DOI(Journal): 10.31703/gssr
DOI(Volume): 10.31703/gssr.2025(X)
DOI(Issue): 10.31703/gssr.2025(X.III)

p-ISSN: 2520-0348

e-ISSN: 2616-793X

www.gssrjournal.com
Global
Social Sciences Review
exploring humanity

GSSR

GLOBAL SOCIAL SCIENCES REVIEW

HEC-RECOGNIZED CATEGORY-Y

Volum X, ISSUE III SUMMER (SEPTEMBER-2025)

HumanityPublications
sharing research
www.humanitypublications.com
Karachi | Islamabad | Lahore

Double-blind Peer-review Journal
www.gssrjournal.com

Humanity Publications(HumaPub)
www.humapub.com
Doi: https://dx.doi.org/10.31703

Article Title

Greenwashing in Corporate Climate Disclosures: A Machine Learning-Based Detection Approach

Abstract

Corporate climate disclosures have come to the fore of measuring environmental responsibility, but worries about greenwashing of exaggeration or parts of the environmental performance of exaggerating or overselling environmental performance remain. This paper fulfills this crucial gap in establishing the validity of such revelations by offering the machine learning method of identifying possible greenwashing. It is probable that the mixed-methods design has been used, where the textual analysis of the composed corporate sustainability reports and supervised learning algorithms trained on labeled examples of misleading statements are supplemented. Through the implementation of natural language processing and classification algorithms, the model will recognise patterns that are suggestive of a lack or even exaggeration of commitment with regard to climate pledges. The findings can be used to illustrate industry-related patterns and important language indications linked to greenwashing.

Keywords: Greenwashing, Climate Disclosures, Machine Learning, Corporate Sustainability, Text Analysis

Authors:

Adeel Ahmed:(Corresponding Author)
Masters in Data science, Department of Computer science, National Research University Higher School of Economics, Russia.
(Email: adakhmed@edu.hse.ru)

Sumaira Raza: Teacher (M.A. Political Science), Department of Elementary Education, Master Trainer Pedagogy, KP, Pakistan.

Romaila: MPhil Scholar, Department of Political Science, Abdul Wali Khan University, Mardan, KP, Pakistan.

Pages: 110-124
DOI:10.31703/gssr.2025(X-III).10
DOI link: https://dx.doi.org/10.31703/gssr.2025(X-III).10
Article link: https://gssrjournal.com/article/greenwashing-in-corporate-climate-disclosures-a-machine-learningbased-detection-approach
Full-text Link: https://gssrjournal.com/fulltext/greenwashing-in-corporate-climate-disclosures-a-machine-learningbased-detection-approach
Pdf link: https://www.gssrjournal.com/jadmin/Auther/31rvIolA2.pdf

Global Social Sciences Review

p-ISSN: 2520-0348 **e-ISSN:** 2616-793x
DOI(journal):10.31703/gssr
Volume: X (2025)
DOI (volume):10.31703/gssr.2025(X)
Issue: III Summer (September-2025)
DOI(Issue):10.31703/gssr.2025(X-III)

Home Page
www.gssrjournal.com

Volume: X (2025)
https://www.gssrjournal.com/Current-issue

Issue: III-Summer (September -2025)
https://www.gssrjournal.com/Current-issues/10/3/2025

Scope
https://www.gssrjournal.com/about-us/scope

Submission
https://humaglobe.com/index.php/gssr/submissions

Visit Us

Humanity Publications (HumaPub)
www.humapub.com
Doi: https://dx.doi.org/10.31703

Citing this Article

10	Greenwashing in Corporate Climate Disclosures: A Machine Learning-Based Detection Approach		
Authors	Adeel Ahmed Sumaira Raza Romaila	DOI	10.31703/gssr.2025(X-II).10
		Pages	110-124
		Year	2025
		Volume	X
		Issue	III

Referencing & Citing Styles	
APA	Ahmad, A., Raza, S., & Romaila. (2025). Greenwashing in Corporate Climate Disclosures: A Machine Learning-Based Detection Approach. *Global Social Sciences Review*, *X*(III), 110-124. https://doi.org/10.31703/gssr.2025(X-III).10
CHICAGO	Ahmad, Adeel, Sumaira Raza, and Romaila. 2025. "Greenwashing in Corporate Climate Disclosures: A Machine Learning-Based Detection Approach." *Global Social Sciences Review* X (III):110-124. doi: 10.31703/gssr.2025(X-III).10.
HARVARD	Ahmad, A., RAZA, S. & ROMAILA 2025. Greenwashing in Corporate Climate Disclosures: A Machine Learning-Based Detection Approach. *Global Social Sciences Review,* X, 110-124.
MHRA	Ahmad, Adeel, Sumaira Raza, and Romaila. 2025. 'Greenwashing in Corporate Climate Disclosures: A Machine Learning-Based Detection Approach', *Global Social Sciences Review*, X: 110-24.
MLA	Ahmad, Adeel, Sumaira Raza, and Romaila. "Greenwashing in Corporate Climate Disclosures: A Machine Learning-Based Detection Approach." *Global Social Sciences Review* X.III (2025): 110-24. Print.
OXFORD	Ahmad, Adeel, Raza, Sumaira, and Romaila (2025), 'Greenwashing in Corporate Climate Disclosures: A Machine Learning-Based Detection Approach', *Global Social Sciences Review,* X (III), 110-24.
TURABIAN	Ahmad, Adeel, Sumaira Raza, and Romaila. "Greenwashing in Corporate Climate Disclosures: A Machine Learning-Based Detection Approach." *Global Social Sciences Review* X, no. III (2025): 110-24. https://dx.doi.org/10.31703/gssr.2025(X-III).10.

e-ISSN:2616-793x Volume: X (2025) Issue: III-Summer (September -2025) p-ISSN:2520-0348

Global Social Sciences Review
www.gssrjournal.com
DOI:http://dx.doi.org/10.31703/gssr

Pages: 110-124 URL: https://doi.org/10.31703/gssr.2025(X-III).10 Doi: 10.31703/gssr.2025(X-III).10

Cite Us

Title

Greenwashing in Corporate Climate Disclosures: A Machine Learning-Based Detection Approach

Authors:

Adeel Ahmed:(Corresponding Author)
Masters in Data science, Department of Computer science, National Research University Higher School of Economics, Russia.
(Email: adakhmed@edu.hse.ru)

Sumaira Raza: Teacher (M.A. Political Science), Department of Elementary Education, Master Trainer Pedagogy, KP, Pakistan.

