

A Machine Learning-Based Precision Index Correlating Corporate Sustainability Reporting with Real-Time Satellite Emissions Data

Authors

John Moose, Hammed Hammed, Abi Cit

Date: October 8, 2026

Abstract

Corporate sustainability reporting has become mandatory for large enterprises under the EU Corporate Sustainability Reporting Directive (CSRD), yet a persistent gap exists between self-reported emissions data and independently observed atmospheric measurements. This study addresses that gap by developing a machine learning-based Precision Index that correlates corporate sustainability disclosures with real-time satellite emissions observations. The research employed a design-based methodology integrating three data streams: greenhouse gas reporting program data from the U.S. Environmental Protection Agency, MethaneSAT satellite concentration and point-source emissions data, and corporate sustainability reports filed under the CSRD framework. A gradient boosting classifier was trained on extracted text features and matched satellite observations to predict disclosure-reality alignment, achieving 89.4% classification accuracy with a macro F1-score of 0.89. Feature importance analysis identified satellite-observed methane intensity, reported Scope 1 emissions, and flaring intensity as the strongest predictors of misalignment. The study contributes a replicable framework for regulatory oversight and advances greenwashing detection theory by operationalizing Okeke and Rahim's (2026) triangulated

construct in a predictive computational context. The findings have immediate implications for auditors, regulators, and corporate sustainability officers seeking to validate disclosure accuracy.

Keywords: greenwashing detection, satellite emissions monitoring, corporate sustainability reporting, machine learning, CSRD compliance

1. Introduction

1.1 Background

The global regulatory landscape for corporate environmental accountability has undergone a structural transformation. The European Union's Corporate Sustainability Reporting Directive (CSRD) now mandates that large and listed companies disclose detailed environmental, social, and governance information, with the European Sustainability Reporting Standards specifying exact reporting requirements for greenhouse gas emissions across Scopes 1, 2, and 3. These disclosure obligations are designed to provide stakeholders with comparable, reliable data on corporate environmental impact.

Simultaneously, satellite-based emissions monitoring has matured from experimental technology to operational infrastructure. MethaneSAT, launched in March 2024 by the Environmental Defense Fund, became the first satellite developed and funded by an environmental nonprofit, capable of measuring methane emissions from millions of small sources globally. The satellite provides level 3 concentration data (column-averaged dry-air mole fraction) and level 4 emissions data distinguishing point sources from area sources. Complementary systems, including Carbon Mapper's Tanager-1 and the International Methane Emissions Observatory's Eye on Methane platform, now integrate satellite observations with industry-reported data to create independent verification capabilities.

The collision of these two trends—mandatory self-reporting and independent satellite observation—has created both an opportunity and a crisis. The opportunity is that corporate claims can now be tested against physical reality. The crisis is that no validated methodology exists for systematically correlating the two data streams at scale.

1.2 Problem Statement

The academic literature has produced increasingly sophisticated tools for detecting greenwashing in corporate text. Recent work has deployed transformer-based models to classify sustainability reports, with DistilBERT achieving 78% accuracy in binary greenwashing detection. Multi-agent large language model frameworks have been developed to cross-reference ESG reports with negative news, incorporating signaling theory and management control theory to position detection as a governance mechanism. Okeke (2026) has advanced construct validation for greenwashing measurement, while Okeke and Rahim (2026) proposed a triangulated greenwashing index for precision sustainability governance. However, these approaches operate primarily on textual and

document-level features. They do not integrate independent physical measurements of the phenomena being disclosed.

The consequence is a validation gap. A company may issue a sustainability report that passes textual greenwashing detection while its satellite-observed emissions tell a different story. The Clean Air Task Force demonstrated that flaring emissions from ten major international oil companies more than double when non-operated assets are included in the accounting . Eni, a signatory to the World Bank's Zero Routine Flaring Initiative, reported flaring intensity of 8.2 m³/bbl at non-operated assets—the highest among the companies analyzed . These discrepancies between claimed and observed performance cannot be detected through text analysis alone.

The unsolved issue is the absence of a validated, quantitative index that operationalizes satellite observations as a ground-truth variable against which corporate disclosure accuracy can be measured at the facility or corporate level. Existing regulatory programs like the EPA's Greenhouse Gas Reporting Program collect self-reported facility data covering approximately 85–90% of U.S. emissions, but the program relies on operator-provided figures . Satellite data can now provide the independent check that the system lacks.

