Information Extraction on sustainability reports using Transformer-based models

Enora Petry

Abstract

Sustainability reports have allowed companies to disclose information regarding environmental, social and governance (ESG) issues. These reports, along with other external data sources, allow rating companies to calculate sustainability scores for each firm. Sustainability reports are very diverse, mix pictures, graphics and text and are difficult to work with. We developed an extraction method that relies on Transformer-based model to evaluate each report on the most pressing key issues. The extracted scores show encouraging results when compared with ESG Scores. The research project has strengthened our understanding of how disclosure orient sustainability scores and has provided more transparency and more reliability about how sustainability scores are calculated.

Contents

Acknowledgments 4

1 Introduction 5

1.1 Background and Context 5

1.1.1 Sustainable development 5

1.1.2 ESG Metrics 5

1.1.3 ESG Reporting 6

1.1.4 Rating agencies 7

2 Literature Review **7**

3 Methodology **8**

3.1 Model selection 9

3.1.1 Zero-shot classification 10

3.2 Comparison to ESG Scores 11

4 Data **11**

4.1 Sustainability reports 11

4.2 ESG Scores 13

4.2.1 MSCI Methodology 13

5 Implementation **14**

5.1 Key Issues 14

5.1.1 Focusing on Environmental and Social key issues 14

5.1.2 Selecting of key issues 14

5.1.3 Translating key issues into topic description 15

5.2 Classifying 16

5.2.1 Pre-processing 16

5.2.2 Inference 16

5.2.3 Output 16

5.3 Processing outputs 16

5.3.1 Probability distribution 17

5.3.2 Aggregation function 18

6 Results **18**

6.1 Opportunities & similar topics 19

6.1.1 Outputs 19

6.1.2 Analysis 20

6.2 Metric scores 21

6.2.1 Outputs 21

6.2.2 Analysis 22

6.3 Other key issues 23

7 Next steps **24**

7.1 Limited sample size 24

7.2 Additional descriptors 25

7.3 Estimating pillar ESG Scores 25

8 Conclusion **26**

A Keywords **27**

Acknowledgments

I would like to acknowledge and give my warmest thanks to my supervisor Florian Berg who made this work possible. His insights and his guidance have guided me throughout this research experience.

I would also like to give thanks to my colleagues Francesco Tripoli and Cristan Foroni for their collaboration and teamwork on exploring MSCI ratings. Your work has certainly added to my understanding of this ESG metrics.

A special thanks to the Aggregate Confusion Project for welcoming me in their team. Their dedication and knowledge have inspired me throughout this project.

Finally, I would like to give a very large thank you to the research chair Finance Durable and Investissement Responsable, and its co-director Patricia Crifo. They have kindly provided financial support and have allowed me to pursue this amazing opportunity abroad.

1 Introduction

1.1 Background and Context

The growing concern for climate change has emphasized the need to shift to a more sustainable economy. It has highlighted the excessive consumption of the planet's resources and the need for a more sustainable development.

1.1.1 Sustainable development

Sustainable development is characterized by the development that meets the needs of the present without compromising the ability of future generations to meet their own needs [1]. Usually broken down into three pillars, sustainable development encompasses a variety of topics across economic, social and environmental issues.

Figure 1: Sustainable Development Venn Diagram

In the financial sector, this idea has been implemented with the concept of green finance or sustainable finance. Sustainable investing was developed in the 1970's under the name Socially responsible investing (SRI). It provided a mean for investors to align their portfolios to their values.

1.1.2 ESG Metrics

Societal impact is difficult to quantify. In order to evaluate the sustainability of corporations, environmental, social and governance (ESG) metrics appeared in the

1990's [2]. Since then, their influence has grown substantially, reaching trillions of dollars of ESG fund flows in 2020.

ESG metrics characterize the sustainability of firms across a large variety of subtopics, as illustrated by the table 2.

Environmental	**Social**	**Governance**
Climate Change	Workplace safety	Board composition
Resource Use	Fair wages	Lobbying
Toxic Emissions	Product Safety	Anti-competitive practices
Biodiversity and Land use	Access to underserved markets	Corruption and Fraud

Figure 2: Examples of ESG factors

1.1.3 ESG Reporting

In order to provide information about societal and planetary impact, companies disclose information through sustainability reports and other reporting frameworks. Sustainability disclosure has become widespread. In 2022, 98% of S&P 500 Companies, the largest publicly traded companies in the US, disclosed sustainability reports [3]. Through ESG reports, corporations provide transparency to stakeholders about how sustainability impacts the companies' decisions, and highlight opportunities and risks that might affect how the company is valued in the future.

Although there are increasing regulations regarding the disclosure of ESG information, the information shared in ESG reports remains at the discretion of the firms, leading to gaps in data and unstructured information.