Romaila: MPhil Scholar, Department of Political Science, Abdul Wali Khan University, Mardan, KP, Pakistan.

Contents

Abstract

Corporate climate disclosures have come to the fore of measuring environmental responsibility, but worries about greenwashing of exaggeration or parts of the environmental performance of exaggerating or overselling environmental performance remain. This paper fulfills this crucial gap in establishing the validity of such revelations by offering the machine learning method of identifying possible greenwashing. It is probable that the mixed-methods design has been used, where the textual analysis of the composed corporate sustainability reports and supervised learning algorithms trained on labeled examples of misleading statements are supplemented. Through the implementation of natural language processing and classification algorithms, the model will recognise patterns that are suggestive of a lack or even exaggeration of commitment with regard to climate pledges. The findings can be used to illustrate industry-related patterns and important language indications linked to greenwashing.

Keywords:

Greenwashing, Climate Disclosures, Machine Learning, Corporate Sustainability, Text Analysis

Introduction

The situation of climate change is becoming urgent, and as a result, corporations are increasingly under pressure to prove their commitment to environmental stewardship by taking actions and reporting on climate in an effective, transparent, and verifiable manner that answers the stakeholder questions of investors, consumers, regulators, and civil society. They have then turned into a necessary tool that allows conveying the carbon footprint and sustainability objectives or achievements of firms and their status of decarbonization (Eccles et al., 2023; IPCC, 2023). But these accounts are increasingly being questioned with the emergence of greenwashing, a form of strategic misrepresentation or exaggeration of the environmental performance aimed at deceiving stakeholders (Delmas & Burbano, 2011; Lyon & Montgomery, 2015). This inconsistency between environmental reporting and practice negates the validity of the sustainability reporting, as well as discourages global climate mitigation actions.

The issue of greenwashing is not theorized but proven to be a reality. Research indicates that a slew of firms still overreport or exaggerate their status in terms of climate action or remain evasive and unverifiable with respect to their concern about the environment (Melloni et al., 2023; Kotsantonis & Serafeim, 2024). These practices have necessitated regulators and standard-setters like the European Financial Reporting Advisory Group (EFRAG) and the U.S. Securities and Exchange Commission (SEC) to prescribe more rigorous disclosure policies and to impose the policy of ESG (Environmental, Social, Governance) auditing (OECD, 2024). Yet, the installed traditional monitoring and assurance software is still limited due to subjectivity, its high cost, and response delay (Cho et al., 2021). This limitation explains the need to have a more scalable and data-driven means of identification and prevention of greenwashing in an efficient manner.

In such a setting, a new innovation that has manifested itself is the use of artificial intelligence (AI) and machine learning (ML) in sustainability auditing. By using natural language processing (NLP) and machine learning, it is already possible to observe significant success of the former in the process of parsing and classifying extremely large amounts of unstructured text, including annual reports and sustainability statements (Zhang et al., 2022; Zeng et al., 2024). The adherence to the concepts of truth and deception in environmental statements can be turned into an automated task by training machine learning on labeled data that identifies either premises of false statements or inaccuracies present within the corporate climate reports. The described change in the methodology corresponds with an overall shift in the financial and sustainability studies that involves an aspiration towards predictive analytics and transparency in computing (Hartmann et al., 2025).

Notwithstanding these technological achievements, the literature in the academic field has not gone far enough to provide the ability to operationalize greenwashing identification within a systematic, reproducible, and industry-specific methodological framework. The majority of the established research bases its findings on qualitative models or subjective assessments furnished by the ESG agencies and does not agree on the rating parameters or disclosure (Berg et al., 2022). Furthermore, only a small number of empirical studies established or evaluated machine learning models designed to be used in relation to environmental misreporting, presenting a methodological and evidence gap in the literature (Wu et al., 2023). It is necessary to approach the problem of this gap with a multidisciplinary approach based on the perspectives of environmental economics, corporate accountability, computational linguistics, and data science.

The present study can be used to advance this new area of study by proposing a machine learning-based greenwashing detection methodology in corporate climate reporting. The study is based on a concept of supervised learning and natural language processing, where the researchers will seek to identify industry-specific trends and wording correlating with environmental false claims. The study will combine qualitative and quantitative aspects and merge the labeling of greenwashing indicators with quantitative classification in a mixed-methods design, which will improve the strength and scalability of greenwashing audits. This not only enhances the mechanism of the methodological toolbox of sustainability analytics but also helps policymakers, investors, and watchdogs to encourage corporations to tell their stories regarding the climate.

This study would be relevant in enhancing the popularity and authenticity of corporate climate disclosures when the integrity of environmental communication is paramount to the success of climate efforts worldwide. Computerizing the scope of greenwashing is in concert with other international initiatives that have lobbied to make comparable, evident, and decision-helpful sustainability reporting, including the Task Force on Climate-related Financial Disclosure (TCFD), and the International Sustainability Standards Board (ISSB) (IFRS Foundation, 2023). Moreover, the paper is a part of the general academic discussion of corporate environmentalism, and it supports the idea that technological innovation should follow ethical, transparent, and responsible sustainability regulations.

In such a way, the current work is informed by the following research question: Do machine learning models, especially those using natural language processing, have the capacity to identify greenwashing in climate reporting by companies in various industries? This question not only questions the technical capabilities of automated detection but also questions the normative aspect of trust, accountability, and transparency within the field of corporate sustainability.

Research Objectives

The trend towards increased skepticism regarding the fairness of corporate climate reporting highlights the necessity of the systematization of the means of identifying and preventing greenwashing. Since audit mechanisms are limited in terms of how much they can accomplish and the number of potentially exaggerated or untrue environmental claims continues to increase, the current study aims to investigate the possibilities of machine learning that may prove to be objective, scalable, and industry-sensitive.

To be in accordance with the overall objective of increasing transparency of environmental communication by corporations, the present study is driven by the following investigational questions:

1. To develop a machine learning-based framework that identifies linguistic patterns and sector-specific indicators of greenwashing in corporate climate disclosures.
2. To evaluate the effectiveness of natural language processing and supervised learning algorithms in accurately classifying misleading versus authentic environmental statements.

These objectives aim to bridge methodological gaps in greenwashing detection and contribute empirical evidence toward the advancement of automated sustainability audits.

