Detecting Greenwashing Signals Through a Comparison of ESG Reports and Public Media

Janna Lipenkova[1*], Guang Lu[2*], Susie Xi Rao[3*]

[1]Equintel GmbH
janna.lipenkova@anacode.de
[2]Lucerne University of Applied Sciences and Arts
guang.lu@hslu.ch
[3]ETH Zurich
srao@ethz.ch

[*]These authors contributed equally to this work.

Abstract

As the requirements on sustainability disclosure have been tightening over the past years, greenwashing has turned into a significant concern for companies and their stakeholders. Greenwashing occurs at the level of corporate communication when companies overemphasize their positive sustainability efforts while downplaying their negative impacts. This bias can mislead even highly experienced investors and stakeholders. The objective of our shared task was to explore natural language processing (NLP) methods to detect greenwashing and create more objective sustainability profiles of companies. To achieve this, we use a dataset that includes both self-reported and third-party data about the ESG practices of a range of companies from the Germany DAX index. We employed cutting-edge NLP techniques to identify signals of greenwashing within this data, enabling us to create more objective evaluations of the considered companies. In pursuit of this goal, we explored the use of large language models (LLMs), sentence and document embeddings, and sentiment analysis. Notably, the scientific reports submitted by participants in the collaborative task yielded intriguing initial findings that will be validated and deepened in future work.

1 Introduction

Greenwashing has been at the top of the social agenda of companies, investors, and society for the past two years (cf. Figure 1), and there is no sign that it will relinquish that position [1]. Greenwashing occurs when companies publish overly positive data and claims about their sustainability efforts while downplaying the negative impacts of their operations [1, 2]. A major reason for this is the subjective nature of most ESG (Environmental, Social, and Governance) information. As public companies face growing pressure to report on their sustainability and ESG efforts, they often succumb to greenwashing, i.e. overly positive, false, or incomplete statements about their impact on society and the planet. Greenwashing can be intentional, but also unintentional due to a lack of ESG expertise at the company. In turn, investors and ESG rating agencies focus on data that is disclosed by companies, such as ESG reports and marketing communications, and often get a distorted and overly positive picture of the company's sustainability and ESG performance. Facing this information asymmetry, even investors who want to invest sustainably can be misled in their investment decisions [3].

In this shared task, we treat greenwashing as an information problem [4] which creeps in through various internal textual communications of a company, such as ESG reports, press releases, and marketing statements. We believe that external media, such as the business press, as well as publications from non-governmental organizations (NGOs) and think tanks, can help address this problem. Our as-

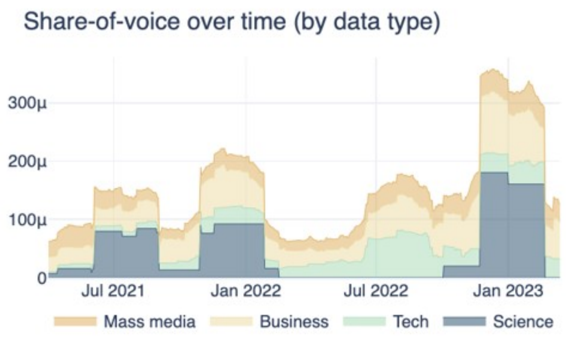

Figure 1: Share-of-voice: Proportion of mentions of the term "greenwashing" compared to all other terms in the data related to the companies under study.

sumption is that most third-party content providers have no interest in promoting the sustainability efforts of a particular company. Therefore, we can get a more objective and critical picture of a company if we consider data from a variety of public media sources.

To detect greenwashing signals, we provide participants in the shared task with a dataset that includes company ESG reports as well as ESG-related public media articles targeting a wide range of stakeholders such as investors, NGOs, regulators, and society. We also provide a dataset with descriptions of the 17 Sustainable Development Goals (SDGs) [5], which serve as a blueprint for global social, economic, and environmental challenges. The task is to develop approaches to address gaps and inconsistencies between company-reported data and external data that may indicate greenwashing, which can be done at the ESG sentiment and/or SDGs level. This is an application-oriented text analysis task for which various natural language processing (NLP) algorithms can be applied [6, 7, 8, 9]. We suggest focusing on the following three aspects:

1. Understand the nature of greenwashing and quantify its extent in internal text data.

2. Develop NLP approaches to detect greenwashing using public documents reflecting various stakeholders.

3. Visualize possible indicators of greenwashing as well as the reliability of these indicators in a transparent way.

2 Data, Subtasks and Participants

2.1 Data

The DAX ESG Media Dataset[1] contains approximately 11,000 external and internal English documents on 38 DAX companies from 2021 to 2023, as well as an additional file with descriptions of the SDGs[2]. The list of fields in the data is as follows:

- *Symbol*: Stock symbol of the company
- *Company*: Company name
- *Date*: Publication date of the document
- *Title*: Document title
- *Content*: Document content
- *Datatype*: Document type
- *Internal*: Document provided by the company (1) or a third-party (0)
- *Domain (optional)*: Web domain where the document was published
- *URL (optional)*: URL where the document can be accessed
- ESG_topics (optional): ESG topics extracted from the data using our internal NLP pipeline

A snapshot of the data frame is shown in Figure 2. The content column is the main focus.

2.2 Subtasks

To overcome the challenge of detecting greenwashing signals from text data, we recommend the following five stages, shown in Figure 3.

Stage 1: Exploratory data analysis (EDA), preprocessing and cleaning. In this step, participants will get an initial overview of the dataset and prepare it for subsequent NLP analysis. This may include the following statistical analyses:

- Average document length by data type
- Word segmentation and word frequencies
- Number of documents by company
- Using term frequency - inverse document frequency (TF-IDF) to find the most characteristic words by company
- Timeseries of ESG topic distributions to analyze patterns over time

[1] https://www.kaggle.com/datasets/equintel/dax-esg-media-dataset
[2] https://sdgs.un.org/goals

	company	content	datatype	date	domain	esg_topics	internal	symbol	title	url
2	Beiersdorf AG	Sustainability Highlight Report CARE BEYOND SK...	sustainability_report	2021-03-31	NaN	['CleanWater', 'GHGEmission', 'ProductLiabilit...	1	BEI	BeiersdorfAG Sustainability Report 2021	NaN
3	Deutsche Telekom AG	Corporate Responsibility Report 2021 2 Content...	sustainability_report	2021-03-31	NaN	['DataSecurity', 'Iso50001', 'GlobalWarming', ...	1	DTE	DeutscheTelekomAG Sustainability Report 2021	NaN
5	Vonovia SE	VONOVIA SE SUSTAINABILITY REPORT 2021 =For a S...	sustainability_report	2021-03-31	NaN	['Whistleblowing', 'DataSecurity', 'Vaccine', ...	1	VNA	VonoviaSE Sustainability Report 2021	NaN
6	Merck KGaA	Sustainability Report 2021 TABLE OF CONTENTS S...	sustainability_report	2021-03-31	NaN	['DataSecurity', 'DataMisuse', 'DrugResistance...	1	MRK	MerckKGaA Sustainability Report 2021	NaN
9	MTU	Our ideas and concepts FOR A SUSTAINABLE FUTUR...	sustainability_report	2020-03-31	NaN	['WorkLifeBalance', 'Corruption', 'AirQuality'...	1	MTX	MTUAeroEngines Sustainability Report 2020	NaN

Figure 2: Snapshot of the DAX ESG Media Dataset.

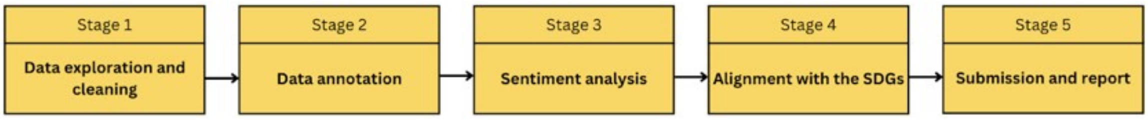

Figure 3: NLP analysis pipeline to detect greenwashing signals from the DAX ESG Media Dataset.

This investigation should also shed light on possible data preprocessing and cleaning steps, such as:

- Word segmentation
- Cleaning up the table of contents
- Cleaning numbers and special characters

The output of Stage 1 could be:

- EDA notebook with visualizations
- Notebook with data cleaning steps

Stage 2: Data annotation. Here, participants use large language models (LLMs) of their choice to annotate the sentiment data (0 = negative, 0.5 = neutral, 1 = positive), which are then used as training data for the sentiment analysis classifier in Stage 3. Participants are free in their choice of an LLM. We recommend considering multiple options, for example:

- OpenAI models[3]
- Open-source LLMs trained with multi-task learning, such as T5[4] and T0[5]
- Other large-scale LLMs, such as BERT[6] and FLAN[7]

In terms of granularity, annotation can be done at document or sentence level. We recommend the sentence level, as this leads to more detailed insights. The steps can be as follows:

- Manual annotation of about 200 documents and/or 500 sentences as a "gold standard". Random selection of sample documents and sentences should take into account the distribution of internal and external data within the entire dataset
- Set up 2-3 LLMs to test annotation against manual annotation
- Experiment with different prompting strategies, e.g., zero-shot and few-shot, for LLMs to annotate the data and compare with the "gold standard"

The outputs of stage 2 are:

- Manually annotated "gold standard" dataset of min. 200 documents or 500 sentences
- Automatically annotated documents for the entire dataset based on the selected LLMs
- Description of the prompting strategies for LLMs

Stage 3: Sentiment analysis and comparison between internal and external data. In this stage, participants create a train/validation/test split (recommended ratio: 70%/15%/15%) of the data annotated in Stage 2 and train a sentiment analysis classifier. The classifier should output sentiment scores on a continuous scale between 0 and 1. Participants then compare the average sentiment of internal and

[3]https://platform.openai.com/docs/introduction/key-concepts
[4]https://huggingface.co/docs/transformers/model_doc/t5
[5]https://huggingface.co/bigscience/T0
[6]https://huggingface.co/bert-base-uncased
[7]https://blog.research.google/2021/10/introducing-flan-more-generalizable.html

external data about a company. They sort the companies based on the difference between internal and external sentiment and conduct a manual follow-up analysis to see if the companies with the largest gap were explicitly involved in greenwashing during the time period under consideration.

The output of Stage 3 could be:

- Code for training the sentiment analysis classifier
- Precision, recall, and accuracy of the classifier in the training and test datasets

Stage 4: Alignment with the SDGs. In this stage, participants try to determine the relevance of specific SDGs for the different companies. The SDGs are described in the supplementary SDGs dataset provided in the shared task. We propose the following two approaches:

- LLM-based approach: Formulate a prompt to directly query the relevance of a particular SDG description to documents about a company. We recommend doing this in the context of an additional query framework, such as the retrievers in LangChain[8].
- Sentence embedding approach: Use an embedding library such as Laser[9] or Sentence-BERT[10] to embed both the SDG descriptions and the documents about a company and compare them for similarity.

Following the analysis, participants can use a visualization of their choice, such as heat maps and bar charts, to demonstrate the relevance of the SDGs to companies.

Stage 5: Report of the results. In this stage, participants prepare a report describing the methodology and results of the different analysis stages.

2.3 Participants

A total of 25 Master students solved the shared task. They all come from the master's program Applied Information and Data Science at Lucerne University of Applied Sciences and Arts in Switzerland. The Stages 1 and 2 were solved by each student individually. The Stages 3 to 5 were solved by the students in a group of two to three persons. Ten groups were formed for the Stages 3 to 5.

[8] https://python.langchain.com/docs/modules/data_connection/retrievers.html

[9] https://github.com/facebookresearch/LASER

[10] https://www.sbert.net/

Figure 4: Word cloud visualization of the topics mentioned in the documents. Work by Tenzin Rungwatsang. Resolution adapted and figure regenerated by Susie Xi Rao.

3 Results and Discussion

This section summarizes the key findings from the various participant groups. The NLP analysis code and scientific reports of all participants are available online[11] with their permission.

3.1 Data Preprocessing and Exploratory Data Analysis (EDA)

Participants first performed systematic data cleaning and preprocessing. This includes:

- Converting strings to lowercase, decoding Unicode
- Removing URL and e-mail address, extra spaces, contact information, tables of contents, named entity, special characters and stop words
- Abbreviation expansion, part-of-speech (POS) tagging, sentiment analysis

Next, basic information about the dataset was reviewed, from average document length to number of documents per company to ESG topics. In addition, words were segmented to calculate word frequency, TF-IDF analysis was performed to find the characteristic word per company, and topics