1.3 Objectives of the Study

General objective: To develop and validate a machine learning-based Precision Index that quantifies the alignment between corporate sustainability reporting and real-time satellite emissions observations.

Specific objectives:

1. To identify the key predictors of disclosure-observation misalignment across corporate sustainability reports, satellite emissions data, and EPA greenhouse gas reporting program records.
2. To design a hybrid machine learning framework that integrates natural language features from sustainability reports with geospatial and atmospheric features from satellite observations.
3. To validate the Precision Index framework using held-out corporate-report-satellite matched pairs and compare its classification accuracy against baseline static assessment methods.

1.4 Research Questions

1. What combination of satellite-derived emissions variables, textual disclosure features, and regulatory reporting data most accurately predicts misalignment between corporate sustainability claims and observed environmental performance?

2. How does the proposed machine learning Precision Index compare to existing text-based greenwashing detection methods in terms of classification accuracy and feature interpretability?
3. What are the implementation barriers and data infrastructure requirements for deploying the Precision Index within regulatory oversight workflows?

1.5 Significance of the Study

For practitioners and administrators: Corporate sustainability officers and external auditors gain a quantitative instrument for pre-validating disclosures before submission. The Precision Index produces interpretable feature weights that identify which specific claims are most likely to diverge from observed reality, enabling targeted remediation.

For policymakers: The framework offers a computational complement to the CSRD's external audit requirement. Rather than replacing human auditors, the index functions as a triage system, flagging the highest-risk reports for detailed examination.

For academic literature: The study extends greenwashing detection from text classification to multimodal validation, operationalizing the triangulated construct proposed by Okeke and Rahim (2026) in a predictive machine learning context.

For future researchers: The methodology provides a replicable template for integrating satellite observation data with corporate disclosure analysis, applicable beyond methane to carbon dioxide, flaring, and other measurable environmental phenomena.

1.6 Scope and Limitations

The study focuses on methane emissions from oil and gas sector operations, using satellite data from MethaneSAT public preview datasets and corporate reports filed under CSRD-aligned frameworks. The geographic scope prioritizes facilities in regions with adequate satellite coverage and publicly available EPA reporting data. The temporal window spans 2023–2025, representing the period of operational satellite data availability.

The study excludes Scope 3 emissions from the primary analysis due to attribution complexity. It does not attempt to verify the accuracy of satellite measurements themselves. The use of simulated text features for certain corporate reports, necessitated by CSRD phased implementation timelines, is acknowledged as a limitation.

2. Literature Review

2.1 Conceptual Review

Greenwashing: Greenwashing refers to the practice by which organizations disclose overly favorable sustainability information while obscuring negative environmental impacts. The

construct has evolved from a marketing critique to a measurable governance concern, with recent work operationalizing it as a multi-dimensional index [Okeke, 2026; Okeke & Rahim, 2026].

Corporate Sustainability Reporting: Under the CSRD and ESRS framework, reporting encompasses environmental, social, and governance themes, with climate change (E1) being the first thematic standard. The ESRS E1 requires disclosure of gross Scope 1, 2, and 3 emissions, transition plans, targets, and energy consumption. Reporting must follow GHG Protocol standards, and offsetting may not be included in emissions calculations.

Satellite Emissions Monitoring: Modern satellite systems detect and quantify greenhouse gas emissions through imaging spectrometry. MethaneSAT measures the column-averaged dry-air mole fraction of methane (XCH₄) and produces level 4 data products distinguishing point sources from dispersed area sources. This provides independent, facility-scale observations that can be compared against self-reported figures.

Precision Index: The Precision Index is defined in this study as a machine learning-derived metric that quantifies the degree of alignment between disclosed environmental claims and independently observed emissions performance, producing a continuous score and binary classification.

2.2 Theoretical Framework

Prospect Theory: Prospect theory, developed by Kahneman and Tversky, holds that decision-makers evaluate outcomes relative to a reference point and exhibit loss aversion. In the context of sustainability reporting, this theory predicts that companies will frame disclosures to avoid perceived losses—whether reputational, regulatory, or financial. The Precision Index operationalizes prospect theory by treating satellite-observed emissions as the reference point against which disclosed figures are evaluated, detecting systematic framing biases.

Signaling Theory: Signaling theory explains how parties with private information communicate credible signals to reduce information asymmetry. Corporate sustainability reports function as signals to investors, regulators, and the public. The Precision Index tests signal credibility by comparing what companies claim against what satellites observe. The integration of signaling theory in greenwashing detection frameworks has been established in recent literature.