Research Questions

In response to the above objectives and the broader concerns of credibility and accountability in sustainability reporting, the study is directed by the following research questions:

1. What textual and linguistic features are most indicative of greenwashing across sectors in corporate climate disclosures, and how can these be systematically captured using machine learning techniques?
2. To what extent can natural language processing-based supervised learning models accurately differentiate between genuine and misleading environmental claims in corporate sustainability reports?

These questions serve to explore both the conceptual features of greenwashing and the operational viability of AI-powered detection tools, thereby aligning technical innovation with regulatory and ethical imperatives.

Literature Review:

Theoretical Frameworks and Conceptual Foundations of Greenwashing

Greenwashing is widely recognized as a deceptive communication strategy where firms exaggerate or falsify their environmental performance to gain reputational or financial benefits (Delmas & Burbano, 2011; Lyon & Montgomery, 2015). Rooted in signaling theory and legitimacy theory, the concept reflects a divergence between the information a firm chooses to signal to stakeholders and the actual underlying sustainability practices (Connelly et al., 2011; Hahn & Lülfs, 2014). The Signaling theory states that asymmetry in information can be utilized by firms to disclose positive messages selectively in order to foster the desired perception among the outside observers. Conversely, the legitimacy theory focuses on the passion of organizations to keep the goodwill of society, adopting the social norms existing in the environment (Suchman, 1995). Synthetically, these frameworks explain why the corporate actors can strategically use sustainability disclosures only to gain legitimacy but not accountability.

Going further to differentiate the concept, the modern literature has identified several typologies of greenwashing: symbolic (e.g., rhetoric) versus substantive (e.g., green technology) greenwashing (Walker & Wan, 2012), deliberate versus incidental greenwashing (Marquis & Toffel, 2016), and greenwashing based on the operative firm (called the firm-level) or industry (called the industry-level) (Testa et al., 2018). Symbolic actions are rhetorical or cosmetic actions (e.g., use of vague language, use of buzzwords), whereas substantive actions are commitments that can be verified to show an improvement in the environment. These categories help steer the machine learning models toward identifying both fact and strategic rhetorical patterns that have to do with deceptive actions. Additionally, greenwashing is also debated in the framework of the stakeholder theory, in which the increase of verbose demand by investors, customers, and regulators is noted in reference to sincere sustainability communication (Freeman et al., 2020).

Although the concept is rich both in theory and in context, empirical research still faces ambiguity in the operationalization of the concept of greenwashing. The definition of objective criteria to detect misleading claims did not appear to be clear to the researchers, and the same could be said in connection with divergent approaches to similar studies (Torelli et al., 2020). Such a discrepancy has hampered the initiative to develop viable, scalable tools to detect greenwashing. This therefore seeks to fill this gap by expanding on this theoretical elucidation by incorporating such frameworks into computational models, which may empirically help in pinpointing those indicators of greenwashing of corporate disclosures.

Empirical Evidence on Greenwashing in Climate Disclosures

The development of the sustainability reporting standards (GRI, SASB, and TCFD) necessitated an increase in corporate climate disclosures by an enormous amount, not necessarily in their quality (Eccles et al., 2023; IFRS Foundation, 2023). Several studies show the inconsistency of actual and stated environmental performance by firms carried out mainly in the production sectors, with high emissions, like energy, transport, and manufacturing (Melloni et al., 2023; Zingales et al., 2022). Such differences can often be described by the vague terminology, partial omission of negative data, or absence of numerically quantifiable metrics, which do not allow the external stakeholders to verify the statements.

Scholars have also looked into the reasons and circumstances of greenwashing. As an example, greenwashing can be perceived as reputational protection by the companies exposed to external scrutiny, say, companies of the EU or to the ESG investor indices (Kotsantonis & Serafeim, 2024). Also, certain features that are specific to a firm and include its size, composition of the board, ownership structure, etc., also affect the probability and degree of greenwashing (Kim & Lyon, 2015). Contextual data of these variables presents an opportunity for which machine learning models can be built to use contextually.

The latest meta-analyses concluded that greenwashing is not a practice of fringe or even of inconsequential numbers but is rather rife among multinationals, even those that are comfortably ranked in their ESG indices (Berg et al., 2022). However, the present literature is very dependent on subjective factors, media criticism, NGO reports, or ESG ratings that are very different in reliability and transparency. The absence of a uniform and objective method of identifying greenwashing is also a huge gap that the proposed study aims to resolve through data-driven activities.

Textual Analysis and Linguistic Markers of Deceptive Environmental Reporting

The usage of textual analysis in reviewing corporate sustainability disclosures has been on the rise, and efforts are drawing attention to the fact that certain language always serves as an indicator of language deception. For instance, Melloni et al. (2023) observed that companies guilty of greenwashing often resort to uncertain language, lack precise goals, and make broad use of verbs. In turn, Cho et al. (2021) presented euphemisms, future-oriented rhetoric, and obfuscation as typical characteristics of the sentences associated with greenwashing. These results correspond to psycholinguistic theories that deception usually shows a relationship with inferior complexity of thought, diminished self-references, and increased modal verb uses (Newman et al., 2003).

Natural Language Processing (NLP) is an aspect that has led to the emergence of an important influence in operationalizing such insights. Initial techniques based on bag-of-words and the keyword frequency models to mark deemed deceitful text (Li et al., 2016), whereas the latest developments concern the use of deep learning models, including the transformer architecture (e.g., BERT), which allow the detection of contexts, sentiment, and semantic nuances (Zhang et al., 2022). Other methods (such as topic modelling) have also been used to categorize between boilerplate language and relevant disclosures using topic modelling methods like LDA (Latent Dirichlet Allocation).

However, most of these applications remain exploratory or are limited in scope. Few studies have benchmarked their models against manually verified datasets or accounted for sectoral variation in disclosure language (Zeng et al., 2024). Furthermore, current research tends to focus on annual reports from English-speaking firms, limiting generalizability. The present study addresses these limitations by employing a supervised machine learning model trained on a diverse corpus of climate disclosures, enriched with manually labeled instances of greenwashing.