1.1.4 Rating agencies

In order to orient their portfolio towards the most sustainable companies, investors rely on sustainability metrics. Rating agencies evaluate corporations' sustainable impact through ESG scoring. They assess and quantify that impact with publicly available and privately sourced data. [4]. ESG ratings are divergent across rating agencies (Berg, Kölbel, and Rigobon), which makes it difficult to evaluate the overall ESG performance of companies.

2 Literature Review

Investors and stakeholders use ESG reports to gather information about firms' risk and opportunities regarding sustainability. However, ESG documents are difficult to parse as they are long documents of textual information with no common format or layout. Numerous researchers have utilized Natural Language Processing techniques to summarize and aggregate data from corporate sustainable disclosures.

The most extensive ESG report analysis was done by Lin et al. They relied on word embeddings (Word2Vec) to create a ESG dictionnary from relevant keywords using a large corpus of over 210,000 annual reports. The dictionary allows them to analyse the evolution of specificity in reports over time and evaluate how sustainability disclosure has been shaped by disclosure regulations.

With the development of Transformer-based models, researchers have fine-tuned models to better decipher ESG-related sentences. Bingler et al. developed Climate-Bert, a fine-tuned Bert model to categorize sustainable related-sentences into four categories. However, they relied on training dataset curated manually from a subset of firms, which hinders scalability in the future.

Sustainability can't be analyzed in their totality with state-of-the-art Transformer models due to the limit of token size. Pasch and Ehnes choose to work with 10k forms instead, due to the more neutral stance in reporting. They fine-tuned BERT models for ESG sentiment analysis. By analysing each paragraphs and averaging the results, they obtained 72% accuracy in determining "good" and "bad" ESG companies.

Past research have explored a large variety of NLP models, from word-embeddings to Transformer models. However, the degree of granularity is very small, as paragraphs are sorted into four categories or firms are categorized using a binary good or bad scale. ESG Reports are very rich by the diversity of topics treated, according to each industry. In this research project, we aim to fill the gap by extracting information using industry-specific key topics. The additional level of granularity allows us to explore the influence the relationship between sustainability disclosures and ESG scores at the sub-topic level.

3 Methodology

In this research project, we developed a method that extracts key sentences from sustainability reports according to sector-specific topics. This allows to compute a ESG report disclosure score for every key issues which we encounter and enable the comparison with the sub-pillar ESG score. Thus, we can explore the relationship between ESG disclosure and ESG ratings for different topics.

3.1 Model selection

Large documents can be analyzed through a variety of text processing methods. For instance, main keywords can be extracted, relevant topics can be clustered or named entities can be located.

Our text data consists of sustainability reports. They are made of large subsets of textual data and graphics that disclose risks and opportunities in a variety of sectors. Topic analysis seemed necessary to divide the sustainability report in clusters and analyse their relationship with the sub-pillars of ESG Scores. In order to align the topics with the ESG metrics provided, we imposed pre-determined topics for text classification.

We decided to rely on Transformer-based models as they have been trained on very large datasets and have proved to efficiently capture the semantic meaning of sentences. Indeed, ESG Reports rely on specific vocabulary that is rapidly evolving and using hard-coded dictionnaries to identify relevant keywords is not adapted.

Although Large Language Models (LLMS) have shown remarkable performances in the recent years, the analysis we wish to produce is a fixed task and does not require the computational power demanded by LLMs. Futhermore, LLMs are prone to hallucinating.

The time-constraint of the project meant that it wasn't possible to create a specific ESG-related dataset and use fine-tuning. Furthermore, text classification with unseen labels, named zero-shot text classification, has shown promising results. We decided to use zero-shot classification for our text extraction.

3.1.1 Zero-shot classification

Zero-shot text classification is a natural language processing problems which aims to correctly label a piece of text with unseen categories.

Yin, Hay, and Roth propose to reframe zero-shot classification as a textual entailment problem, or natural language inference (NLI) problem. The main idea of natural language inference consists of looking at a text and evaluating the relationship of the text with a succeeding hypothesis. If the text allows the hypothesis to be true, we have a high entailement score.

Natural language inference can be used for textual inference with the following framework. Indeed, after looking at the text, the NLI model evaluates the probability of the following hypothesis being true: "This example is [topic] ?".

The entailement score, which captures the probability of the hypothesis being true given the text, can be interpreted as the probability of the text belonging to the topic.

Classification problem		Entailment problem
Probability of text being in topic	$\longrightarrow$	Given the text, is the following hypothesis true ? *text* : text *hypothesis* : "This example is topic"

Figure 3: Reframing text classification into natural language inference

Therefore, zero-shot text classification can be implemented using state-of-the art Natural Language Inference (NLI) models.