[11] https://drive.google.com/drive/folders/126J34mGwCgEZ8MKVYdqTSr_MLuHTqzxu

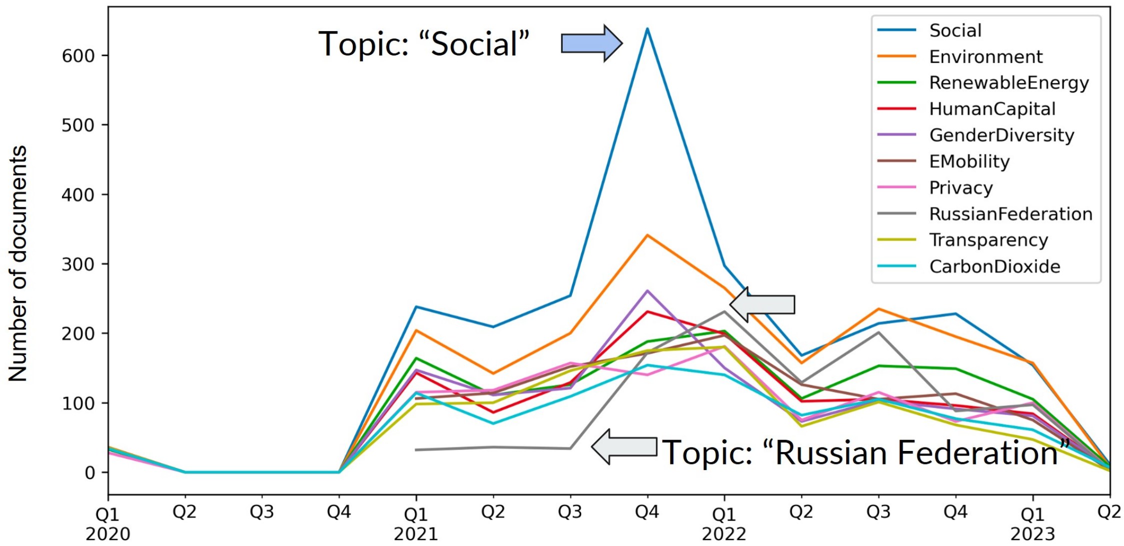

Figure 5: Development and trends in ESG topics. Work by Tenzin Rungwatsang.

were extracted from documents using topic modeling methods and visualized with word clouds (e.g., Figure 4). Finally, the pattern of ESG topic distribution was analyzed in a time series. For example, Figure 5 shows the number of documents per topic as a function of time in quarters. In the fourth quarter of 2021, the focus was on the topics of social affairs, environment and gender diversity. Since the beginning of 2022, the topic of the Russian Federation has been included in the analysis, likely due to the ongoing conflict. In addition to topic distribution, it is worth noting that some participants at this stage have already explored other interesting patterns in the data, such as changes in text polarity over time and visualization of word embedding.

3.2 Annotation of Text Sentiment with Large Language Models (LLMs)

Manual annotation of text sentiment was mostly done at the document and sentence level, while sentence-level annotation worked better in our context, as shown by participants' results. For example, 200-1,000 annotations were made at the sentence level, with sentiment indicated as negative (0), neutral (0.5), and positive (1). The annotations were made independently of the ESG topics, as we believe this already provides a good baseline for testing the annotation performance of the LLMs. The imbalance of the data caused by the different number of internal and external documents in our dataset was taken into account, so that the text samples drawn reflect this imbalanced distribution. Different LLMs were selected to re-annotate the sentences previously annotated by humans (“gold standard”). Participants chose different annotation strategies with different LLMs, such as:

- Zero-shot: BERT[12], T5[13], DistilBERT[14], RoBERTa[15]
- Few-shot: GPT-3.5[16] (e.g., text-davinci-002, text-davinci-003, GPT-3.5-turbo), FLAN-T5[17]

We then compared human and LLMs annotations and determined the best LLMs performance. It can be concluded that manual annotation should follow a well-defined strategy to minimize the impact on subsequent LLMs annotation. In general, the GPT-3.5 model performs better on sentiment annotations compared to the other models, but incurs more cost and computation time.

Figure 6 shows a comparison of the confusion matrix of human and LLMs annotations of text sentiment based on zero-shot and few-shot of the GPT-3.5-turbo model. It can be clearly seen that learning with few-shot performs better than learning with zero-shot in predicting the sentiments. However,

[12] https://huggingface.co/docs/transformers/model_doc/bert
[13] https://huggingface.co/docs/transformers/model_doc/t5
[14] https://huggingface.co/docs/transformers/model_doc/distilbert
[15] https://huggingface.co/docs/transformers/model_doc/roberta
[16] https://platform.openai.com/docs/models/gpt-4
[17] https://huggingface.co/docs/transformers/model_doc/flan-t5