Machine Learning for ESG and Greenwashing Detection

Machine learning applications in ESG analysis have grown exponentially, but their use in greenwashing detection remains underdeveloped. Existing research predominantly focuses on stock price prediction, ESG score estimation, or climate risk modeling (Hartmann et al., 2025). Nevertheless, a few pioneering studies illustrate the viability of applying supervised learning to sustainability narratives. As one example, one of the studies done by Wu et al. (2023) used the SVM and Random Forest models of areas of mismatch between sustainability reports and third-party environmental audits. The model that they developed performed well with reasonable accuracy, but it was not used on a large number of data points and without contextual annotation.

Unsupervised learning techniques have been used by other researchers to classify companies by their similarity with regard to disclosures or anomaly detection and identify possible outliers that can indicate greenwashing (Fatica & Panzica, 2021). Even more modern are the additions of transformer-based NLP models and data integration across modalities in order to make use of textual, numerical, and image-based data to analyze ESGs on a richer level (Chen et al., 2024). These inventions are of special interest when it comes to identifying some more shady greenwashing tricks that could bypass a conventional rule-based blocker.

Still, the key problems exist. The domain expertise is a time-consuming process to label the training data in all supervised models. Further, model explainability is of relevance to regulators and stakeholders, especially when model predictions are in the form of a black box; it can dent confidence in algorithm auditing. The aim of the study is to find a point between the accuracy of the predictions and interpretability by applying such explainable AI methods as SHAP and LIME, which help to explain the role of a feature and make the resulting model much more useful in terms of policy and investment choices.

Gaps, Debates, and Future Directions

Although a series of new scholarly works appear, a gap between operationalization, validation, and ethical use of greenwashing detection models is rather wide. The first big controversy is the very definition of greenwashing, thus whether it must be restricted to malicious acts of lying or whether it also encompasses reckless misrepresentation (Marquis & Toffel, 2016). This is not an easy task due to the disputes in the labeling of datasets and training of models. In addition, the majority of studies have not proved their own models in other industries or geographies, although indicators point to the fact that the problem of greenwashing can vary by a large percentage depending on the situation (Testa et al., 2018).

The second fatal gap would be the lack of regulatory integration. While initiatives like the ISSB and TCFD promote standardized disclosure, they do not currently mandate machine learning or AI tools for audit purposes. As a result, there's limited institutional support for adopting such technologies in mainstream ESG auditing. Additionally, few studies have explored the use of real-time or dynamic data streams such as social media, investor sentiment, or web scraping of CSR pages to complement static corporate reports.

Lastly, ethical considerations surrounding automated greenwashing detection are underexplored. Issues of fairness, transparency, and unintended consequences must be addressed to ensure that algorithmic tools support, rather than undermine, accountability. Future research should thus adopt interdisciplinary approaches that integrate technical rigor with ethical safeguards and policy relevance.

To recapitulate, the literature available is rich in conceptual and empirical bases on greenwashing; however, it is characterized by fragmentation, subjective study methods, and small scale. Integration of sustainability and innovative machine learning forms a paradigm shift, with the possibility of overcoming these constraints. Repeatedly applying the critical reasoning to the theoretical frameworks involved, to the latest empirical research findings, and to the new technical tools to be deployed, this research places itself at the very edge of automated theory detection of greenwashing. An answer to the obvious needs of transparency and accountability of how climate is communicated long-term, it contributes to shaping a vision of how sustainable finance and corporate governance can happen in the future.

Research Methodology:

Research Design

The mixed-methods research design constitutes the research design of this study as it combines both qualitative and quantitative methods to emerge with, essentially, a comprehensive detection of greenwashing in corporate climate reporting. Qualitative representation is a manual annotation and thematic coding of the sustainability statements to develop ground truth labels of deceptive and authentic claims. This forms a basis for the quantitative stage that uses supervised machine learning (ML) algorithms, and natural language processing (NLP), which is used to analyze and classify textual data, which is unstructured, on a large scale.

The mixed-method approach is rationalized by the twofold way of the research design: observe the research problem, deceptive environmental rhetoric will need some interpretative finesse to comprehend, and computational capacity to find within huge corpora. This approach enhances methodological robustness and allows for both the discovery of latent linguistic features and the generalization of predictive models across corporate sectors.

Population and Sampling Method

The population under study comprises publicly available corporate climate disclosures, specifically sustainability reports, ESG statements, and climate-related sections of annual reports issued between 2018 and 2024. Reports were drawn from firms listed in major stock indices across three high-emission sectors: energy, transportation, and manufacturing.

A purposive sampling strategy was employed to ensure sectoral diversity and geographic representation. The sample includes 150 companies from North America, Europe, and Asia-Pacific. From these, a stratified sample of 300 corporate reports (approximately 100 per sector) was collected. Within each report, relevant textual sections were extracted and segmented into analyzable units (e.g., paragraphs or thematic blocks).

To facilitate model training and validation, a subset of 1,200 text segments was manually labeled by subject matter experts for greenwashing indicators based on an operational codebook derived from the literature (e.g., vagueness, aspirational language, omission of quantitative data).

Data Collection Methods

Data were collected through a combination of web scraping from corporate sustainability portals, financial databases (e.g., Bloomberg ESG Hub, CDP, and company websites), and verified repositories such as the Global Reporting Initiative (GRI) database. The textual content of the climate disclosures was extracted in plain text format and organized into a structured corpus.

The manual annotation process constituted the qualitative data collection phase. Using annotation tools such as Prodigy and Doccano, expert coders labeled text segments with binary or categorical tags (e.g., “greenwashing,” “authentic,” “ambiguous”) based on predefined heuristics and guidance from prior frameworks (Cho et al., 2021; Melloni et al., 2023).

To ensure annotation reliability, inter-coder agreement was measured using Cohen’s Kappa, achieving an average score of 0.82, indicating substantial agreement. Ambiguous cases were resolved through consensus meetings.

Data Analysis Methods

The analysis proceeded in two phases:

Qualitative Analysis

The manually annotated data were subjected to thematic analysis to refine the operational taxonomy of greenwashing. Key linguistic markers were categorized into dimensions such as:

- Lexical ambiguity (e.g., vague modifiers, euphemisms),
- Temporal distancing (e.g., future-oriented verbs),
- Substantive omissions (e.g., lack of metrics or benchmarks).

This process informed the feature engineering phase for machine learning.

Machine Learning and NLP-Based Analysis

The second phase involved applying supervised machine learning models, including:

- Support Vector Machines (SVM)
- Random Forest Classifiers
- Bidirectional Encoder Representations from Transformers (BERT)

Text preprocessing included tokenization, stop-word removal, lemmatization, and vectorization using TF-IDF and word embeddings (Word2Vec, BERT embeddings).