In this project, we use a NLI model called BART-mnli. BART is a Transformer-based model trained on removing noise from text to reconstruct the original data. The BART-mnli model is a pretrained checkpoint of the BART model trained on the MultiNLI dataset by William and al [10]. This dataset is one of the most extensive natural language inference dataset for training models on the task of textual entailment. BART has been shown to yield state-of-the-art results on zero-shot text classification [11].

3.2 Comparison to ESG Scores

After classifying every paragraph of the report, we would like to find an aggregation function that can transform the probability distribution of paragraphs into an overall disclosure score for the ESG report.

This disclosure score, calculated for every topic, can be compared with ESG Scores. We can then explore patterns in that relationship.

4 Data

4.1 Sustainability reports

For this project, 1,863 ESG reports were analysed. This consists of 909 different firms over three years (2020, 2021, 2022).

Sustainability reports are often in PDF format, and can range from 20 to 100 pages. They provide information regarding the corporation's priorities, policies that have been implemented and future targets regarding sustainability. They are primarily text, with added pictures, graphics and tables for quantitative data.

The sustainability reports are clustered according to the Global Industry Classification Standard into eleven sectors : *Communication Services, Consumer Discretionary, Consumer Staples, Energy, Financials, Health Care, Industrials, Information Technology, Materials, Real Estate, Utilities.*

Sector	2020	2021	2022	Total
Industrials	105	149	78	332
Financials	68	95	66	229
Information Technology	77	101	51	229
Consumer Discretionary	63	93	56	212
Materials	78	86	46	210
Health Care	49	66	45	160
Energy	64	71	23	158
Utilities	40	47	26	113
Real Estate	27	38	17	82
Consumer Staples	28	32	15	75
Communications Services	22	26	15	63
Total	621	804	438	1,863

Figure 4: Description of ESG Reports per sector

4.2 ESG Scores

The ESG metrics used were MSCI (Morgan Stanley Capital International) ESG Ratings. We decided on MSCI Ratings as it is one of the largest independent provider of ESG ratings and there scores are extensive and widely-used. The ratings look at 37 key issues, divided across three pillars.

4.2.1 MSCI Methodology

MSCI evaluates firms differently according to their sub-industry. They rely on their own sub-industry classifications to group companies into 77 clusters.

Key issues. Companies are evaluated on 5 to 7 key issues according to their sub-industry.

E, S, G Pillars. The key issues scores are aggregated into an Environmental Pillar Score, a Social Pillar Score and a Governance Pillar Score.

Weighted Average Score. The weighted sum of the three pillar scores forms a Weighted Average Score.

Industry Adjusted Score. The MSCI Methodology aims to compare similar firms to highlight the most sustainable companies relative to their industry. Thus, the Weighted Average Score is normalized by sub-industry to compute a Industry Adjusted Score. This Industry Adjusted Score determines the leading companies and laggard companies per sub-industry.

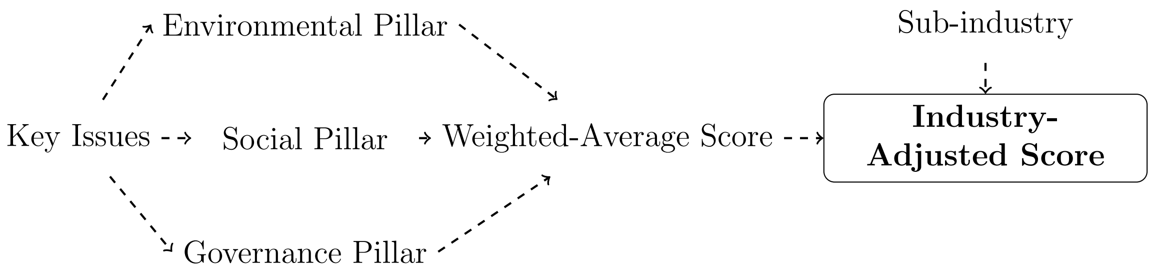

Figure 5: Overview of MSCI ESG Methodology

5 Implementation

Throughout the project, we grouped companies by GICS sector. Reducing the number of clusters from 77 MSCI Industries to 11 GICS Sector increased the sample size per cluster and facilitated the analysis of results.

5.1 Key Issues

5.1.1 Focusing on Environmental and Social key issues

The MSCI Methology for the Governance Key Issues uses a system of deductions to remove points from a perfect score. This methology is harder to translate into topic analysis on ESG reports. For this reasons, we restricted our analysis to the Environmental and Social Pillar.

5.1.2 Selecting of key issues

To determine the key issues, we selected the six most recurrent key issues for each sector. This allowed to select key issues present in the majority of sector's firms and excluded outliers.

