Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Measuring What Companies Say about Net Zero and Climate Transition: A Weighted-Lexicon Framework with ClimateBERT Corroboration

Srikanth Potharla

srikanthyadav444p@gmail.com

ICFAI Foundation for Higher Education

Research Article

Keywords: Net-Zero Disclosure, Climate-Transition Disclosure, Computer-Aided Text Analysis, Weighted Lexicon, ClimateBERT, India, BRSR

Posted Date: November 16th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11478503/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Measuring What Companies Say about Net Zero and Climate Transition: A Weighted-Lexicon Framework with ClimateBERT Corroboration

Srikanth Potharla
ICFAI Business School (IBS), Hyderabad
A constituent of ICFAI Foundation for Higher Education (IFHE). India,
Email: srikanthyadav444p@gmail.com
Institutional Email: psk@ibsindia.org
https://orcid.org/0000-0003-3176-9675
Mobile No. +91- 9703664124

Measuring What Companies Say about Net Zero and Climate Transition: A Weighted-Lexicon Framework with ClimateBERT Corroboration

Abstract

What companies disclose about their net-zero and climate transition is not the same as what they achieve, yet existing text-analytic approaches face a trade-off between scalable but undifferentiated keyword counts and contextually sensitive but costly manual coding. This paper develops Net-Zero/Climate-Transition Disclosure (NZCTD), a transparent, auditable measure of the volume and specificity of net-zero and climate-transition language in corporate annual reports — not a measure of actual climate performance, emissions reduction, implementation, compliance, or sincerity. NZCTD is built from an eighteen-category, 273-pattern lexicon with evidentiary weights from 0 to 4, applied at the paragraph level with contextual safeguards, and corroborated by an independent, model-based ClimateBERT layer applied only to qualifying paragraphs. Applied to 468 Indian listed companies and 4,048 company-year observations spanning FY2013–FY2025, the measure has a mean score of 319.58 (median 140.0) with pronounced right-skewness and a mean rising from approximately 32.2 in FY2013 to approximately 909.7 in FY2025, an independently recomputed 28.25-fold increase. Weighted and unweighted scores correlate highly (Pearson r = 0.9855; Spearman rho = 0.9919), and principal component analysis indicates a coherent common structure alongside meaningful multidimensionality (PC1 explains 46.85% of variance; r = 0.8655 with the total score), while ClimateBERT confirmation ratios reach approximately 0.684 in later years. NZCTD offers a scalable, transparent framework for measuring climate-transition disclosure that preserves contextual insights and the disclosure-versus-performance distinction throughout.

Keywords: Net-Zero Disclosure; Climate-Transition Disclosure; Computer-Aided Text Analysis; Weighted Lexicon; ClimateBERT; India; BRSR

1. Introduction

Corporate annual reports describe net-zero ambitions and climate-transition activities through statements that vary considerably in scope and specificity. General aspirations appear alongside emissions inventories, quantified targets, transition plans, and reported progress, creating a central research problem: how can this heterogeneous language be measured consistently without treating disclosure as evidence that the underlying commitments have been implemented? The problem concerns both the comparability of textual measures and the conclusions drawn from them. Research has traced the institutionalization of corporate carbon reporting, its connections with sustainability initiatives, and the subsequent differentiation of climate-related disclosure frameworks and measurement approaches (Kolk et al., 2008; Ngo et al., 2023; Tumewang et al., 2025; Agustini et al., 2026). As this literature expands, researchers and users of corporate reports need measures with clearly specified informational boundaries. India provides a relevant setting because reporting expectations have evolved through the introduction of Business Responsibility and Sustainability Reporting (BRSR) requirements (Khalid & Rawat, 2025; Kajal & Bansal, 2025). These developments sharpen the need to measure transition-related language while distinguishing changes in reporting from changes in environmental outcomes.

The resulting research gap concerns integrating a bounded-disclosure construct, differentiated textual evidence, contextual interpretation, and transparent aggregation into a scalable measurement framework. Three connected theoretical perspectives explain why this integration matters. Disclosure–performance decoupling establishes that environmental communication and outcomes can diverge; organizational fragmentation and implementation uncertainty also caution against interpreting every divergence as deliberate deception (Walker & Wan, 2012; Kudłak, 2025). Signaling and credibility perspectives explain why readers examine disclosed commitments for information about transition intentions, while research on emissions, transition pathways, and greenwashing shows why credibility requires evidence beyond the presence of detailed language (Ding et al., 2023; Chan et al., 2024; David et al., 2025). Research on governance-related determinants further motivates the examination of disclosure as an organizational reporting output with its own antecedents (Gebhardt et al., 2024; Wahyuningrum et al., 2024). Together, these perspectives imply that measurement should preserve distinctions between disclosed content, its contextual relevance, and independently observed performance. Methodological development has progressed from manual content analysis and dictionary-based approaches to transformer-based classification, offering different capacities for coding reliability, reproducibility, contextual interpretation, and scalability (Milne & Adler, 1999; Loughran & McDonald, 2016; Garrido-Merchán et al., 2026). Related advances in measuring firm-level climate exposure demonstrate the value of corporate narratives but address constructs distinct from transition-disclosure content (Sautner et al., 2023; Dell'Atti et al., 2026). The literature reviewed, therefore, motivates a framework that combines explicit content categories and contextual safeguards while keeping supplementary classification separate from disclosure scoring. This study addresses that need by developing Net-Zero/Climate-Transition Disclosure (NZCTD), a measure of the weighted volume of qualifying transition-related language, whose categories differentiate the types of disclosed

content and their specificity. Its intended conceptual contribution is to translate the disclosure–performance distinction into explicit measurement rules.

The study applies NZCTD to an archive-based sample of 468 listed Indian companies and 4,048 company-year observations spanning fiscal-year starts 2013–2025, with the final period corresponding to FY2025–26. Annual-report PDFs undergo native-text extraction or optical character recognition, followed by cleaning and paragraph segmentation; the resulting measurement depends on the completeness and fidelity of recovered text (Hickman et al., 2022; Gutiérrez Ponce et al., 2023). The instrument contains 273 matching patterns across eighteen categories, with fixed, judgment-based weights ranging from zero to four. Paragraph-level safeguards govern weak-context expressions, ambiguous acronyms, and bare numerical matches, while external-actor tagging supplies an annotation without excluding otherwise qualifying matches. This specification makes category definitions, weighting choices, and contextual rules explicit, drawing on work on domain-sensitive dictionaries, instrument development, content analysis limitations, and methodological transparency (Loughran & McDonald, 2011; Short et al., 2010; Dumay & Cai, 2015; Reid et al., 2023). Qualifying category counts are weighted and aggregated into an unnormalized company-year score. A pretrained ClimateBERT detector separately assesses lexicon-selected paragraphs and generates a company-year climate-relevance ratio, following the broader use of complementary approaches in textual accounting and finance research (Fisher et al., 2016; Huang et al., 2023; Bochkay et al., 2023). Model classifications do not alter NZCTD. Evaluation comprises arithmetic reconciliation, weighted–unweighted comparisons, inter-category correlations, and principal component analysis.