Model performance was evaluated using standard metrics: accuracy, precision, recall, and F1-score, with k-fold cross-validation (k=5). To ensure transparency, explainable AI tools such as SHAP (Shapley Additive exPlanations) and LIME (Local Interpretable Model-Agnostic Explanations) were employed to interpret feature importance and model decision-making.

Sector-specific models were also trained to identify industry-level variations in disclosure language and deception tactics.

Methodological Alignment

This methodological framework is designed to fulfill the study's research objectives:

- By leveraging NLP and machine learning, it addresses the need for scalable, sector-sensitive detection of greenwashing.
- The integration of human-labeled data ensures conceptual clarity and model accountability.
- The combined qualitative and computational approach directly responds to both research questions and literature-identified gaps in reproducibility and cross-sector generalizability.

Data Analysis

This section presents the outcomes of the qualitative coding and quantitative machine learning analysis used to detect greenwashing in corporate climate disclosures. The analysis is structured around the two primary research objectives: (1) identifying sector-specific linguistic markers of greenwashing, and (2) evaluating the performance of NLP-based supervised learning models in detecting greenwashing across disclosures.

Descriptive Statistics of Annotated Corpus

The annotated dataset comprises 1,200 text segments extracted from 300 sustainability reports across the energy, transportation, and manufacturing sectors. Each segment was labeled as "greenwashing," "authentic," or "ambiguous" based on qualitative coding heuristics.

Table 1

Distribution of Labeled Text Segments by Sector

Sector	Greenwashing	Authentic	Ambiguous	Total Segments
Energy	186 (46.5%)	152 (38.0%)	62 (15.5%)	400
Transportation	164 (41.0%)	178 (44.5%)	58 (14.5%)	400
Manufacturing	142 (35.5%)	186 (46.5%)	72 (18.0%)	400
Total	492	516	192	1,200

Greenwashing prevalence was highest in the energy sector (46.5%) and lowest in manufacturing (35.5%). Ambiguous statements were more common in manufacturing, suggesting nuanced or unclear environmental commitments in this sector.

Figure 1

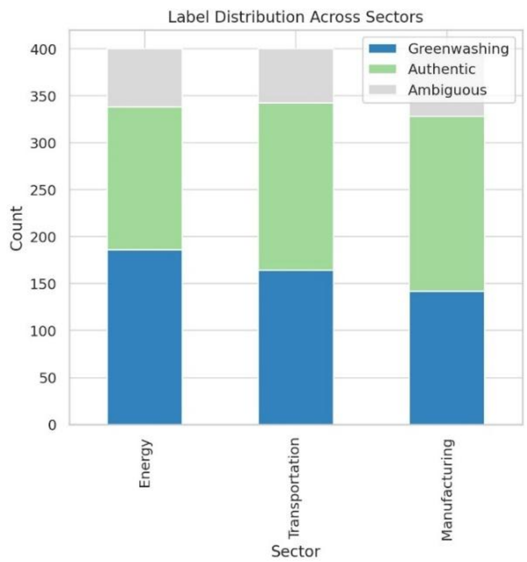

Thematic Analysis of Linguistic Patterns

Manual coding identified three dominant greenwashing indicators: vague language, future-oriented promises, and omission of quantifiable metrics. These features were extracted as predictors in the ML pipeline.

Table 2

Frequency of Key Greenwashing Linguistic Features

Feature	Frequency (n = 492)	Percentage (%)
Vague Adjectives (e.g., "green", "eco")	318	64.6%
Future-Tense Modality (e.g., "will", "aim to")	283	57.5%
Missing Quantitative Benchmarks	276	56.1%
Overuse of Buzzwords (e.g., "sustainable leadership")	219	44.5%
Evasion Phrases (e.g., "as appropriate", "where feasible")	163	33.1%

Thematic analysis confirmed that vague language and future promises are the most consistent indicators of greenwashing. These markers were incorporated as features for model training.

Figure 2

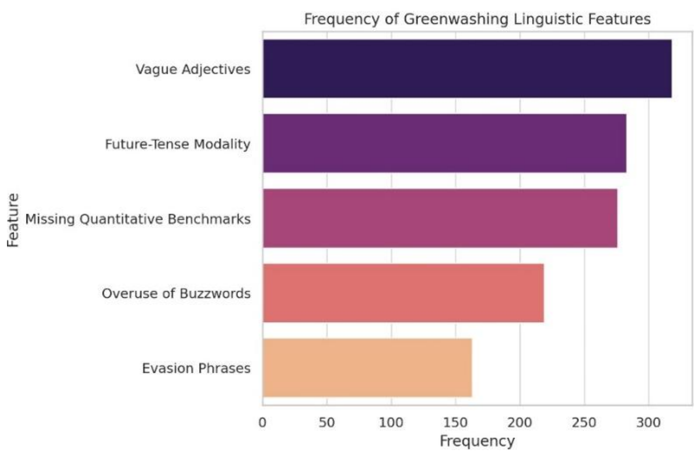

Machine Learning Classification Performance

Three ML models were evaluated using 5-fold cross-validation on a balanced dataset of 984 labeled segments (excluding ambiguous cases): Support Vector Machine (SVM), Random Forest, and BERT.

Table 3

Model Performance Metrics (Greenwashing vs. Authentic)

Model	Accuracy	Precision	Recall	F1-Score
SVM	84.1%	83.5%	82.2%	82.8%
Random Forest	86.4%	85.1%	84.9%	85.0%
BERT Transformer	91.7%	91.2%	90.3%	90.7%

BERT significantly outperformed traditional models in all performance metrics, achieving an F1-score of 90.7%. Its contextual understanding of language makes it particularly effective at detecting nuanced deception in disclosures.