In total, 22 key issues were selected. Some issues were common to all sectors such as "Carbon Emissions", whilst more than half the key issues were unique to the sector.

Key Issue	Sectors
Carbon Emissions	11
Human Capital Dev.	10
Corruption	6
Water Stress	5
Toxic Emission and Waste	4
Privacy and Data Security	4
Health and Safety	4
Product Safety and Quality	3
Anti-competitive practices	3
Business Ethics Fraud	2
Opps in Clean Tech	2

Key Issue (cont.)	Sectors
Product Carbon Footprint	2
Access To Finance	1
Access To Healthcare	1
Biodiversity and Land Use	1
Controversial Sourcing	1
Financing Env Impact	1
Finance Product Safety	1
Opps in Green Building	1
Opps in Nutrition and Health	1
Opps in Renewable Energy	1
Raw Material Sourcing	1

Figure 6: Overlap of key issues across sectors

5.1.3 Translating key issues into topic description

Each key issues is transformed into a topic description by selecting four to five keywords. These words are intended to better illustrate the relevant issues that would be described in the sustainability reports.

The keywords were manually extracted using the MSCI Key Issues documentation. The keywords used for the classification are described in Appendix A.

5.2 Classifying

5.2.1 Pre-processing

ESG reports were initially transcribed from PDF format to txt files. Each ESG report was split into paragraphs to most optimally respect the original PDF layout.

5.2.2 Inference

The classification was implemented using the `bart-large-mnli` pre-trained Bart model using the HuggingFace interface.

We use the parameter `multi-label = True` to independently calculate the probability of belonging to each topic. In fact, each paragraph can belong to multiple labels and we do not need the sum of all the probabilities to equal one.

The classification was parallelized on three T4 GPU and took 10 hours.

5.2.3 Output

For each paragraph, the text classification algorithm estimates the probability of the text fragment belonging in each of the six sector-specific key issues.

5.3 Processing outputs

After running the algorithm, we would like to aggregate scores to create a ESG report disclosure score for each key issue that evaluates the quality and quantity of the topic in that report.

5.3.1 Probability distribution

For every key issue, the probability distribution resembles a logarithmic distribution. Indeed, the majority of values are very close to zero which indicated that the retrieval method is very selective.

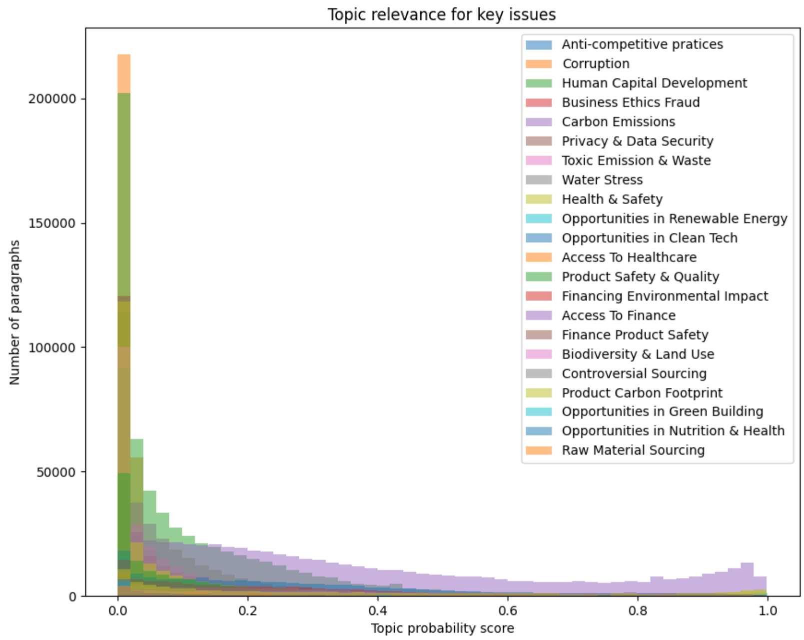

Figure 7: Topic Probability distribution

To evaluate the presence of relevant paragraphs, we want to focus on the few paragraphs with high probability. Indeed, selecting the top paragraphs allows us to avoid the noise from the large quantity of irrelevant information.

5.3.2 Aggregation function

Various aggregation function were explored, including selecting all paragraphs above a threshold, or looking at the top N most relevant paragraphs.

The method that seemed give the best results when comparing with ESG scores was the following : select the top 20% of paragraphs and average their topic probability scores.