The application reveals substantial heterogeneity in qualifying disclosure: the mean NZCTD score is approximately 319.6, compared with a median of 140.0. Annual average scores are higher in later samples, although changing company coverage and text volume prevent interpreting this pattern as equivalent growth within firms or as a regulatory effect. Category prevalence also differs from weighted contribution, demonstrating the additional information available from category profiles. Weighted and unweighted scores yield closely corresponding rankings, with a Spearman correlation of 0.9919; however, this comparison depends on both relative weights and the inclusion of zero-weight categories. ClimateBERT provides supplementary evidence about the climate relevance of selected text, conditional on lexicon selection. These findings support a conceptual contribution through clearer construct boundaries, a methodological contribution through explicit scoring and separate validation, and an empirical contribution through longitudinal characterization of an Indian reporting corpus. The framework also offers a basis for disclosure research and preliminary screening. Its classification performance remains to be established against independently coded text, which requires a documented annotation protocol and an appropriate reliability assessment (O'Connor & Joffe, 2020; Cole, 2024). Accordingly, the present evidence supports internal consistency and descriptive interpretation while leaving external validation and text-length sensitivity as priorities for further development.

Section 2 develops the conceptual foundations, reviews the literature, and specifies the research gap. Section 3 describes the methodology. Section 4 presents the results, diagnostic evidence,

and discussion. Section 5 concludes with implications, limitations, and directions for future research.

2. Theoretical Background and Literature Review

2.1 Conceptual Background and Literature Review

The literature on corporate climate disclosure provides both conceptual foundations and methodological approaches for measuring net-zero and climate-transition disclosure. Early research established carbon and climate reporting as an institutionalized corporate practice rather than a neutral technical exercise (Kolk, Levy, & Pinkse, 2008). Subsequent reviews have situated frameworks such as the Task Force on Climate-related Financial Disclosures (TCFD) within the broader sustainability-reporting literature (Tumewang, Ntim, & Haque, 2025). Climate-risk disclosure has also been examined in relation to global sustainability initiatives (Ngo, Le, Ullah, & Trinh, 2023), underscoring its treatment as a distinct research focus within sustainability communication. A recent bibliometric mapping documents a diverse, rapidly growing climate-related disclosure literature with multiple measurement approaches (Agustini, Hartono Mustakini, Arifa, & Barokah, 2026). This diversity highlights the importance of specifying the scope and interpretation of any proposed disclosure measure.

NZCTD focuses on the volume and specificity of net-zero and climate-transition language in corporate annual reports. Its scope is narrower than broad ESG or sustainability disclosure, which encompasses environmental, social, and governance topics. It also differs in measurement purpose from TCFD-aligned climate-related financial disclosure, which addresses governance, strategy, risk management, and metrics and targets relating to climate-related financial risks and opportunities (Tumewang, Ntim, & Haque, 2025). These constructs overlap: emissions inventories, transition plans, governance arrangements, and targets may contribute to both NZCTD and TCFD-aligned disclosure. However, NZCTD measures qualifying transition-related language rather than comprehensiveness of alignment with the TCFD recommendations. The distinction, therefore, concerns scope and measurement purpose, rather than mutually exclusive disclosure content.

Textual analysis in accounting and finance has historically relied extensively on dictionary- and keyword-based methods. Loughran and McDonald (2016) document their transparency and reproducibility, as well as the challenges of interpreting context-dependent language. Undifferentiated keyword counts may give similar treatment to specific commitments and vague or boilerplate statements. Manual content analysis also requires explicit attention to

coding reliability. Milne and Adler (1999) examine intercoder reliability in annual-report social and environmental disclosure analysis, showing that reliability assessment and coder preparation are particularly relevant to detailed subcategory analysis. A related line of research develops firm-level climate measures from other textual sources. Sautner, van Lent, Vilkov, and Zhang (2023), for example, construct a firm-level climate change exposure measure from earnings call transcripts, illustrating the value of scalable textual measurement while examining a construct distinct from the specificity of companies' net-zero and transition disclosures. Dell'Atti, Foglia, and Onorato (2026) extend this literature by constructing a climate-risk exposure index from corporate narratives and examining its relationship with firm valuation and disclosure tone. Their study further illustrates why the interpretation of a textual measure depends on the construct and research question it is designed to address.

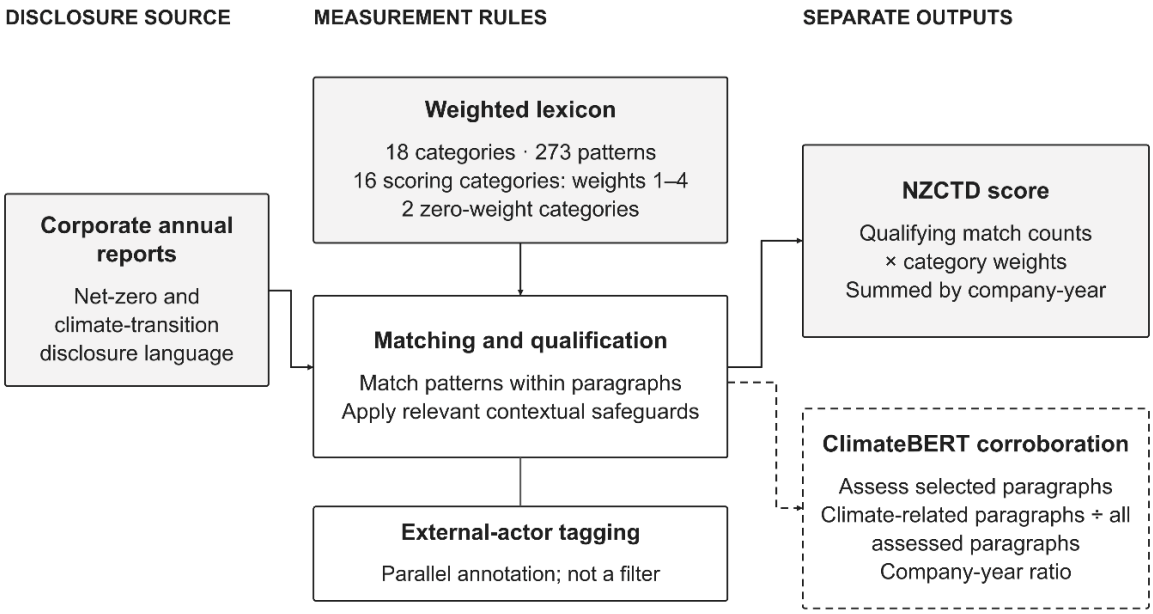


Figure 1. NZCTD measurement architecture and interpretive boundaries

Note: Figure 1 shows how NZCTD aggregates qualifying category-level match counts using fixed weights. Sixteen categories have positive weights; weak context and the external-actor flag have a weight of zero. ClimateBERT assesses only lexicon-qualifying paragraphs and produces a separate company-year corroboration ratio; its classifications do not alter the NZCTD score. External-actor tagging is an annotation, not a qualification filter. Neither output verifies implementation, emissions reductions, commitment credibility, or regulatory compliance. The dashed route denotes supplementary model-based corroboration.