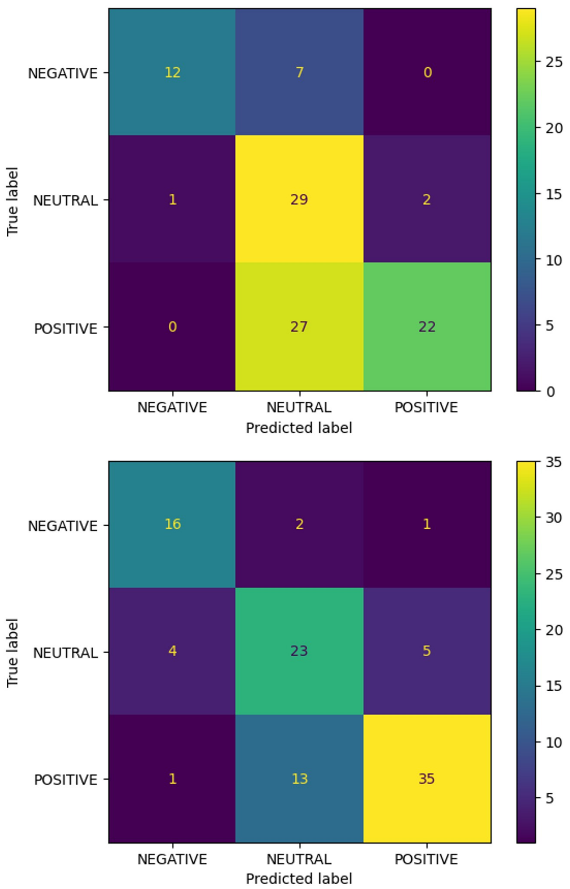

Figure 6: Comparison of human and LLMs annotations of text sentiment based on zero-shot and few-shot of the GPT-3.5-turbo model. Top: Zero-shot; Bottom: Few-shot. Work by Zal Paulina Aleksandra.

the model can still be improved as it tends to predict positive texts into neutral ones.

3.3 Sentiment Analysis of Internal and External Documents

This stage aims to calculate and identify a large discrepancy in sentiment scores (1: positive, 0: negative) between internal and external documents of the companies under study. The greenwashing signal is reported when the internal scores are significantly higher than the external ones. This discrepancy could be indicative of greenwashing signals. Most of the participants have analyzed that greenwashing occurs when internal sentiment scores are substantially higher.

The analysis has been conducted at the document level, with the possibility of aggregating sentiment scores from sentence-level analyses using either simple averaging or count-based metrics. To address potential class imbalance issues, techniques such as upsampling, downsampling, or stratified *k*-fold cross-validation have been explored.

To assess the accuracy of the sentiment analysis, participants made use of the ground truth data, which could involve human annotation generated in Stage 2, machine-generated sentiment analysis, or a combination of both. Furthermore, different scales of sentiment have been considered, including categorical (0, 1), numerical (0, 0.5, 1), and continuous (ranging from 0 to 1), to capture the full spectrum of sentiment variations.

The approach to sentiment analysis has encompassed various models, including out-of-the-box LLMs, fine-tuned LLMs, and traditional Machine Learning (ML) models. Model selection has involved evaluating options such as RoBERTa, BART, FinBERT, GPT-3.5, FLAN-T5, among others, based on individual evaluations for the task (such as budget, accuracy). For evaluating the sentiment analysis, a range of metrics, including accuracy, precision, recall, F1-score, Mean Squared Error (MSE), Mean Absolute Error (MAE), Mean Relative Error (MRE), and R-squared (R2), have been employed.

In addition to sentiment analysis, an exploration of the phenomenon of greenwashing has been undertaken by conducting analyses across different sectors, topics, companies, and over time to identify patterns and trends. As part of the research, detailed case studies and manual analyses have been performed. These analyses have identified common entities such as Beiersdorf AG, Deutsche Bank AG, and the automobile industry as the main source of greenwashing in our dataset.

By integrating both automated sentiment analysis and manual examination of specific cases, a comprehensive understanding of sentiment discrepancies and greenwashing practices in various contexts has been sought. As shown in Figure 7 we observe a negative difference in sentiment (external-internal), meaning that external documents have a lower average sentiment than internal documents, could be indicative of greenwashing.