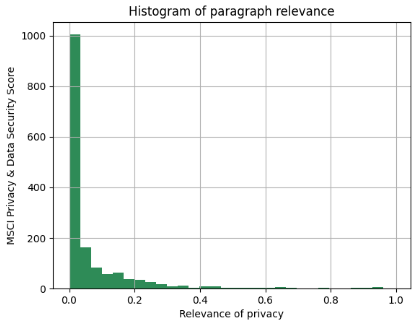

Score from ESG report:

0.28

Average probability score
of the top 20% of text paragraphs.

Figure 8: Aggregating to a single disclosure score

6 Results

To evaluate our calculated ESG disclosure scores were relevant, we chose to compare them with actual ESG scores. Our hypothesis is that higher ESG disclosure score would mean the company disclosed extensive information related to the ESG topic, and might result in higher ESG Scores.

Our results show a variety of patterns between the quantity and relevance of information shared by the firm and the resulting ESG score.

To evaluate the correlation, we plot the relationship between our calculated ESG disclosure score and the MSCI ESG score of the same Key Issue category. We cluster points of the same sector to understand how the sector influences the relationship.

6.1 Opportunities & similar topics

6.1.1 Outputs

Our results illustrated that there is a positive correlation for six key issues : Opportunities in Clean Tech, Opportunities in Nutrition and Health, Opportunities in Green Buildings, Opportunities in Renewable Energy, Access to Health Care and Product Carbon Footprint.

The resulting plots are very noisy due to the small sample size and the noisy nature of ESG scores. In fact, linear correlations methods do not converge except for the key issue Opportunities in Green Building.

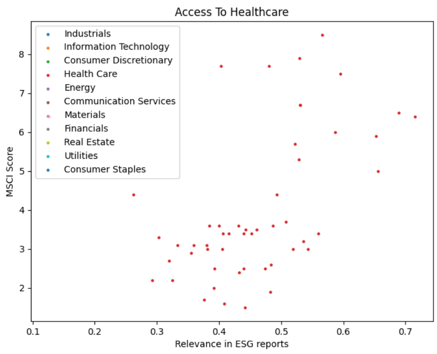

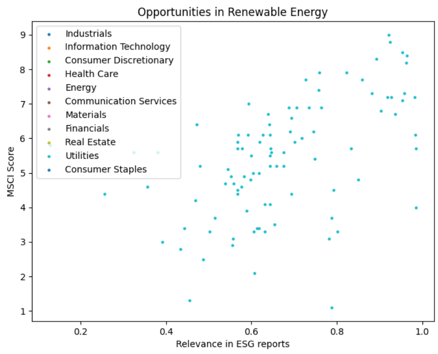

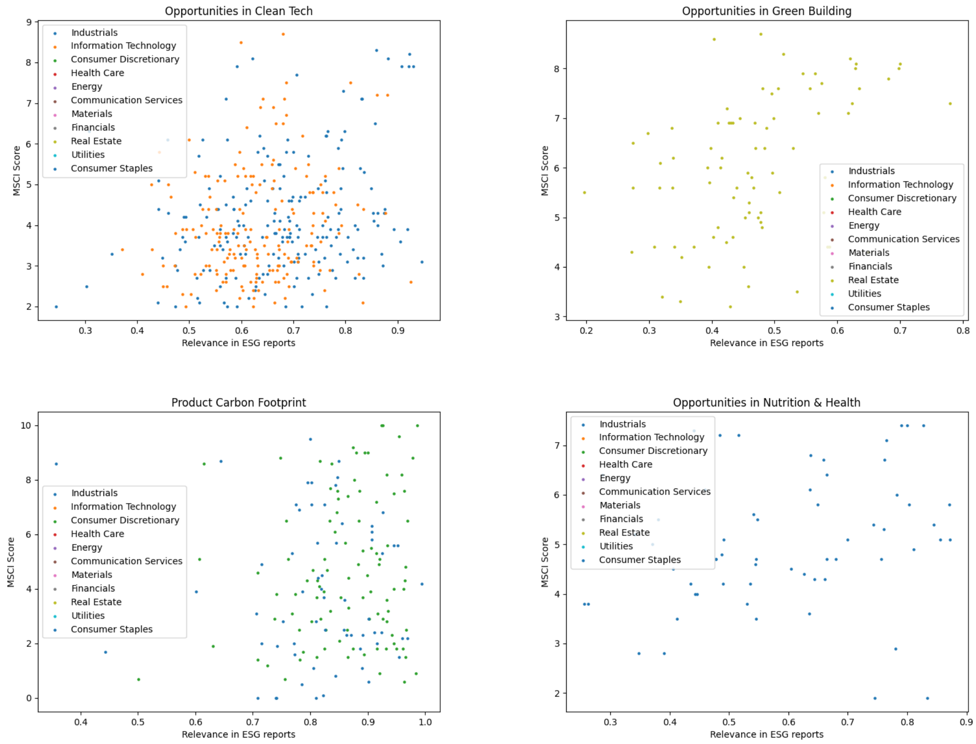

Figure 10: Scatter plots of disclosure score and ESG key issue score

6.1.2 Analysis

The correlation found in this analysis are encouraging as it illustrates there is a link between how much the sustainability reports mentions a topic and the resulting score.