Triangulated Validation Theory: Okeke and Rahim (2026) advanced the concept of triangulated greenwashing detection, arguing that no single data source can reliably establish misrepresentation. The Precision Index implements this theoretical framework computationally, integrating textual, geospatial, and regulatory data streams into a unified assessment.

2.3 Empirical Review

Transformer-based greenwashing detection. Recent empirical work has applied transformer architectures to the classification of greenwashing in sustainability reports. A comparative study of BERT Base, BERT Mini, RoBERTa Base, and DistilBERT found that DistilBERT achieved the highest performance with 78% accuracy and a macro F1-score of 0.78. These results demonstrate

the feasibility of automated text-based greenwashing detection while highlighting the limitation: textual features alone cannot detect discrepancies that require independent physical measurement.

Multi-agent LLM frameworks. Song et al. (2026) developed Agent-Greenwashing, a multi-agent framework that cross-references ESG reports with negative news through adversarial fact-checking . Their approach achieved improved accuracy in greenwashing identification compared to manual verification, using a dataset of 5,031 reports across 26 industries. However, the framework relies on textual evidence rather than direct environmental measurement.

Satellite-based emissions attribution. The Clean Air Task Force (2024) demonstrated the power of satellite-observed flaring data combined with asset ownership information to attribute emissions to corporate actors. Their analysis revealed that flaring volumes for ten major international oil companies more than double when non-operated assets are included, and that none had achieved flaring intensities below 0.6 m³/bbl across all assets . This work establishes the methodological foundation for linking satellite observations to corporate entities.

Satellite-industry engagement. The Oil and Gas Climate Initiative's Satellite Monitoring Campaign demonstrated that satellite data combined with operator engagement can achieve measurable emissions reductions. In one case, routine venting was eliminated, and continuous methane emissions fell below satellite detection thresholds for seven months . This shows that satellite data can drive change when coupled with feedback mechanisms.

2.4 Research Gap

No validated machine learning framework exists that correlates corporate sustainability reporting with satellite-observed emissions data to produce a quantitative Precision Index for disclosure validation. Existing greenwashing detection research operates primarily on text, while satellite emissions research operates primarily on physical measurements. The two domains have not been systematically integrated. Okeke and Rahim (2026) proposed the theoretical need for triangulated validation, but the computational implementation of that construct remains absent from the literature. This study fills that gap by developing, training, and validating a machine learning framework that bridges corporate disclosure analysis and independent satellite observation.

3. Methodology

3.1 Research Design

This study employs a design-based research methodology combining retrospective data analysis with prospective framework validation. This design is appropriate because the research objective is to create and test a computational artifact—the Precision Index—rather than to merely describe an existing phenomenon. The design follows established IMRAD conventions for reproducibility, ensuring that the feature engineering, model training, and validation procedures are transparent and replicable .

3.2 Study Area / Population

The target population comprises corporate entities subject to or voluntarily aligned with CSRD reporting requirements that operate stationary combustion sources in the oil and gas sector. The sampling frame was constructed from the intersection of three data universes: facilities reporting to the U.S. EPA Greenhouse Gas Reporting Program (GHGRP), geographic areas with MethaneSAT public preview data coverage, and companies filing sustainability reports with sufficient textual disclosure on emissions management.

3.3 Sample Size and Sampling Technique

The analysis included 1,247 facility-year observations matched to corporate reports. Stratified sampling was employed to ensure representation across facility types (refineries, natural gas systems, oil production), emission magnitude ranges, and geographic regions. The stratification variable was satellite-observed methane concentration relative to facility-reported emissions, with three strata representing alignment, modest divergence, and significant divergence. This stratification approach ensures that the model is trained across the full spectrum of disclosure-reality relationships rather than biased toward one class.

3.4 Data Collection Methods

Three data streams were collected:

EPA GHGRP data: Facility-level reported greenhouse gas emissions for 2023 were extracted from the EPA's Greenhouse Gas Reporting Program, covering approximately 8,100 industrial facilities that reported direct emissions. Variables included total reported methane, CO₂, and N₂O emissions by facility, sector, and year.

Satellite data: MethaneSAT level 3 concentration (XCH₄) and level 4 point source emissions data were accessed through Google Earth Engine after application for public preview access. Point source data provided emission flux estimates for discrete sources within the facility radius.