Transformer-based language models, including the ClimateBERT family, offer another approach by incorporating contextual information into text classification. Garrido-Merchán, González-Barthe, and Coronado-Vaca (2026) fine-tune BERT and ClimateBERT on

ClimaText for climate-related text detection, providing an example of task-specific adaptation for disclosure analysis. This approach differs from the present study's use of a pretrained climate detector as supplementary validation of paragraphs selected by a separately specified lexicon. In the applications identified during this review, transformer models generally serve as primary classifiers rather than as supplementary checks on a lexicon-based measure. The literature reviewed here suggests an opportunity to combine a weighted, category-structured lexicon, explicit safeguards against generic language, ambiguous acronyms, and unrelated numeric text, and model-based validation within a multi-year annual-report application.

The interpretation of the resulting measure depends on maintaining the distinction between disclosure and actual climate performance. The disclosure–performance decoupling literature provides the conceptual basis for this distinction. Walker and Wan (2012) show that environmental communication and underlying environmental performance are separable constructs with potentially different financial implications. Subsequent studies examine carbon emissions alongside TCFD-aligned disclosure (Ding, Liu, & Chang, 2023), investigate whether ESG disclosure functions as credible signaling or greenwashing in BRICS equity markets (David, Wang, Luo, & Cisse, 2025), and assess the credibility of corporate net-zero transition pathways as a distinct research problem (Chan, Cheung, & Shen, 2024). The study by Kudłak (2025) offers an organizational explanation for loose coupling between environmental commitments and outcomes, highlighting how fragmented structures and uncertain causal relationships can impede implementation. This perspective cautions against treating divergence between commitments and outcomes as sufficient evidence of deliberate deception. Research on governance-related determinants of climate reporting (Gebhardt, Schneider, Siedler, Ottenstein, & Zülch, 2024) and carbon-emission disclosure (Wahyuningrum, Ihlashul'amal, Utami, Djajadikerta, & Sriningsih, 2024) further positions disclosure as an object of analysis in its own right.

Accordingly, a higher NZCTD score indicates greater qualifying disclosure under the specified category weights; it does not establish that transition actions have been implemented, that emissions have declined, or that commitments are credible or fulfilled. Assessing actual climate performance requires separate evidence on emissions levels, intensity, reductions, or progress against defined targets. ClimateBERT provides supplementary validation of the climate relevance of qualifying paragraphs and does not enter the NZCTD score or verify the underlying actions described. Figure 1 illustrates this measurement architecture and the boundary between disclosed language and actual climate performance.

2.2 Research Gap and Contribution

Four gaps follow from this review. First, the literature does not appear to provide a climate-transition disclosure measure that combines a fully weighted, category-structured lexicon with explicit, paragraph-level contextual safeguards against the specific false-positive risks posed by such a lexicon, while remaining fully auditable to the exact source passage. Second, published applications of ClimateBERT-type models typically use the model as the primary classifier over narrow, pre-selected text, rather than as an independent corroborative layer over a separately specified, full-corpus lexicon score; the architecture developed here inverts that relationship. Third, net-zero/climate-transition disclosure is frequently conflated with broad ESG or sustainability disclosure, motivating a sharper operational boundary than is typically drawn in practice. Fourth, we identify a gap in the availability of a comparably large, fully auditable NZCTD measure for an emerging-market annual-report corpus spanning a period of substantial disclosure-regime change — in this study's setting, the years surrounding India's Business Responsibility and Sustainability Reporting mandate.

The study's contribution follows directly from these four gaps. Methodologically, it develops a transparent, editable, and auditable architecture — a weighted, category-structured lexicon; three explicit contextual safeguards and one non-gating attribution tag; and an independent language-model corroboration layer — that other researchers can adopt or adapt to other jurisdictions, languages, or disclosure domains. Empirically, it produces and descriptively characterizes a large, multi-year, India-specific NZCTD measure covering 468 listed firms across thirteen fiscal years. Conceptually, it draws a sharper operational line between net-zero/climate-transition disclosure and the broader ESG/sustainability disclosure literature with which it is frequently blended. Practically, it offers a reusable lexicon and a documented, resumable pipeline that can serve as a low-cost screening infrastructure for researchers, analysts, or regulators.

3. Methodology

3.1 Measurement Framework and Construct Definition

NZCTD measures qualifying net-zero and climate-transition language in corporate annual reports at the company-year level. Following the construct boundary established in Section 2, it captures disclosed content rather than verified implementation, emissions performance, or commitment credibility. The instrument combines a category-weighted lexicon, paragraph-level contextual safeguards, and a supplementary ClimateBERT assessment. Category weights

differentiate types of disclosure language, consistent with explicit instrument design in computer-aided text analysis (Short, Broberg, Cogliser, & Brigham, 2010). They are judgment-based, not estimated from environmental outcomes or externally validated. ClimateBERT assesses the climate relevance of selected paragraphs but does not enter the NZCTD score.

3.2 Data and Sample

The analytical sample comprises 468 listed Indian companies and 4,048 company-year observations across thirteen fiscal-year starts, 2013-14 to 2025-26. The panel is unbalanced, averaging 8.65 observations per company (median: 11). Coverage comprises 66 companies with 1–3 observations (14.1%), 49 with 4–6 (10.5%), 54 with 7–9 (11.5%), and 299 with 10–11 (63.9%). No duplicate company–fiscal-year records were identified. Reports were drawn from the project archive available at the time of extraction, yielding a convenience sample rather than a probability-based sample. All analyses use the final scored dataset consisting of 4,048 company-year observations.

3.3 Text Extraction and Preprocessing

Processing begins with a fixed, read-only PDF archive. A persistent manifest records each file as pending, completed, or failed, supporting resumable processing. Pages undergo native-text extraction or, when extractable text falls below a minimum-character threshold, optical character recognition (OCR) using Tesseract at 300 DPI. Cleaning removes headers and footers, rejoins hyphenated words, and normalizes whitespace. Text is segmented at blank-line boundaries into paragraphs, the unit used for contextual safeguards, following text-preprocessing guidance for organizational research (Hickman, Thapa, Tay, Cao, & Srinivasan, 2022).

A PDF that encounters an extraction error is marked as failed and contributes no analytical observations; an extraction failure is not assigned a zero-disclosure score. Native extraction and OCR can introduce character errors and imperfect reading orders. Developmental spot-checking was undertaken, but no systematically sampled extraction-quality audit is documented. These risks require consideration when interpreting scores, particularly where recovered text may incompletely represent a report (Gutiérrez Ponce, Chamizo González, & Al-Mohareb, 2023). Figure 2 summarizes extraction, scoring, and the separate corroboration branch.