These topics are some of the less complex key issues. Indeed, they are very one-dimensional and relevant paragraphs can easily be found using the keywords. This means that the precision and the recall for the paragraph extraction are very good. Secondly, four of the six topics relate only to opportunities, which is positively correlated with ESG scores. In other topics, we would find extensive reporting on other aspects such as risks and efforts to minimize them, which is less correlated with higher ESG scores.

6.2 Metric scores

6.2.1 Outputs

Another set of interesting plots are the two key issues that relate to emissions scores : Carbon Emissions and Toxic Emission & Waste.

Here the resulting graphs seem to indicate a negative correlation between our calculated ESG report score and the MSCI Scores. Companies that talks more extensively about these issues are also the companies that have the lower ESG scores.

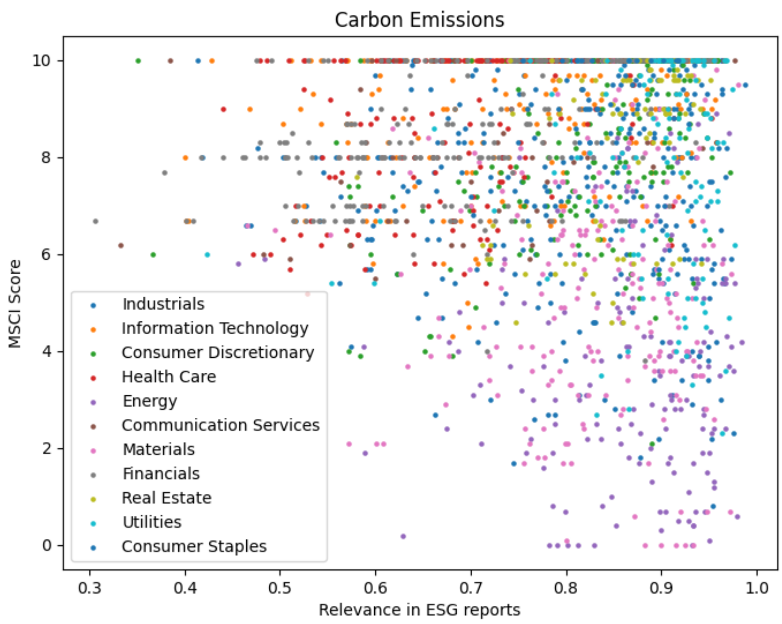

Figure 11: Scatter plots of Carbon Emissions

6.2.2 Analysis

The two key issues are mostly determined by emissions metrics. The correlation we observe seem to indicate that companies with lower scores seem compelled to disclose more information about the risks they encounter and the policies and decisions they are taking to manage their current situations.

If we plot sectors separately, we can observe some positive correlation between disclosing information and higher ESG scores.

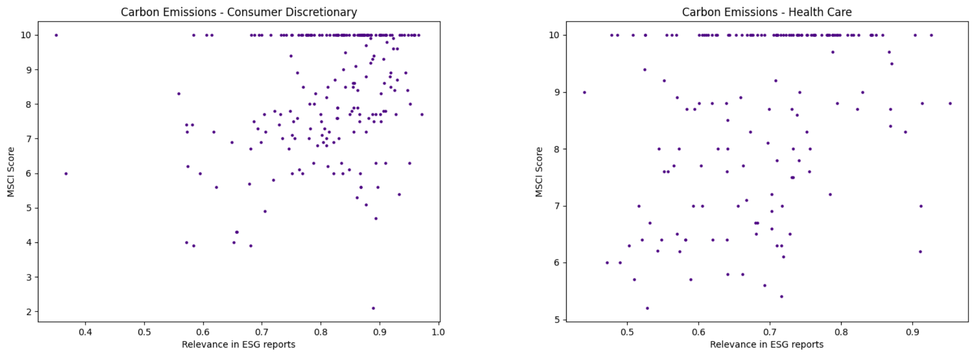

Figure 12: Sector-specific plots for Carbon Emissions

Additionnally, we can notice natural clusters in the scatter plots. For the Toxic Emission & Wate, sectors that obtain lower ESG score are the ones who disclose more information in ESG reports regarding this key issue.

