance change of a company with section 1) increased, it is difficult to consider the sentence as a causal rationale since the reason is not explained. On the other hand, sentences at location 12 are regarded as a causal rationale because the reason for the increment in product sales is specified and clear in the previous context. Note that, since the document and location information are also provided in the final version of the file, it is possible on modeling and inference for multiple sentences as in the experiments in section 4.7.

E Comparison Between FinCausal and CR-COPEC

In this section, we compare CR-COPEC with the most relevant dataset FinCausal (Mariko et al., 2020), specifically FinCausal task 1. Both of two datasets target to extract causality detection in the financial domain and a binary classification.

- Goal
 - FinCausal: If a text contains a causal scheme
 - CR-COPEC: If a sentence contains a causal rationale of changes in corporate financial performance.
- Data Source
 - FinCausal: 2019 financial news provided by Qwam
 - CR-COPEC: 1997 and 2017 10-K reports
- Data Size
 - FinCausal: 31,580 text section
 - CR-COPEC: 1,584 10-K reports with 283,490 sentences
- Class Bias
 - FinCausal: 7.24%
 - CR-COPEC: 3.93%

However, there are some fundamental differences between the two datasets that make our data unique. First of all, the Goal of the two datasets are different. FinCausal aims to find **a causal scheme in a text**, while our dataset targets to identify more specific information that is **a causal rationale of changes of financial performance**. It indicates that CR-COPEC solely concentrates on causal rationales of **accounting items considering the unique characteristics of various industries**. Secondly, FinCausal collected their data from financial news. Otherwise CR-COPEC utilizes 10-K reports. It means that **CR-COPEC is written in a formal manner** because they should comply with regulatory rules. Furthermore, the data size of CR-COPEC is larger than FinCausal. Finally, both of datasets are biased but the causal class distribution of the two datasets is different.

F Details on Multi-Sentence Modeling Experiments

This section explains in detail experimental settings for the multi-sentence modeling in Section 4.7. We first briefly introduce the sentence n-gram setting. Then, we describe how we constructed inputs for sentence n-grams. Note that, since sentence n-grams are often longer than the maximum input token size (512 tokens), we utilized Longformer Base, which can handle longer input tokens (4,096 tokens), for the experiments. We used the same hyper-parameter setting of BERT fine-tuning.

In this paper, sentence n-grams consist of 1 target sentence and n-1 previous context sentences. For example, let's suppose the i^{th} sentence (s_i) is a target sentence for which we want to predict whether it is a causal sentence or not. A sentence uni-gram only consists of s_i itself, which is the same as the general single sentence modeling. Otherwise, a sentence bi-gram is made up of s_i and the s_{i-1} context sentence. Similarly, a sentence tri-gram contains s_i and two context sentences, s_{i-1} and s_{i-2} .

$$s_1, s_2 \dots s_{i-2}, s_{i-1}, s_i$$

The target sentence and the context are concatenated and input to Longformer. Herein, we mark the boundary between the separate target sentence and context by utilizing Longformer's separator token $\langle /s \rangle$.

$$\{\textit{target sentence place holder}\} \langle /s \rangle \{\textit{context place holder}\}$$

G Details of the Experimental Setting on ChatGPT

Herein, since it is difficult to obtain a probability distribution in the prompting-based ChatGPT inference, we set the output as a binary format; 'yes' or 'no' whether the given input text contains causal rationale. The BERT Base model employs a threshold probability of 0.5. Finally, we set an evaluation metric as the F1 score known to be beneficial for evaluating imbalanced data. Furthermore, temperature is set to 0 for the consistent performance.