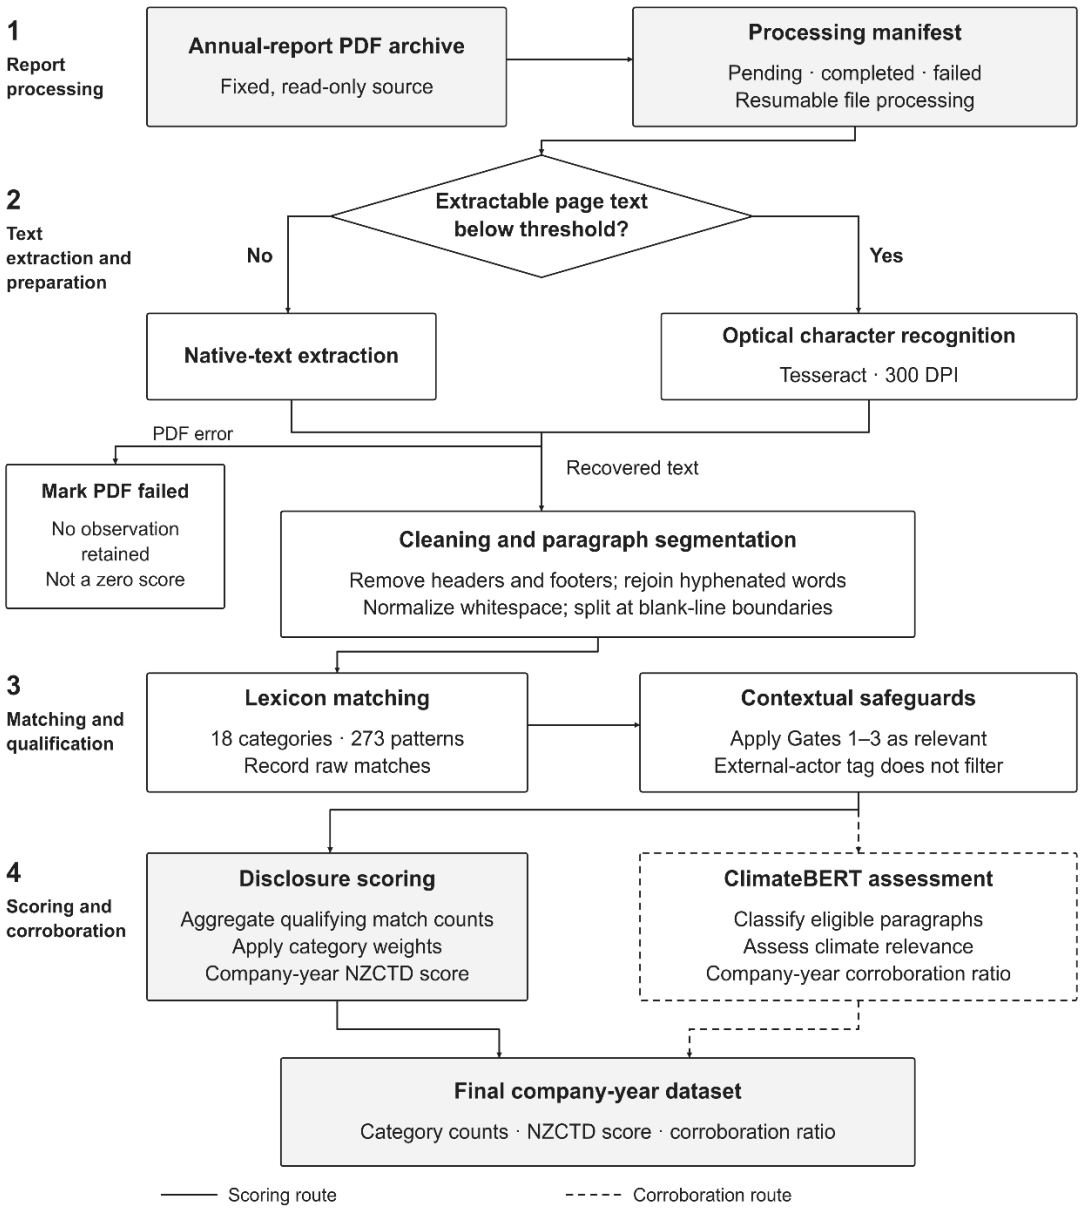


Figure 2. Annual-report processing and construction of the company-year NZCTD dataset

Note: The extraction route is determined at the page level using the specified minimum-character threshold. OCR denotes optical character recognition; DPI denotes dots per inch. A PDF extraction error results in a failed processing status and no retained analytical observation, rather than a zero-disclosure score. Relevant contextual safeguards determine match qualification; external-actor tagging does not filter matches. Qualifying matches are aggregated for scoring, while eligible paragraphs are assessed separately with ClimateBERT. Both outputs enter the company-year dataset. The dashed route denotes corroboration and does not contribute to the score calculation. Detailed settings and qualification rules are described in Sections 3.3–3.6.

3.4 Lexicon Development and Architecture

The lexicon is maintained in a version-controlled JSON file (v1.3) containing 273 case-insensitive regular expression patterns across 18 categories. Each category has a fixed weight applied to qualifying matches. This structured specification addresses the contextual limitations of generic dictionaries (Loughran & McDonald, 2011) and makes category assignments and weighting choices inspectable, thereby supporting methodological transparency in dictionary-based text analysis (Reid, McKenny, & Short, 2023). Categories distinguish disclosure dimensions without implying mutually exclusive concepts or textual expressions.

Weights range from 0 to 4. Weight 4 covers core net-zero and emissions-inventory language; weight 3 covers emissions scopes, targets, transition planning, and reported progress. Weight 2 covers temporal references, quantification, offsets and removals, governance, assurance, and finance/materiality. Weight 1 covers energy, operational, and value-chain actions, as well as references to reporting frameworks. Weak context and the external-actor flag receive weight 0, serving qualification and attribution functions without directly increasing the score. The weights express intended differentiation of disclosure content, not a validated ordering of disclosure quality. Table 1 provides category definitions, weights, pattern counts, and illustrative terms.

India-specific coverage includes policy, institutional, reporting, and financial terminology, such as the PAT Scheme, BRSR, and sovereign green bonds. Numerical patterns accommodate lakh- and crore-tonne expressions, as well as tCO2e and MtCO2e. The available changelog documents v1.2 and v1.3 refinements to matching patterns, generic expressions, India-specific terminology, and ambiguous acronyms. This contextual adaptation supports application to the study setting but does not itself establish classification accuracy. Version history preceding v1.2 is not documented in the available archive.