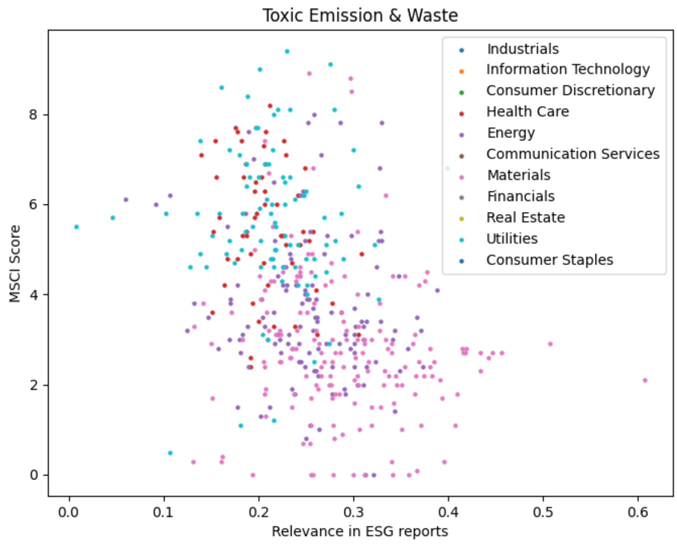

Figure 13: Scatter plots of Toxic Emissions & Waste

6.3 Other key issues

The other key issues do not yield interpretable patterns. This can be due to a variety of reasons. First, the keyword description does not select the relevant paragraphs in the sustainable report. Further research might be necessary to target other major issues that are most relevant to the topic.

Secondly, the ESG scores might rely on external data. For instance, two key issues appear to have discrete ESG scores. These are two Governance pillar topics : Anti-Competitive practices & Business Ethics Fraud.

The scatter plots illustrate that the data is discrete and calculated from deductions if the company has been flagged for these issues. ESG reporting does not seem to have an influence on the ESG score. This would indicate that the data is provided by another data source, such as a controversies dataset.

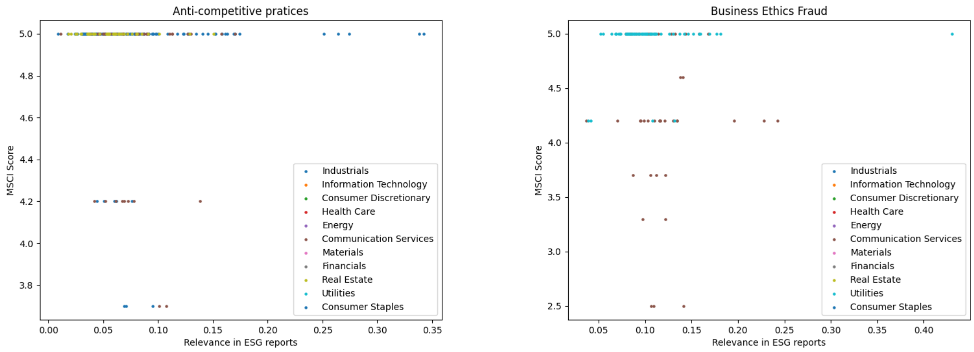

Figure 14: Discrete key issues scores

7 Next steps

The results from this analysis illustrate that the method developed in this research project does indicate some correlation between ESG disclosure and the resulting ESG scores. To futher explore this patterns, we would like to increase the sample size, add additional descriptors and explore external datasets.

7.1 Limited sample size

The sample size for this analysis was limited to reports in a time period of three years. This choice was due to limited time and computational resources. The time period will be extended from 2016 to 2022 to fully utilize the datasets available. This step will double the number of sustainability analysed. Futhermore, it will allow us to explore the evolution of the relationship between ESG reports and ESG scores throughout the years.

7.2 Additional descriptors

Paragraphs relate a variety of information : they can inform stakeholders of risks, they can higlight emerging opportunities or they can describe how the company is reacting to the key issue. These different information about the companies' exposure and behaviour have distinct influences on the ESG scores. Indeed, according to the MSCI Methodology, large risks should to be negatively correlated with ESG scores and favorable management decisions should positively correlated.

To better understand what kind of reporting is done on the key issues, we want to add three topics categories : Risk, Opportunities, Management. This will help dissociate the three topics and avoid aggregating together conflicting information about the firm's disclosure.

7.3 Estimating pillar ESG Scores

The final step of the project would be to calculate the Environmental and Pillar Score and compare these estimations to ratings' agency scores. We would need to look into additional datasets, such as the controversy reports, to fully explore how ESG scores are estimated. Then, using sector specific weights, we can calculate estimated ESG pillar scores from external data. This would conclude our external assessment of ESG scores with data.