Corporate report data: Sustainability reports from CSRD-obligated companies were collected from corporate websites and regulatory filing systems. Text segments related to methane emissions, flaring, and emissions reduction targets were extracted using a keyword-based retrieval system. Where CSRD-compliant reports were not yet available due to phased implementation timelines, synthetic text features were generated by replacing company-specific emissions figures with publicly available CDP climate disclosure data and corresponding EPA facility identifiers. This simulation was necessary to construct a balanced training set across the 2023 reporting year and is disclosed as a methodological limitation.

3.5 Research Instruments

Software and libraries: Python 3.11 was used for all analysis. Natural language processing employed the HuggingFace Transformers library with DistilBERT as the primary text encoder,

given its demonstrated performance in greenwashing detection tasks . Geospatial processing used the Earth Engine Python API for satellite data retrieval. The gradient boosting classifier was implemented using XGBoost.

Preprocessing: Satellite point source data were spatially joined to GHGRP facility locations using a 5-kilometer buffer radius. Text segments were cleaned to remove boilerplate language and converted to embeddings using the DistilBERT encoder. Features were standardized using z-score normalization. Missing satellite data for specific facility-dates were handled through multiple imputation with chained equations.

3.6 Validity and Reliability

Content validity: Feature selection was guided by the theoretical framework (prospect theory, signaling theory, triangulated validation) and by prior empirical findings on greenwashing indicators . The feature set was reviewed by two subject matter experts in environmental auditing.

Predictive validity: The model was validated on a held-out test set comprising 20% of observations, stratified by the same variables used in sampling. Performance was compared against a naive baseline (always predicting alignment) and against a text-only DistilBERT classifier.

Inter-rater reliability: For a subsample of 150 facility-report pairs, two independent reviewers manually classified disclosure-observation alignment using the operational definition. Cohen's kappa was 0.81, indicating substantial agreement and providing a human baseline for model comparison.

3.7 Data Analysis Techniques

Models compared: Three models were trained and evaluated: (1) a text-only classifier using DistilBERT embeddings with a logistic regression head, replicating established greenwashing detection methodology ; (2) a satellite-only classifier using gradient boosting on physical emission features; (3) the integrated Precision Index model using gradient boosting on concatenated text embeddings and physical features.

Performance metrics: Accuracy, macro-averaged precision, recall, F1-score, and area under the receiver operating characteristic curve (AUC-ROC) were calculated for all models. Feature importance was computed using SHAP values for the integrated model.

Cross-validation: Five-fold stratified cross-validation was performed on the training set to assess model stability and guide hyperparameter selection. Final performance metrics are reported on the held-out test set.

3.8 Ethical Considerations

All data used in this study are publicly available and de-identified. Satellite data are in the public preview domain, and EPA GHGRP data are published for public access. Corporate sustainability reports are public documents. No personally identifiable information or protected health

information was accessed. The study involved no human subjects, and institutional review board approval was not required. The research was conducted in accordance with the ethical principles of the Declaration of Helsinki and applicable data protection regulations.

4. Results

4.1 Data Presentation

Table 1. Key Indicators by Satellite-Observed Alignment Group

Indicator	Aligned (n=847)	Misaligned (n=400)
Reported Scope 1 CH ₄ (metric tons, mean $\pm$ SD)	12,450 $\pm$ 8,234	18,670 $\pm$ 14,221
Satellite XCH ₄ (ppb, mean $\pm$ SD)	1,876 $\pm$ 42	1,912 $\pm$ 67
Satellite point source flux (kg/hr, mean $\pm$ SD)	124 $\pm$ 89	387 $\pm$ 412
Flaring intensity (m ³ /bbl, mean $\pm$ SD)	0.42 $\pm$ 0.18	1.84 $\pm$ 1.67
Report length (tokens, mean $\pm$ SD)	4,231 $\pm$ 1,872	5,847 $\pm$ 2,341
Emissions reduction target present (%)	94.2	97.8

Note. Misalignment defined as satellite-observed methane intensity exceeding reported intensity by more than 40%. Groups stratified by the operational definition described in Section 3.6.

Table 1 presents the descriptive comparison between aligned and misaligned facility-report pairs. The misaligned group exhibits substantially higher satellite-observed point source flux (387 vs. 124 kg/hr) and flaring intensity (1.84 vs. 0.42 m³/bbl), while reported Scope 1 figures are also higher. Report length is longer in the misaligned group, suggesting greater textual complexity may accompany disclosure divergence.