[Table 1 about here]

3.5 Contextual Matching, Safeguards, and Scoring

Each paragraph is checked against all lexicon patterns. Raw matches are recorded during processing with their category, matched text, and weight; a separate pass determines qualification using paragraph-level safeguards. Qualifying counts and weighted contributions remain distinct: qualifying zero-weight matches do not increase the score. Explicit contextual

rules address risks associated with unguarded content analysis (Dumay & Cai, 2015) but do not guarantee semantic accuracy.

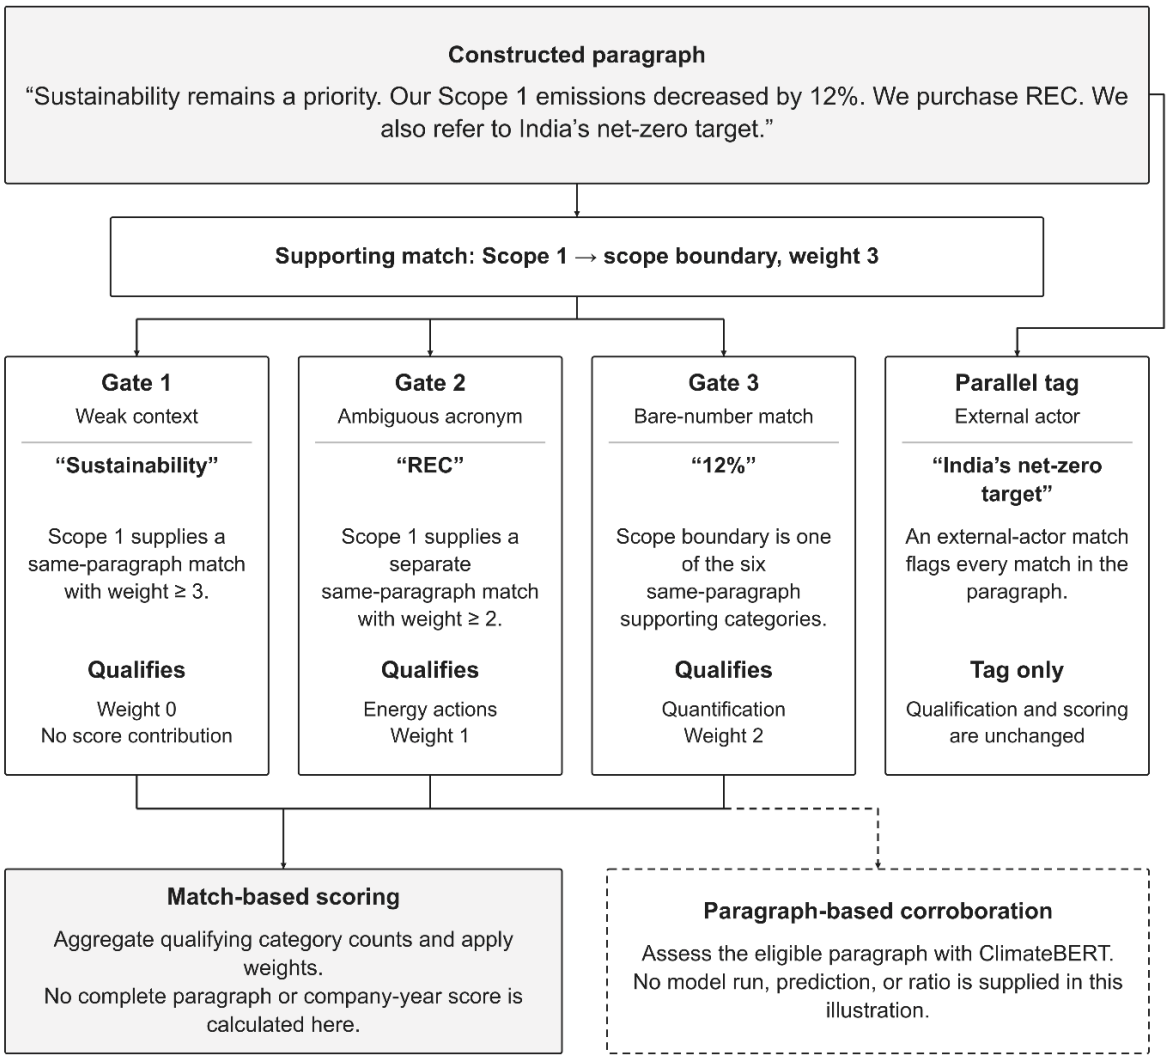


Figure 3. Illustrative application of contextual safeguards and separate ClimateBERT corroboration

Note: The paragraph is constructed for explanation and is not a quotation from a sampled report or a verified pipeline output. The illustration assumes that the displayed expressions are detected and assigned to the indicated categories, including a bare-number quantification match and an external-actor pattern. Scope 1 provides sufficient in-paragraph support for the applications displayed in Gates 1–3. The figure illustrates one sufficient condition for Gate 1, not its complete rule. Only selected matches and their weights are shown; other possible matches, including net-zero and reported-progress language, are not enumerated. No complete paragraph or company-year score is calculated. External-actor tagging applies to all matches in the paragraph without changing qualification or scoring. No ClimateBERT run, prediction, confidence score, or corroboration ratio is supplied. Complete rules are provided in Section 3.5 and the contextual safeguards table.

As shown in Figure 3, Gate 1 applies to weak-context matches: qualification requires a match with a weight ≥ 3 or matches from at least two other distinct substantive categories in the same paragraph. Gate 2 applies to PPA, REC, I-REC, CSO, CDR, VCS, and TPT, etc., an acronym match requires a separate match with weight ≥ 2 in the paragraph. Gate 3 applies to bare-number quantification matches: qualification requires a same-paragraph match in core net-zero, emissions inventory, scope boundary, target commitment, energy actions, or performance progress. Co-occurrence reduces ambiguity but does not conclusively link an expression or number to a climate claim. External-actor tagging operates separately. Detection of an external-actor pattern flags every match in that paragraph without changing qualification or scoring eligibility. The tag does not identify ownership of each individual claim; externally attributed language can therefore contribute when it otherwise qualifies. Table 2 summarizes the safeguards and tagging rules.

[Table 2 about here]

Qualifying matches are aggregated by category and company year. The score is calculated as:

$$\mathrm{NZCTD}_{it} = \sum_{c=1}^{18} w_c \, Q_{itc}$$

Here, Q_{itc} is the qualifying-match count for company (i), year (t), and category (c), and w_c is its assigned weight. Counts represent textual matches, not patterns or paragraphs. Company-years without qualifying positive-weight matches receive zero and remain in the sample, including those with only zero-weight matches.

Internal consistency is assessed by recomputing weighted totals from category counts. Diagnostic procedures comprise Pearson and Spearman correlations, an unweighted sum across all 18 categories, pairwise category correlations, and principal component analysis of standardized category counts. The unweighted comparison changes both the relative weights and the contributions of the two zero-weight categories; it does not isolate the effect of weighting alone. PCA likewise includes all eighteen categories. Section 4 reports these diagnostics and the arithmetic reconciliation, distinguishing internal consistency from construct validation. Match-level safeguard details were not retained in the delivered records, limiting retrospective examination of individual qualification decisions. Recalculating scores from the category counts confirms that they agree with the stored scores.