We also can find official evidence of these incidences. In 2002, Beiersdorf AG faced allegations of making false claims about achieving CO2 emission neutrality in the production of their products. This incident aligns with the disparity we observed in our data concerning internal and external documents. In 2022, Deutsche Bank AG was compelled to part ways with one of its executives due to a scandal linked to supposedly sustainable funds that

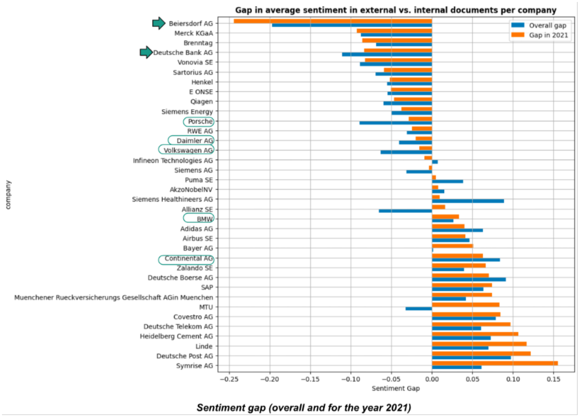

Figure 7: Sentiment gap (external minus internal) in average sentiment. Work by Patrick Fox, Yasmine Mohamed.

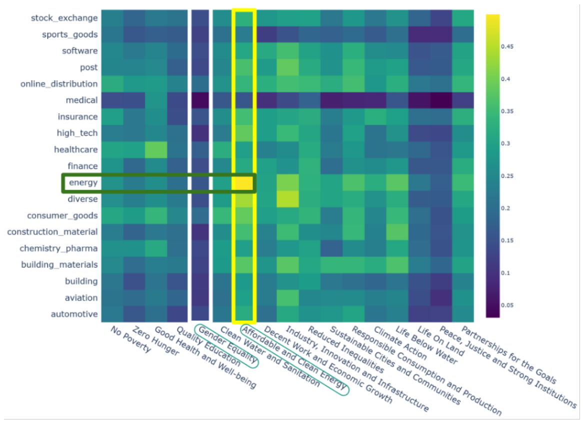

Figure 8: Heatmap of similarities of SDGs per industry. Work by Palten Stefanie, Freimann Elisabeth, and Zal Paulina Aleksandra.

failed to meet the promised sustainability criteria. This situation underscores the idea that when a company publishes documents filled with commitments, the discrepancies can ultimately become evident in the data.

3.4 Alignment with Sustainable Development Goals (SDGs)

In Stage 4 of the project, the primary objective was to assess the alignment between the company reports and the descriptions of the SDGs, while also identifying the discrepancies between the internal and external reports. The methodology involved encoding company documents and SDGs descriptions using sentence embeddings and calculating their similarity through cosine similarity.

As suggested by the organizers, an interesting direction could be to employ LLMs to create prompts for direct inquiries regarding the relevance of specific SDGs descriptions to company documents, although this avenue remains relatively unexplored due to the limited time before result submission.

The assessment of SDGs relevance was extended to individual industries, revealing variations in the importance of these goals across sectors. Furthermore, this stage delved into the relationship between sentiment analysis, the distinction between internal and external documents, and SDGs topic analysis.

Some notable findings included the varying degrees of relevance of specific SDGs, with “affordable and clean energy,” “industry, innovation and infrastructure,” and “responsible consumption and production” standing out as highly pertinent for companies, while “gender equality” was deemed less relevant. Additionally, these findings underscored the sector-specific nature of SDGs alignment, with significant differences observed, particularly within the automotive industry. This stage aimed to provide a comprehensive understanding of how companies engage with SDGs, taking into account specific goals, industry contexts, and the role of sentiment analysis and document types in shaping this alignment.

As seen in Figure 8, company documents exhibited the strongest alignment with “affordable and clean energy”, while the weakest alignment was observed with “gender equality”. Notably, the energy industry displayed the most robust alignment with “affordable and clean energy”, surpassing other sectors in this regard. Figures 9 and 10 show for most of the companies, internal documents lie closer to SDGs. In the radar plot (Figure 11) we observe the alignment of German car company profiles to SDGs. The car companies have shown similar patterns and they are the nearest neighbors in the sentence embedding space. The scatter plot in Figure 12 depicts the alignment between the sentiment of the internal document and the SDGs Sum suggesting a higher level of consistency between the sentiment of the internal document and the overall alignment with the SDGs.

4 Discussion, Conclusion and Outlook

During the course of this project, several limitations and challenges were encountered, primarily stemming from computational constraints and manual annotation difficulties. These limitations, in turn, influenced the project’s scope and outcomes.