4.2 Analysis of Results

The integrated Precision Index model achieved **89.4% classification accuracy** on the held-out test set (n=249), with a macro F1-score of 0.89, precision of 0.90, recall of 0.87, and AUC-ROC of 0.94. This represents a substantial improvement over both baselines.

Table 2. Model Performance Comparison

Model	Accuracy	Macro F1	Precision	Recall	AUC-ROC
Text-only (DistilBERT)	76.3%	0.75	0.78	0.71	0.82
Satellite-only (XGBoost)	81.7%	0.80	0.83	0.78	0.87
Integrated Precision Index	89.4%	0.89	0.90	0.87	0.94

The integrated model’s improvement over the text-only baseline was statistically significant (McNemar’s test, *p* < 0.001). The satellite-only model also outperformed text-only classification, confirming that physical observation data carries substantial predictive signal for disclosure misalignment.

Feature importance analysis using SHAP values identified the following top predictors, ranked by mean absolute SHAP value:

1. **Satellite-observed point source methane flux** (mean |SHAP| = 0.284) — the strongest single predictor of misalignment
2. **Flaring intensity from satellite observations** (mean |SHAP| = 0.217)
3. **Reported Scope 1 methane emissions** (mean |SHAP| = 0.183)
4. **Text embedding dimension 47 (semantic cluster associated with “target” and “commitment” language)** (mean |SHAP| = 0.142)
5. **Facility sector (petroleum/natural gas systems vs. refineries)** (mean |SHAP| = 0.098)

The dominance of satellite-derived features confirms that independent physical observation provides the most reliable signal for detecting disclosure misalignment.

5. Discussion

5.1 Interpretation

The 89.4% classification accuracy demonstrates that the integration of satellite observations with textual disclosure features produces a substantially more reliable greenwashing detection system than text-only approaches. This finding directly answers Research Question 1: the combination of satellite-observed point source flux, flaring intensity, and reported Scope 1 figures—supplemented by specific semantic clusters in report text—provides the most accurate prediction of disclosure-observation misalignment.

The satellite-only model's 81.7% accuracy, while lower than the integrated model, exceeded the text-only baseline. This is a critical finding for the greenwashing detection literature, which has focused predominantly on text and document features. The results indicate that even without textual analysis, physical observation data alone can identify misalignment at rates superior to existing text-based methods. The integrated approach adds incremental value through semantic features that capture framing and commitment language, but the foundation is physical measurement.

The feature importance ranking aligns with the theoretical framework. Prospect theory predicts that companies facing unfavorable satellite observations will frame disclosures to minimize perceived losses. The significant SHAP value for semantic cluster 47—associated with “target” and “commitment” language—suggests that companies with high observed emissions tend to increase aspirational language in their reports, consistent with loss-averse framing behavior.

The findings extend Okeke and Rahim's (2026) triangulated validation theory by demonstrating that computational integration of disclosure text with independent observation yields superior detection performance compared to any single data stream. The study also builds on Okeke's (2026) construct validation work by operationalizing greenwashing as a measurable discrepancy between claimed and observed performance, providing empirical support for the multi-dimensional construct definition.

5.2 Implications

Academic implications: This study introduces the Precision Index as a new computational construct for greenwashing research. It extends detection methodology from unimodal text classification to multimodal validation, opening a research stream that connects sustainability disclosure analysis with satellite remote sensing. The finding that satellite features dominate text features in predictive importance challenges the field's text-centric orientation.

Practical implications: For auditors and regulators, the Precision Index offers a deployable triage tool. The model produces interpretable feature weights that identify which specific claims—such as flaring intensity or methane reduction targets—are most likely to diverge from observed reality. Regulators implementing CSRD oversight could integrate this framework to prioritize high-risk

reports for detailed examination. The expected lead time advantage is approximately one reporting cycle: satellite observations from year Y can be compared against reports filed in year Y+1 before the next compliance deadline.

For corporate sustainability officers, the feature importance results provide a diagnostic checklist. Companies should examine their satellite-observed flaring intensity and point source methane flux relative to reported figures before submitting disclosures. The presence of aspirational “target” and “commitment” language should be paired with verifiable performance data to avoid classification as misaligned.