8 Conclusion

This research project tackles information extraction from sustainability reports using Transformer-based models. The method developed in this project scores sustainability reports according to the industry's main key issues. The results are encouraging as the calculated scores shows some correlations with MSCI ESG scores. This illustrates the score's capability to grasp the quality and the quantity of disclosed information concerning the topic. The next steps for the project are to run the algorithm on a larger set of sustainability reports and distinguish paragraphs according the following criteria: risk, opportunities or management. This will add an extra layer of granularity when exploring the data.

By exploring extraction methods, this project helps increase the transparency of ESG metrics and explain some of the discrepancies we find in sustainability scores. It opens the road for more advanced and reliable extraction methods to successfully utilize the information contained in sustainability reports, and convince investors to trust sustainability scores to orient their portfolios.

A Keywords

MSCI Key Issue	Keywords
Access To Finance	expand financial services
Access To Healthcare	improve health access developing countries
Anti-competitive practices	anti-competitive practices
Biodiversity and Land Use	impact operations biodiversity land
Business Ethics Fraud	fraud conflict of interest
Carbon Emissions	manage carbon related risks and opportunities
Controversial Sourcing	efforts sourcing traceability and certification
Corruption	corruption risks bribery
Financing Env Impact	capitalize on opportunities green finance
Finance Product Safety	regulations financial products
Health and Safety	workplace safety standards
Human Capital Development	employee training leadership productivity
Opps in Clean Tech	strategy and revenue clean technology
Opps in Green Building	building regulations and performance real estate
Opps in Nutrition and Health	improve nutritional health profile
Opps in Renewable Energy	renewable power development
Privacy and Data Security	privacy regulations information security
Product Carbon Footprint	reduce carbon footprint
Product Safety and Quality	product safety quality management
Raw Material Sourcing	materials traceability certification
Toxic Emission and Waste	toxic contamination management
Water Stress	water stress risks opportunities

References

[1] international Institute for Sustainable Development. *Sustainable Development.* URL: https://www.iisd.org/mission-and-goals/sustainable-development.

[2] Elizabeth Pollman. "The Making and Meaning of ESG". In: *U of Penn, Inst for Law & Econ Research Paper* No. 22-23 (2022). URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4219857.

[3] Governance & Accountability Institute. *2023 Sustainability Reporting in Focus.* URL: https://www.ga-institute.com/research/ga-research-directory/sustainability-reporting-trends/2023-sustainability-reporting-in-focus.html.

[4] KPML. *ESG data and ratings.* URL: https://kpmg.com/xx/en/home/insights/2023/03/esg-data-and-ratings.html#:~:text=ESG%20data%20and%20ratings%20encompass,emissions%20and%20risk%20exposure%20data..

[5] Florian Berg, Julian F Kölbel, and Roberto Rigobon. "Aggregate Confusion: The Divergence of ESG Ratings*". In: *Review of Finance* 26.6 (May 2022), pp. 1315–1344. ISSN: 1572-3097. DOI: 10.1093/rof/rfac033. eprint: https://academic.oup.com/rof/article-pdf/26/6/1315/47018560/rfac033.pdf. URL: https://doi.org/10.1093/rof/rfac033.

[6] Yan Lin et al. "Global Evolution of Environmental and Social Disclosure in Annual Reports". In: ().

[7] Julia Anna Bingler et al. “Cheap talk and cherry-picking: What ClimateBert has to say on corporate climate risk disclosures”. In: *Finance Research Letters* 47 (2022), p. 102776. ISSN: 1544-6123. DOI: https://doi.org/10.1016/j.frl.2022.102776. URL: https://www.sciencedirect.com/science/article/pii/S1544612322000897.

[8] Stefan Pasch and Daniel Ehnes. “NLP for Responsible Finance: Fine-Tuning Transformer-Based Models for ESG”. In: (2022), pp. 3532–3536. DOI: 10.1109/BigData55660.2022.10020755.

[9] Wenpeng Yin, Jamaal Hay, and Dan Roth. “Benchmarking Zero-shot Text Classification: Datasets, Evaluation and Entailment Approach”. In: *CoRR* abs/1909.00161 (2019). arXiv: 1909.00161. URL: http://arxiv.org/abs/1909.00161.

[10] Adina Williams, Nikita Nangia, and Samuel Bowman. “A Broad-Coverage Challenge Corpus for Sentence Understanding through Inference”. In: *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. Ed. by Marilyn Walker, Heng Ji, and Amanda Stent. New Orleans, Louisiana: Association for Computational Linguistics, June 2018, pp. 1112–1122. DOI: 10.18653/v1/N18-1101. URL: https://aclanthology.org/N18-1101.

[11] Lei Shu et al. “Zero-Shot Aspect-Based Sentiment Analysis”. In: *CoRR* abs/2202.01924 (2022). arXiv: 2202.01924. URL: https://arxiv.org/abs/2202.01924.