- Manual Annotation Challenges: The manual

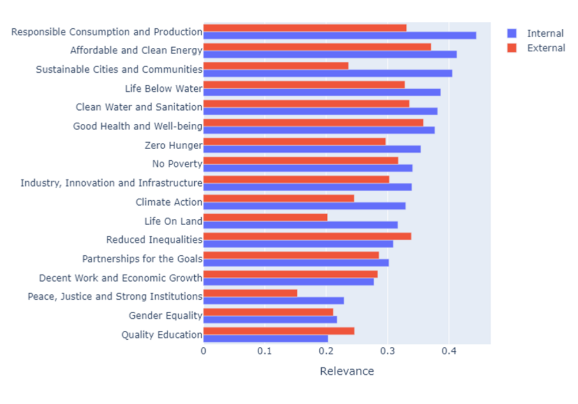

Figure 9: SDG relevance for Beiersdorf AG. Work by Palten Stefanie, Freimann Elisabeth, and Zal Paulina Aleksandra.

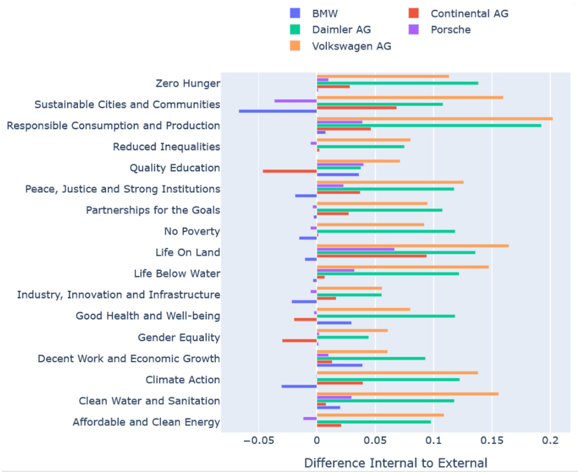

Figure 10: Difference in SDG relevance of internal and external documents. Work by Palten Stefanie, Freimann Elisabeth, and Zal Paulina Aleksandra.

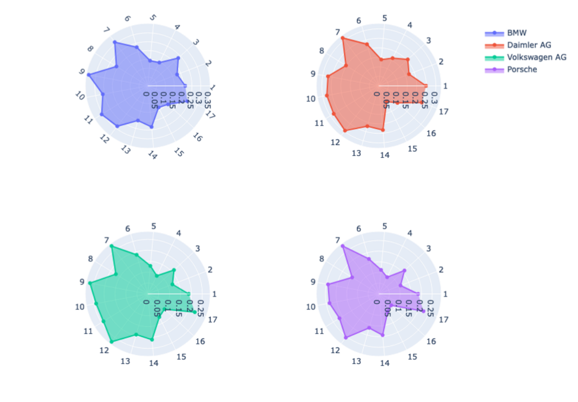

Figure 11: Radar plot indicating alignment of companies of the German car industry and the SDGs (1-17 SDG goals). Work by Thomas Radinger, Marco Rieder.

annotation of documents, especially in identifying instances of greenwashing, posed a challenge due to limited subject knowledge.

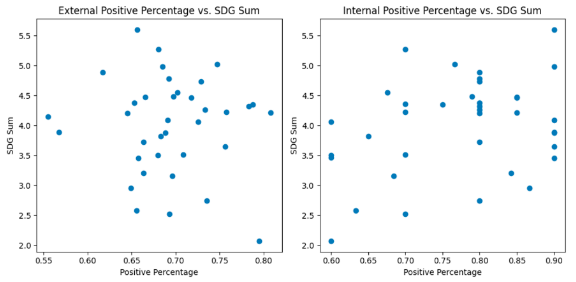

Figure 12: Sentiment and its correlation with internal and external documents. Work by Julian Kalis, Kimberly Kent.

This required meticulous consideration to ensure accurate document labeling. However, there was no guarantee of 100 percent accuracy, potentially affecting subsequent project stages.

- Computational Challenges: The project faced computational limitations that prevented the use of reliable LLMs. Instead, a self-tuned/trained model with weaker accuracy was employed, leading to less reliable predictions across the entire dataset.

- Heterogeneous Data Sources: While combining company-reported and third-party data provides a more objective image of a company, it also makes the application of NLP algorithms, especially sentiment analysis, more challenging due to the very different styles of the texts. For more reliable results, separate sentiment classifiers should be trained for each data type.

- Scope for Further Development: Future developments in this area may benefit from comparing documents produced under the same reporting standards, such as the Global Reporting Initiative (GRI), International Financial Reporting Standards (IFRS), Sustainability Disclosure Standards, or Sustainability Accounting Standards Board (SASB) Standards. Similar structures in these documents would facilitate meaningful company comparisons.

Despite these challenges, the trained model demonstrated high accuracy in sentiment analysis. However, a thought-provoking observation arose regarding the complexity of the methods used. After significant computational and temporal investments in earlier stages, the question emerged: could simpler approaches like logistic regression yield comparable results? This raises questions about the necessity of using larger models and complex methods.