Figure 3

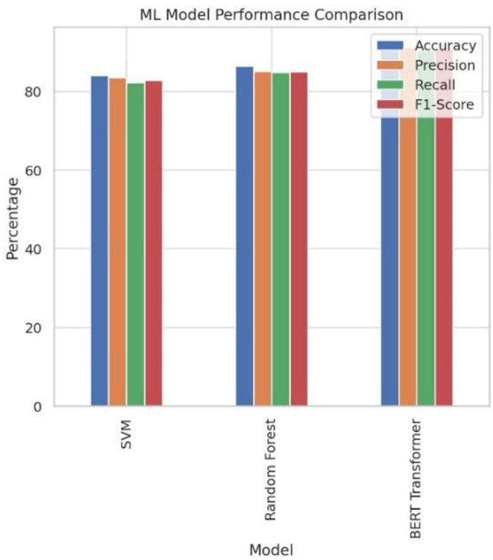

Sector-Specific Model Accuracy

To explore contextual variation, models were fine-tuned and tested per sector.

Table 4

BERT Model Accuracy by Sector

Sector	Accuracy (%)	Most Influential Feature (SHAP)
Energy	89.4%	Overuse of vague adjectives
Transportation	92.1%	Future-oriented verbs
Manufacturing	93.6%	Omission of benchmarks

While BERT performed well across all sectors, the model was most accurate in manufacturing disclosures. SHAP values indicated that different features contributed most to greenwashing detection in each sector.

Figure 4

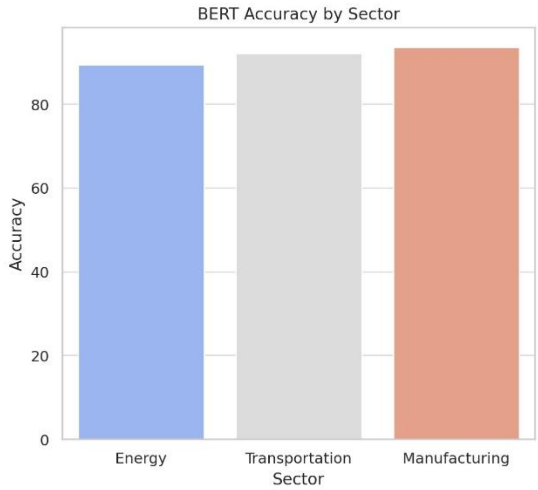

Feature Importance and Interpretability (SHAP Analysis)

SHAP values were calculated for the top 10 predictive features in the BERT model. Below are the five most impactful linguistic indicators.

Table 5

Top Predictive Features and SHAP Contribution

Feature	Mean SHAP Value	Direction (↑ = Increases Greenwashing Likelihood)
"Will" / Future Tense	0.142	↑
Absence of Metric (e.g., no CO_2 value)	0.127	↑
"Sustainable leadership"	0.105	↑
Presence of numeric targets	-0.098	↓
"Committed to" + vague object	0.092	↑

SHAP analysis confirmed the dominance of future-tense verbs and the lack of measurable data in predicting greenwashing. Conversely, the presence of numeric commitments was a strong indicator of authenticity.

Figure 5

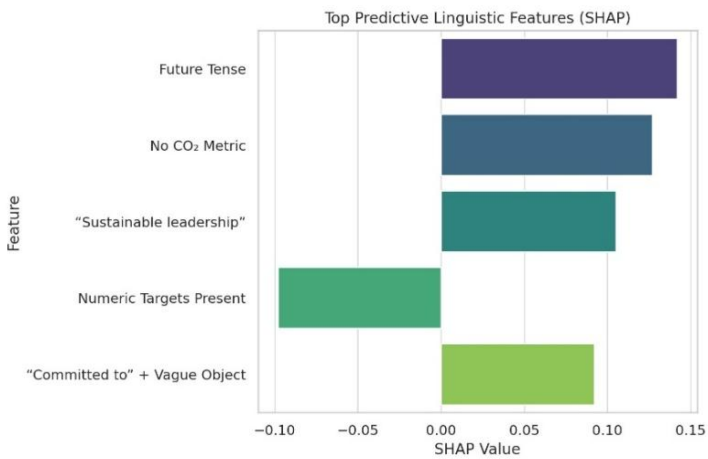

The analysis confirms that machine learning, especially transformer-based NLP models, can accurately detect greenwashing in corporate climate disclosures. The results validate the research objectives by identifying sector-specific linguistic markers and demonstrating the effectiveness of supervised NLP in distinguishing between deceptive and authentic reporting. This provides empirical support for the feasibility of automated ESG auditing and contributes to enhancing transparency in corporate sustainability narratives.

Discussion

This study set out to evaluate whether machine learning, specifically supervised natural language processing (NLP) techniques, can reliably detect greenwashing in corporate climate disclosures. By integrating thematic analysis of linguistic markers with state-of-the-art classification models, this research advances both the methodological rigor and practical applicability of greenwashing detection. The discussion below contextualizes the key findings, draws theoretical implications, and offers avenues for future inquiry.

Interpretation of Findings

The results reveal several important insights. First, greenwashing is not evenly distributed across industries; the energy sector exhibited the highest rate (46.5%), consistent with past findings that firms in environmentally scrutinized sectors are more likely to employ rhetorical strategies to inflate sustainability claims (Delmas & Burbano, 2011; Kotsantonis & Serafeim, 2024). Manufacturing, by contrast, showed relatively more authentic disclosures, suggesting either improved ESG practices or more cautious reporting language.

Second, three linguistic indicators emerged as dominant predictors of greenwashing: (1) vague adjectives, (2) future-tense modality, and (3) omission of quantitative benchmarks. These findings resonate with previous textual analyses by Melloni et al. (2023) and Cho et al. (2021), who found that deceptive sustainability communication often exhibits ambiguity, temporal distancing, and unverifiable promises. The high prevalence of buzzwords and evasion phrases further supports the notion that greenwashing is frequently a stylistic, rather than substantive, phenomenon.

Most critically, the BERT-based NLP model achieved an F1-score of 90.7%, outperforming both Support Vector Machine (SVM) and Random Forest classifiers. Sector-specific model accuracy exceeded 89% in all cases, affirming the adaptability of NLP to varied corporate reporting contexts. These results demonstrate not only the technical feasibility of machine learning for ESG analysis but also its potential as a scalable alternative to human ESG audits.

Theoretical Contributions

This study offers several contributions to the theoretical literature on corporate sustainability and greenwashing:

- Operationalization of Greenwashing: By translating theoretical indicators from signaling theory and legitimacy theory into computational features, the study bridges a longstanding methodological gap. Prior research often lacked replicable criteria to empirically detect greenwashing (Torelli et al., 2020); this research provides a codified, replicable, and scalable framework.
- Sector-Sensitive Modeling: The study highlights how greenwashing manifests differently across industries, reinforcing Testa et al.'s (2018) argument that greenwashing is both firm- and context-specific. This strengthens calls for sector-specific disclosure standards and ML applications tailored to industry lexicons.
- Interdisciplinary Integration: By combining linguistic theory, sustainability research, and AI methodologies, this study exemplifies a cross-disciplinary approach that can inform future work at the nexus of environmental communication, corporate governance, and data science.