3.6 ClimateBERT-Assisted Corroboration

ClimateBERT supplements lexicon matching with contextual classification, following the use of complementary rule-based and model-based approaches in textual accounting and finance research (Fisher, Garnsey, & Hughes, 2016; Huang, Wang, & Yang, 2023). The pretrained model climatebert/distilroberta-base-climate-detector assesses only lexicon-qualifying paragraphs, rather than the entire corpus, to maintain computational feasibility on CPU hardware. Selected text is truncated to 2,000 characters before classification. The model returns a predicted label and confidence score; the climate-relevance label supplies a binary indicator. Neither the label nor the confidence score changes NZCTD scoring.

For each company-year, the corroboration ratio divides the number of assessed paragraphs classified as climate-related by the total number of eligible paragraphs assessed. The dataset stores this measure as bert_confirmed_ratio. A ratio of zero means that none of the assessed paragraphs were classified as climate-related; a ratio of one means that all were. If no paragraphs are eligible, the ratio is undefined and recorded as unavailable, not zero. Descriptive summaries exclude unavailable ratios but retain valid zeros. Overall and annual averages are calculated from the available company-year ratios, giving each ratio equal weight regardless of how many paragraphs were assessed.

Corroboration is conditional on lexicon selection. A different classification mechanism does not make the assessment independent of paragraph selection: ClimateBERT cannot identify relevant passages that the lexicon omits or achieve full-corpus recall. A higher ratio indicates greater model-classified climate relevance among assessed paragraphs, not necessarily greater specificity, credibility, or actual climate performance. The ratio is not interpreted as a measure of accuracy, precision, recall, or agreement with human coding. Section 4 distinguishes this supplementary evidence from arithmetic consistency and human or external-criterion validation.

4. Results, Validation, Robustness, and Discussion

4.1 Descriptive Characteristics and Diagnostic Evidence

The analytical sample comprises 4,048 company-year observations from 468 listed Indian companies, as described in Section 3.2. Table 3 summarizes the distribution of NZCTD scores. The mean is 319.6, compared with a median of 140.0 and a standard deviation of 493.7. Scores range from 0 to 8,674, with a skewness of 4.87. The 25th and 75th percentiles are

approximately 53.8 and 395.0, which locate the middle half of the observations within a broad interval. The mean–median difference and positive skewness indicate that relatively high scores occur alongside substantially lower values across much of the sample. These statistics describe variation in weighted qualifying match counts and do not establish differences in verified climate performance.

The ClimateBERT corroboration ratio is available for 3,969 company years. Its mean is 0.566, median 0.609, and standard deviation 0.219. These statistics summarize the company-year proportions of assessed paragraphs classified as climate-related, calculated as described in Section 3.6. The mean gives equal weight to each available company-year ratio; it is not the proportion of all selected paragraphs across the corpus that ClimateBERT classifies as climate-related.

[Table 3 about here]

Seventy-nine company-years (1.95%) have zero NZCTD scores. These observations correspond exactly to the 79 unavailable ClimateBERT ratios, with no other unavailable ratios reported. This issue is observed in the dataset rather than required by the scoring formula, because qualifying matches confined to zero-weight categories also produce zero scores. Of the zero-score observations, 55 occur in fiscal-year-start 2024, representing 69.6% of all zero scores and 12.0% of that year's 458 observations. These percentages distinguish the concentration of zero scores in one year from their incidence within that year's sample. The available evidence does not establish whether this concentration reflects report content, incomplete text recovery, or other processing characteristics.

Table 4 shows higher annual mean NZCTD scores in later years: 32.2 at the fiscal year start in 2013, 146.3 in 2019, and 909.7 in 2025. The endpoint ratio is approximately 28.3 but compares five observations in 2013 with 203 in 2025. Median scores also rise, from 26.0 in 2013 to 683.0 in 2025, although these medians likewise refer to different annual samples. Changes in company coverage and differences in report length or extracted text volume may therefore influence the pattern; it does not establish equivalent growth across companies. Annual mean ClimateBERT ratios follow a different, non-monotonic pattern. The mean is 0.614 in 2013, ranges from 0.441 to 0.581 during 2014–2021 and lies between 0.671 and 0.684 during 2022–2025. These later averages indicate higher company-year proportions of assessed text classified as climate-related, conditional on the lexicon selection and the availability of ratios.

[Table 4 about here]

Table 5 distinguishes prevalence, the percentage of company-years containing a qualifying category match, from contribution, the category's share of the aggregate weighted score. Contribution percentages use the sum of scores across company-years as their denominator, rather than averaging individual company-year contribution percentages. Categories can co-occur, so prevalence percentages are not additive. Contribution depends on both qualifying-match frequency and assigned weight. Assurance and validation are the most prevalent positive-weight category, appearing in 97.16% of company-years but contributing 9.51% of the aggregate score. Emissions inventory appears in 75.94% and contributes the largest share of the score, 28.03%. Quantification contributes 13.75%, while core net-zero appears in 38.93% and contributes 5.89%. Thus, widespread occurrence does not necessarily imply a similarly large contribution. Assurance-related matches, in particular, do not establish that climate disclosures were independently assured. The distinction between prevalence and score contribution is particularly clear for the two zero-weight categories. Weak-context matches occur in 77.50% of company-years and external-actor flags in 2.89%. These categories identify generic sustainability language and references to external actors' targets, respectively. They support contextual assessment and attribution but do not directly increase the NZCTD score because both carry zero weight.

[Table 5 about here]

Table 6 reports score comparisons, inter-category correlations, and principal component analysis. Across 153 category pairs, the mean Pearson correlation is 0.401; the highest is between emissions inventory and scope boundary (0.882), and the lowest is between quantification and the external-actor flag (0.031). PCA uses all eighteen standardized category counts, including the two zero-weight categories. PC1 explains 46.9% of total variance, PC2 explains 8.1%, and the first three components together explain 60.6%. PC1 correlates with NZCTD ($r = 0.8655$). These results describe relationships among the counts and their aggregate summaries.

[Table 6 about here]

4.2 Validation and Robustness

The reported evidence addresses internal consistency, supplementary classification, and sensitivity to score construction. Recomputing NZCTD from category counts and weights produces no discrepancies across all 4,048 observations, supporting arithmetic consistency with the stored score. The pipeline also incorporates duplicate-record, negative-count, score-

equality, and ratio-bounds checks. Their inclusion documents specified integrity controls, while the reported recomputation provides an explicit numerical result.

ClimateBERT supplies a second classification mechanism, but its assessment remains conditional on the lexicon-selected paragraphs. Developing one would require an explicit annotation protocol, independent coding, and documented assessment and reconciliation of coding disagreements to strengthen transparency and consistency (Cole, 2024). Intercoder reliability and the instrument's accuracy against the resulting benchmark are distinct evaluations (O'Connor & Joffe, 2020). The ratio, therefore, provides supplementary evidence about selected paragraphs, with the interpretive limits specified in Section 3.6.