In the project's final stages, no greenwashing indications were detected in the provided data based on the methods applied. The alignment between internal and external communications and similarities with SDGs supported these findings. Nonetheless, the possibility exists that unexplored methods may yield different results, encouraging further exploration and evaluation of alternative models and approaches to gain deeper insights.

In conclusion, while the project faced computational and manual annotation challenges, it ultimately provided valuable insights into sentiment analysis, greenwashing detection, and alignment with SDGs. The findings suggest the absence of greenwashing based on the applied methods but also underscore the need for ongoing exploration and refinement in this field.

Acknowledgements

The authors thank the organizers from SwissText2023 for hosting our workshop. Special acknowledgment goes to Guerne Jonathan and Emmanuel de Salis from the University of Applied Sciences and Arts of Western Switzerland for their great support throughout the workshop organization. The authors thank Peter Egger (Chair of Applied Economics, Department of Management, Technology, and Economics at ETH Zurich) for his support. The authors would also like to thank the Applied Information and Data Science master's program at the Lucerne University of Applied Sciences and Arts for supporting this workshop.

References

[1] Santos, C., Coelho, A., & Marques, A. (2023). A systematic literature review on greenwashing and its relationship to stakeholders: state of art and future research agenda. Management Review Quarterly, 1-25.

[2] de Freitas Netto, S. V., Sobral, M. F. F., Ribeiro, A. R. B., & Soares, G. R. D. L. (2020). Concepts and forms of greenwashing: A systematic review. Environmental Sciences Europe, 32(1), 1-12.

[3] Kleffel, P., & Muck, M. (2023). Aggregate confusion or inner conflict? An experimental analysis of investors' reaction to greenwashing. Finance Research Letters, 53, 103421.

[4] Moodaley, W., & Telukdarie, A. (2023). Greenwashing, Sustainability Reporting, and Artificial Intelligence: A Systematic Literature Review. Sustainability, 15(2), 1481.

[5] Assembly, G. (2015). Sustainable development goals. SDGs Transform Our World, 2030, 6-28.

[6] Boelders, F. B. (2020). The relationship between firm characteristics and greenwashing in the European energy sector; An NLP approach (Master's thesis).

[7] Woloszyn, V., Kobti, J., & Schmitt, V. (2021, December). Towards Automatic Green Claim Detection. In Proceedings of the 13th Annual Meeting of the Forum for Information Retrieval Evaluation (pp. 28-34).

[8] Kim, J. S., Sim, J. B., Kim, Y. J., Park, M. K., Oh, S. J., & Doo, I. C. (2022, December). Establishment of NLP-Based Greenwashing Pattern Detection Service. In International Conference on Computer Science and its Applications and the International Conference on Ubiquitous Information Technologies and Applications (pp. 253-259). Singapore: Springer Nature Singapore.

[9] tammbach, D., Webersinke, N., Bingler, J., Kraus, M., & Leippold, M. (2023, July). Environmental Claim Detection. In Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers) (pp. 1051-1066).

A List of participants to the workshop

We thank our workshop participants for valuable feedback, contributions, and suggestions.

Janna Lipenkova, Equintel GmbH (Organizer)
Guang Lu, Lucerne University of Applied Sciences and Arts (Organizer)
Susie Xi Rao, ETH Zurich (Organizer)
Yihan Deng, Swiss Post
Celien Donze, University of Applied Sciences and Arts of Western Switzerland
Hatem Ghorbel, University of Applied Sciences and Arts of Western Switzerland
Flurin Gishamer, BSI Business Systems Integration AG
Anna Rogers, IT-University of Copenhagen
Peter Egger, ETH Zurich (Principal Investigator)
Participants from the master's program Applied Information and Data Science at Lucerne University of Applied Sciences and Arts in Switzerland: Claire Bussat, Arian Contessotto, Camille Cosandier, Patrick Fox, Elisabeth Freimann, Tim Giger, Michaela Havlickova, Martina Heidemann, Julian Alexander Kalis, Kimberly Kent, Luca Laboranti, Christopher T. Loo, Adrian Meier, Yasmine Mohamed, Nithusya Mohanadasan, Milica Pajkic, Stefanie Palten, Andreia Sofia Pereira Almeida, Kelly Queiroga, Thomas Radinger, Levin Reichmuth, Marco Rieder, Tenzin Rungwatsang, Ushasri Thouti, Paulina Aleksandra Zal