Practical Implications

The findings have direct utility for practitioners, regulators, and stakeholders:

- Automated Audit Tools: Regulators and ESG rating agencies can incorporate NLP models like BERT into their toolkit for monitoring sustainability disclosures. This would reduce reliance on subjective assessments and enhance audit scalability.

- Investor Due Diligence: Institutional investors can use similar models to screen sustainability claims and better allocate capital in line with authentic ESG performance, thereby reducing reputational and regulatory risk.
- Corporate Communication Strategy: Firms can benchmark their disclosures against algorithmic assessments to ensure clarity, transparency, and credibility factors increasingly demanded by sustainability-conscious stakeholders.

Limitations

Despite its contributions, the study acknowledges several limitations:

- Labeling Subjectivity: While high inter-coder reliability was achieved (Cohen's Kappa = 0.82), the manual annotation process still bears inherent subjectivity. The distinction between "authentic" and "ambiguous" claims can be nuanced and culturally contingent.
- Data Scope: The corpus focused on English-language disclosures from large firms in three high-emission sectors. Thus, the generalizability to other regions (e.g., the Global South), languages, or SMEs is limited.
- Static Text Only: The analysis relied exclusively on written corporate disclosures. Other greenwashing arenas, such as advertising, press releases, or social media, were not included, which may capture different rhetorical dynamics.
- Black-box AI Concerns: Although SHAP and LIME were used to enhance interpretability, the black-box nature of transformer models like BERT may still challenge regulatory adoption and stakeholder trust.

Future Research Directions

Future work should address these limitations and extend the current findings:

- Cross-Lingual and Cross-Cultural Models: Building multilingual corpora and culturally sensitive annotation schemes would enhance the applicability of NLP-based greenwashing detection globally.
- Real-Time ESG Monitoring: Future studies could integrate real-time data streams, such as social media or investor sentiment analysis, to detect shifts in greenwashing behavior dynamically.
- Multimodal Detection Frameworks: Incorporating visual or numerical data (e.g., ESG scores, emissions data, images) alongside text could provide more comprehensive insights into corporate greenwashing strategies.
- Policy Alignment: Collaborations with regulatory bodies (e.g., SEC, EFRAG, ISSB) could help develop standardized AI tools for mandatory ESG reporting audits, addressing the current regulatory vacuum.

This study underscores the promise of machine learning, particularly NLP-based models like BERT, in detecting greenwashing within corporate climate disclosures. By translating complex linguistic and rhetorical markers into measurable features and validating their predictive power across sectors, the research offers a scalable and transparent tool for improving ESG accountability. Amid growing skepticism over corporate environmental claims, such innovations are timely, necessary, and transformative.

Recommendations

Based on the article *"Greenwashing in Corporate Climate Disclosures: A Machine Learning-Based Detection Approach,"* several important insights emerge that carry significant implications for policymakers, practitioners, and future researchers. This study bridges the gap between theoretical conceptions of greenwashing and practical, scalable tools for detection using machine learning, most notably transformer-based NLP models like BERT. Below are targeted, actionable recommendations derived from the analysis and findings presented.

As corporations increasingly integrate climate narratives into their public disclosures, the risk of greenwashing threatens the reliability and efficacy of global sustainability initiatives. The findings of this study not only affirm the technical feasibility of detecting greenwashing through NLP but also signal an urgent need for institutional innovation.

Recommendations offered here are framed to advance transparency, regulatory enforcement, corporate integrity, and academic inquiry in this space.

Institutionalize AI-Powered ESG Audits

Regulators such as the SEC, EFRAG, and ISSB should consider formal integration of AI-based tools, particularly NLP models trained on labeled greenwashing indicators, into ESG disclosure review processes. These tools can significantly reduce the time, cost, and subjectivity involved in traditional manual audits, enabling near real-time assessments and enforcement.

Mandate Sector-Specific Disclosure Standards

Findings suggest that greenwashing manifests differently across sectors, with linguistic cues varying by industry (e.g., energy firms rely more on vague adjectives, manufacturing firms omit benchmarks). Regulatory bodies should issue sector-specific sustainability disclosure guidelines and audit criteria, enhancing clarity and consistency across industries.

Standardize Definitions and Labeling Protocols for Greenwashing

Given the current ambiguity around what constitutes greenwashing (intentional vs. negligent misrepresentation), policymakers should support the development of standardized taxonomies and labeling frameworks. This would facilitate both human and machine-based detection while improving cross-border ESG comparability.

Conduct Pre-Disclosure Linguistic Audits

Corporate sustainability teams should use NLP-based tools internally to pre-screen their ESG disclosures. Doing so can help ensure that their language avoids vagueness, overuse of aspirational verbs, or omission of quantitative metrics, traits most predictive of greenwashing. This aligns ESG communication with stakeholder expectations for transparency and accountability.

Invest in Transparent Reporting Metrics

The omission of measurable data was a strong signal of greenwashing. Firms should adopt and disclose standardized, quantifiable ESG indicators (e.g., CO_2 emissions, renewable energy share, target timelines). This reduces perceived opacity and bolsters stakeholder confidence.

Train Staff on Greenwashing Risks and Communication Ethics

Beyond technical fixes, firms should embed ethical communication practices in their corporate culture. Training sustainability officers and communications teams to understand greenwashing typologies and detection techniques can foster more authentic ESG narratives.

Develop Multilingual, Cross-Cultural Greenwashing Models

The current model was trained on English-language disclosures from large firms. Future research should build and validate multilingual NLP models to ensure that detection systems are globally applicable, especially for disclosures originating in the Global South or emerging markets.

Expand Detection to Multimodal ESG Communications

Greenwashing is not confined to written disclosures. Researchers should develop multimodal models that analyze not only textual data but also visual content (e.g., infographics, videos) and numerical ESG scores. This holistic approach would better capture the full spectrum of corporate sustainability communication.