The score comparison in Table 6 yields Pearson $r = 0.9855$ and Spearman $\rho = 0.9919$ between NZCTD and the unweighted sum of all eighteen category counts. These associations indicate strong linear correspondence and broadly similar company-year rankings under the two specified constructions. However, the comparison simultaneously changes relative weights and includes weak-context and external-actor counts that contribute zero to NZCTD. It cannot isolate the influence of differential weighting or establish that category inclusion is inconsequential. The result supports a conclusion of limited sensitivity rather than the optimality of the chosen weights.

4.3 Discussion

The temporal pattern indicates that the observed corpus contains higher average volumes of weighted qualifying language in later years. This finding is relevant to research on changing corporate reporting practices, but its interpretation depends on the archive's composition and textual coverage. Evolving reporting expectations provide a possible institutional explanation that the descriptive design cannot identify separately from other influences. The immediate contribution is therefore a documented pattern that subsequent studies can examine using comparable-company samples, text-length adjustments, and designs appropriate for assessing regulatory effects. Those extensions would test explanations rather than infer them from the annual averages alone.

The category findings show why NZCTD benefits from interpretation at both aggregate and category levels. Prevalence identifies how widely a type of language appears, while contribution reflects its weighted volume across the corpus. Their divergence means that the composite score cannot by itself describe the composition of a company's disclosure. Category profiles make that composition more accessible, provided the categories are understood as

textual classifications rather than verified activities. This approach is consistent with transparent, theoretically grounded instrument development in computer-aided text analysis (Short, Broberg, Cogliser, & Brigham, 2010). The dimensional diagnostics complement these profiles by describing shared variation without establishing that all categories are interchangeable indicators of one underlying trait.

The similarity between weighted and unweighted rankings offers a useful reference point for evaluating the instrument. It indicates that the two constructions summarize much of the same variation within this sample. However, their comparison cannot determine whether weights, category inclusion, or their combination accounts for ranking differences. Nor do similar rankings imply similar category contributions or theoretical equivalence between the constructions. Researchers adapting NZCTD should therefore make their category and weight choices explicit and evaluate their consequences for the intended application. The present evidence supports transparency about these choices rather than a claim that one specification is universally preferable.

ClimateBERT adds a complementary perspective on the climate relevance of selected passages, consistent with combining dictionary-based and machine-learning approaches in accounting and finance research (Bochkay, Brown, Leone, & Tucker, 2023). Its role is most useful when the distinction between the two outputs remains explicit: NZCTD quantifies qualifying language, while corroboration summarizes model classifications of selected paragraphs. The framework's contribution lies in making disclosure measurement inspectable and clarifying its evidentiary boundaries.

5. Conclusion and Implications

This study develops NZCTD as a transparent framework for measuring net-zero and climate-transition disclosure language in corporate annual reports. Its eighteen-category weighted lexicon combines paragraph-level contextual safeguards with supplementary ClimateBERT assessment, keeping disclosure scoring separate from model-based corroboration. An application to 468 Indian-listed companies and 4,048 company-year observations reveals substantial variation in qualifying disclosure and higher annual average scores in later samples. Category-level results show that the widespread occurrence of a type of language need not translate into a proportionately large contribution to the aggregate score. Weighted and unweighted constructions produce similar rankings, although their comparison changes both category weights and inclusion. These findings establish descriptive patterns within the

observed corpus; changes in sample composition and text volume limit the interpretation of temporal differences. The study's contribution is an explicit, inspectable measurement architecture that organizes heterogeneous disclosure language while preserving the distinction between reported content and evidence of implemented commitments or improved climate performance.

Conceptually, NZCTD specifies a bounded disclosure construct within broader sustainability reporting and distinguishes it from independently measured environmental outcomes. Methodologically, explicit categories, weights, and qualification rules allow researchers to inspect measurement choices and adapt them to different reporting settings. Keeping ClimateBERT corroboration separate also makes each approach's contribution visible. In practice, the framework could support comparative disclosure research, preliminary screening, and the selection of passages for closer examination by analysts or regulators. Category profiles can help users identify which language forms contribute to a score, rather than relying solely on the aggregate value. These applications require validation appropriate to their setting; the score should not be treated as a disclosure-quality rating, compliance test, or certification of credible transition commitments.

The main limitations are the absence of a human-coded benchmark. An archive-based annual reports sample with changing coverage, and potential sensitivity to text volume and measurement choices. Future research should prioritize independently coded selected and unselected passages, extraction audits, and preservation of match-level records. Subsequent tests should examine text-length normalization and distinguish sensitivity to category weights from sensitivity to category inclusion. Evaluation across firms, countries, and reporting regimes should precede broader application. Relationships with independently measured environmental outcomes warrant separate examinations. Together, these extensions would establish classification performance and identify the conditions under which NZCTD can support more robust comparisons of corporate climate-transition disclosure.

Declaration

Funding
No funding was received for conducting this study or preparing this manuscript.

Ethics declaration

Not applicable. This study analyzed corporate annual reports at the company level and did not involve human participants, human biological material, or animals.

Competing interests

The author declares no competing interests.

References

Agustini, A. T., Mustakini, J. H., Arifa, C., & Barokah, Z. (2026). Investigating climate-related disclosure: bibliometric insights and a critical evaluation of measurement methods. *International Journal of Disclosure and Governance*, 1-28. https://doi.org/10.1057/s41310-026-00388-6

Bochkay, K., Brown, S. V., Leone, A. J., & Tucker, J. W. (2023). Textual analysis in accounting: What's next?. *Contemporary accounting research*, *40*(2), 765-805. https://doi.org/10.1111/1911-3846.12825

Chan, K. J. D., Cheung, B., & Shen, L. Y. (2024). An economic foundation for assessing the credibility of corporate net zero transition pathways. *Business Strategy and the Environment*, *33*(8), 8868-8881.https://doi.org/10.1002/bse.3951

Cole, R. (2024). Inter-rater reliability methods in qualitative case study research. *Sociological Methods & Research*, *53*(4), 1944-1975. https://doi.org/10.1177/00491241231156971

David, L. K., Wang, J., Luo, M., & Cisse, I. I. (2025). Green signaling or greenwashing? ESG disclosure, firm performance and capital allocation in BRICS equity markets. *Journal of Capital Markets Studies*, *9*(2), 169-190. https://doi.org/10.1108/JCMS-06-2025-0080

Dell'Atti, S., Foglia, M., & Onorato, G. (2026). From Text to Value: Measuring and Pricing Firm Climate Risk Exposure. *Business Strategy and the Environment*. https://doi.org/10.1002/bse.70896