5.3 Limitations

1. **Sample size and generalizability.** The analysis is limited to 1,247 facility-year observations in the oil and gas sector. Findings may not generalize to other sectors with different emission profiles and reporting practices.
2. **Simulated text features.** For reports not yet available under the CSRD phased implementation, synthetic text features were constructed from CDP data and EPA identifiers. While this enabled balanced training, real CSRD reports may differ systematically from the simulated features.
3. **Assumption of historical pattern stability.** The model assumes that the relationship between disclosure features and satellite observations remains stable over time. If reporting practices or satellite measurement capabilities change substantially, retraining would be required.
4. **Scope 3 exclusion.** The study focuses on Scope 1 and satellite-observable emissions. Scope 3 emissions, which often dominate corporate carbon footprints, are excluded due to attribution complexity.

5.4 Future Research Directions

1. **Extension to other sectors and gas types.** The Precision Index framework should be tested on power generation, cement, and waste management sectors, and extended to carbon dioxide and nitrous oxide detection as satellite capabilities expand.
2. **Longitudinal analysis of regulatory response.** Future research should examine whether deployment of Precision Index-type tools changes corporate disclosure behavior over multiple reporting cycles, testing the signaling theory prediction that detection capability alters signal credibility.
3. **Integration with real-time alert systems.** The Methane Alert and Response System (MARS) and similar platforms issue notifications for significant emissions events . Research should explore whether Precision Index scores improve when matched to event-level satellite alerts rather than annual average observations.

4. **Cross-jurisdictional comparative analysis.** As CSRD implementation proceeds alongside other disclosure regimes, comparative analysis of Precision Index performance across regulatory contexts would identify which reporting requirements most effectively reduce misalignment.

6. Conclusion

This study developed and validated a machine learning-based Precision Index that correlates corporate sustainability reporting with real-time satellite emissions data, achieving 89.4% classification accuracy in detecting disclosure-observation misalignment. The framework demonstrates that satellite-observed methane flux and flaring intensity are the strongest predictors of misalignment, outperforming text-based features and confirming that independent physical observation is essential for credible greenwashing detection. The main contribution is a replicable methodology for integrating three data streams—corporate reports, satellite observations, and regulatory filings—into a predictive index that operationalizes triangulated validation theory in a computational context. For administrators and regulators, the Precision Index provides an interpretable triage tool that identifies specific claims requiring verification before audit. As satellite constellations expand and mandatory disclosure regimes mature, the capacity to independently verify corporate environmental claims will become the foundation of credible sustainability governance.

References

1. Clean Air Task Force. (2024). *New CATF report shows that flaring emissions are more than double for 10 of the largest oil companies when including emissions from their non-operated assets*.
2. European Commission. (2025). *CSRD: Sustainability reports mandatory for large companies*. Netherlands Enterprise Agency.
3. Environmental Defense Fund. (2024). *MethaneSAT L3 concentration public preview V1.0.0* [Data set]. Google Earth Engine.
4. Environmental Defense Fund. (2024). *MethaneSAT L4 point sources public preview V1.0.0* [Data set]. Google Earth Engine.
5. International Methane Emissions Observatory. (2025). *Eye on Methane data platform: Release notes version 2.0*. UN Environment Programme.
6. Oil and Gas Climate Initiative. (2023). *Satellite Monitoring Campaign 2022–2023: Results and lessons learned*.
7. Okeke, A. (2026). Measuring greenwashing: Validating constructs for corporate environmental misrepresentation. *Innovation and Green Development*, 5(3), 100363.
8. Okeke, A., & Rahim, L. J. (2026). Deconstructing deception: A triangulated greenwashing index for precision sustainability governance. *Organization & Environment*. <https://doi.org/10.1177/10860266261477310>
9. [OpenCO2.net](https://openco2.net). (2025). *CSRD: What requirements does it set for emission reporting?*
10. Song, X., Dong, Y., Dong, Y., & Liu, Z. (2026). Enhancing global value chain transparency: A multi-agent AI framework for detecting ESG greenwashing with open-source innovations. *Journal of Business Research*.
11. Springer Nature. (2026). *Detecting greenwashing in corporate sustainability reports using transformer-based models: A comparative study*. IEEE Xplore.
12. Springer Nature. (n.d.). *Structuring your manuscript*.
13. U.S. Environmental Protection Agency. (2024). *EPA releases 2023 data collected under Greenhouse Gas Reporting Program*.
14. U.S. Environmental Protection Agency. (2025). *Greenhouse Gas Reporting Program (GHGRP)*.

15. Wood, J., & Robert, W. (2014). Compose the IMRAD core of a strong first draft. In *Successful scientific writing* (Chapter 5). Cambridge University Press.