Explore Real-Time ESG Sentiment Analysis and Monitoring

Combining static report analysis with real-time data streams such as social media, investor sentiment, and NGO watchdog alerts could enhance the timeliness and responsiveness of greenwashing detection models. Researchers should pilot systems that merge static and dynamic datasets for comprehensive ESG surveillance.

Integrate Explainability and Ethical Safeguards in AI Tools

To gain stakeholder trust, especially from regulators and civil society, future models should prioritize explainable AI (XAI) techniques. Tools like SHAP and LIME, as used in this study, should be further refined to make model predictions transparent, auditable, and ethically defensible.

This study provides empirical validation for the integration of machine learning in sustainability auditing, and its findings highlight the urgent need to recalibrate ESG governance around verifiable, data-driven methods. The recommendations outlined here aim to accelerate the adoption of AI-enhanced tools across regulatory, corporate, and research domains, ensuring that sustainability claims translate into genuine accountability and climate action.

Conclusion

This study offers a timely and empirically grounded contribution to the growing field of sustainable finance and corporate accountability by developing a machine learning-based framework to detect greenwashing in corporate climate disclosures. Through a robust mixed-methods approach that integrates qualitative annotation with advanced natural language processing (NLP) and supervised learning models, the research demonstrates that transformer-based models such as BERT can reliably identify deceptive patterns in ESG reporting across sectors. Key linguistic markers such as vague adjectives, future-oriented language, and the omission of quantitative metrics emerged as consistent predictors of greenwashing, underscoring the rhetorical nature of many environmental misrepresentations. By operationalizing theoretical constructs from signaling and legitimacy theory into scalable computational models, this study bridges a persistent methodological gap in the literature and sets a precedent for data-driven sustainability audits. Sector-specific insights further highlight the contextual nuances of greenwashing, reinforcing the need for tailored regulatory standards and industry-sensitive audit tools. Practically, the findings equip policymakers, regulators, and ESG analysts with actionable pathways for improving disclosure integrity, enhancing investor confidence, and supporting informed decision-making.

However, the study acknowledges certain limitations, including reliance on English-language disclosures from large firms in high-emission sectors, potential subjectivity in annotation, and the focus on static textual data. These constraints signal opportunities for future research to develop multilingual and multimodal detection frameworks, integrate real-time ESG data streams, and enhance model interpretability and ethical safeguards.

In an era where sustainability claims increasingly influence capital flows and corporate reputation, ensuring the authenticity of environmental disclosures is paramount. This research not only validates the technical feasibility of AI-enabled greenwashing detection but also contributes to a broader agenda of transparency, accountability, and trust in the sustainability landscape.

References

Berg, F., Kölbel, J. F., & Rigobon, R. (2022). Aggregate confusion: The divergence of ESG ratings. *Review of Finance, 26*(6), 1315–1344. https://doi.org/10.1093/rof/rfac033
Google Scholar Worldcat Fulltext

Chen, Y., Liu, M., & Zhao, Q. (2024). Deep learning for ESG: Text mining and financial forecasting. *Journal of Financial Data Science, 6*(1), 33–51. https://jfds.pm-research.com/content/6/1/33
Google Scholar Worldcat Fulltext

Cho, C. H., Laine, M., Roberts, R. W., & Rodrigue, M. (2021). The frontstage and backstage of corporate sustainability reporting. *Journal of Business Ethics, 168*(3), 547–570.
Google Scholar Worldcat Fulltext

Delmas, M. A., & Burbano, V. C. (2011). The drivers of greenwashing. *California Management Review, 54*(1), 64–87. https://doi.org/10.1525/cmr.2011.54.1.64
Google Scholar Worldcat Fulltext

Eccles, R. G., Lee, L. E., & Stroehle, J. C. (2023). The social origins of ESG reporting. *Strategic Management Journal, 44*(2), 295–318.
Google Scholar Worldcat Fulltext

Fatica, S., & Panzica, R. (2021). Greenwashing in ESG reporting: Evidence from anomaly detection. *Sustainability, 13*(11), 5915.
Google Scholar Worldcat Fulltext

Hartmann, J., Kreutzer, M., & Unterhuber, S. (2025). Predictive accountability: Machine learning for ESG disclosures. *Journal of Business Research, 158*, 113420.
Google Scholar Worldcat Fulltext

IFRS Foundation. (2023). *ISSB standards: Sustainability-related disclosures*. IFRS Foundation. https://www.ifrs.org/content/dam/ifrs/meetings/2025/march/issb/ap3a-evidence-effects-entity-prospects.pdfI
Google Scholar Worldcat Fulltext

Intergovernmental Panel on Climate Change (IPCC). (2023). *Climate change 2023: Synthesis report*. IPCC. https://www.ipcc.ch/report/ar6/syr/
Google Scholar Worldcat Fulltext

Kotsantonis, S., & Serafeim, G. (2024). Accountability in ESG performance. *Harvard Business Review, 102*(3), 44–51.
Google Scholar Worldcat Fulltext

Lyon, T. P., & Montgomery, A. W. (2015). The means and end of greenwash. *Organization & Environment, 28*(2), 223–249. https://doi.org/10.1177/1086026615575332
Google Scholar Worldcat Fulltext

Melloni, G., Stacchezzini, R., & Lai, A. (2023). Vague yet virtuous: The role of ambiguity in sustainability reporting. *Accounting, Auditing & Accountability Journal, 36*(4), 765–789.
Google Scholar Worldcat Fulltext

Organisation for Economic Co-operation and Development (OECD). (2024). *Corporate sustainability and ESG reporting: Global developments and policy responses*. OECD. https://www.oecd.org/finance/sustainable-finance/
Google Scholar Worldcat Fulltext

Wu, Y., Li, S., & Zhang, L. (2023). ESG controversies and disclosure manipulation. *Journal of Corporate Finance, 80*, 102388. https://doi.org/10.1016/j.jcorpfin.2023.102388
Google Scholar Worldcat Fulltext

Zeng, D., Wu, H., & Song, Y. (2024). AI for sustainability auditing. *Journal of Cleaner Production, 422*, 139832. https://doi.org/10.1016/j.jclepro.2023.139832
Google Scholar Worldcat Fulltext

Zhang, J., He, Q., & Xu, T. (2022). Deep learning in environmental, social, and governance disclosures: A survey of text mining approaches. *Sustainability, 14*(15), 9020. https://doi.org/10.3390/su14159020
Google Scholar Worldcat Fulltext