Ding, D., Liu, B., & Chang, M. (2023). Carbon emissions and TCFD-aligned climate-related information disclosures. *Journal of Business Ethics*, *182*(4), 967-1001. https://doi.org/10.1007/s10551-022-05292-x

Dumay, J., & Cai, L. (2015). Using content analysis as a research methodology for investigating intellectual capital disclosure: A critique. *Journal of Intellectual Capital*, *16*(1), 121-155. https://doi.org/10.1108/JIC-04-2014-0043

Fisher, I. E., Garnsey, M. R., & Hughes, M. E. (2016). Natural language processing in accounting, auditing and finance: A synthesis of the literature with a roadmap for future research. *Intelligent Systems in Accounting, Finance and Management*, *23*(3), 157-214. https://doi.org/10.1002/isaf.1386

Garrido-Merchán, E. C., González-Barthe, C., & Coronado-Vaca, M. (2026). Fine-tuning ClimateBert transformer with ClimaText for the disclosure analysis of climate-related issues in corporates' financial and non-financial reports. *Neural Computing and Applications*, *38*(1), 12. https://doi.org/10.1007/s00521-025-11820-3

Gebhardt, M., Schneider, A., Siedler, F., Ottenstein, P., & Zülch, H. (2024). Climate reporting in the fast lane? The impact of corporate governance on the disclosure of climate-related risks and opportunities. *Business Strategy and the Environment*, *33*(7), 7253-7272. https://doi.org/10.1002/bse.3852

Gutiérrez Ponce, H., Chamizo González, J., & Al-Mohareb, M. (2023). Annual reports readability from linguistic and communication perspectives: Systematic literature review. *Business and Professional Communication Quarterly*, *86*(4), 446-497. https://doi.org/10.1177/23294906231151893

Hickman, L., Thapa, S., Tay, L., Cao, M., & Srinivasan, P. (2022). Text preprocessing for text mining in organizational research: Review and recommendations. *Organizational Research Methods*, *25*(1), 114-146. https://doi.org/10.1177/1094428120971683

Huang, A. H., Wang, H., & Yang, Y. (2023). FinBERT: A large language model for extracting information from financial text. *Contemporary Accounting Research*, *40*(2), 806-841. https://doi.org/10.1111/1911-3846.12832

Kajal, A., & Bansal, S. (2025). Analysing impact of corporate attributes on sustainability disclosures through India's new BRSR framework. *International Journal of Law and Management*, *67*(2), 282-298. https://doi.org/10.1108/IJLMA-02-2024-0043

Khalid, A. M., & Rawat, P. S. (2025). Corporate environmental disclosures and role of top management: Evidence based on the business responsibility and sustainability reporting in India. *Corporate Social Responsibility and Environmental Management*, *32*(4), 4619-4635. https://doi.org/10.1002/csr.3205

Kolk, A., Levy, D., & Pinkse, J. (2008). Corporate responses in an emerging climate regime: The institutionalization and commensuration of carbon disclosure. *European accounting review*, *17*(4), 719-745. https://doi.org/10.1080/09638180802489121

Kudłak, R. (2025). Greenwashing or striving to persist: An alternative explanation of a loose coupling between corporate environmental commitments and outcomes: Greenwashing or striving to persist. *Journal of Business Ethics*, *197*(2), 355-370. https://doi.org/10.1007/s10551-024-05778-w

Loughran, T., & McDonald, B. (2011). When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. *The Journal of Finance*, *66*(1), 35-65. https://doi.org/10.1111/j.1540-6261.2010.01625.x

Loughran, T., & McDonald, B. (2016). Textual analysis in accounting and finance: A survey. *Journal of accounting research*, *54*(4), 1187-1230. https://doi.org/10.1111/1475-679X.12123

Milne, M. J., & Adler, R. W. (1999). Exploring the reliability of social and environmental disclosures content analysis. *Accounting, auditing & accountability journal*, *12*(2), 237-256. https://doi.org/10.1108/09513579910270138

Ngo, T., Le, T., Ullah, S., & Trinh, H. H. (2023). Climate risk disclosures and global sustainability initiatives: A conceptual analysis and agenda for future research. *Business Strategy and the Environment*, *32*(6), 3705-3720. https://doi.org/10.1002/bse.3323

O'Connor, C., & Joffe, H. (2020). Intercoder Reliability in Qualitative Research: Debates and Practical Guidelines. *International Journal of Qualitative Methods*, *19*. https://doi.org/10.1177/1609406919899220

Reid, S. W., McKenny, A. F., & Short, J. C. (2023). Synthesizing best practices for conducting dictionary-based computerized text analysis research. Vol. 14, pp. 43–78. https://doi.org/10.1108/S1479-838720220000014004

Sautner, Z., Van Lent, L., Vilkov, G., & Zhang, R. (2023). Firm-level climate change exposure. *The Journal of Finance*, *78*(3), 1449-1498. https://doi.org/10.1111/jofi.13219

Short, J. C., Broberg, J. C., Cogliser, C. C., & Brigham, K. H. (2010). Construct validation using computer-aided text analysis (CATA) an illustration using entrepreneurial orientation. *Organizational research methods*, *13*(2), 320-347. https://doi.org/10.1177/1094428109335949

Tumewang, Y. K., Ntim, C. G., & Haque, F. (2025). Task Force on Climate-Related Financial Disclosures: A Systematic Literature Review and Future Research Agenda. *Business Strategy and the Environment*, *34*(8), 10033-10058. https://doi.org/10.1002/bse.70112

Wahyuningrum, I. F. S., Ihlashul'amal, M., Utami, S., Djajadikerta, H. G., & Sriningsih, S. (2024). Determinants of carbon emission disclosure and the moderating role of environmental performance. *Cogent Business & Management*, *11*(1). https://doi.org/10.1080/23311975.2023.2300518

Walker, K., & Wan, F. (2012). The harm of symbolic actions and green washing: Corporate actions and communications on environmental performance and their financial implications. *Journal of Business Ethics*, *109*(2), 227-242. https://doi.org/10.1007/s10551-011-1122-4

Tables

Table 1: NZCTD lexicon categories, weights, and illustrative terms

Category	Content captured	Weight	Patterns (n)	Illustrative terms
Scoring categories				
Core net-zero	Net-zero and decarbonization language	4	11	net zero; carbon neutral
Emissions inventory	Emissions and inventory-related language	4	24	GHG; carbon footprint
Scope boundary	Emissions-scope and boundary references	3	16	Scope 1; Scope 3
Target commitment	Climate-target and target-setting language	3	18	science-based target; SBTi
Time and baseline	Temporal and baseline references	2	9	baseline year; by 2030
Quantification	Numerical and measurement expressions	2	8	tCO2e; MWh
Transition plan	Transition-planning language	3	13	transition plan; decarbonization plan
Energy actions	Energy-related action language	1	23	renewable energy; green hydrogen
Operational actions	Operational and product-related action language	1	18	EV fleet; green building
Value-chain actions	Value-chain action language	1	12	supplier engagement; logistics emissions
Offsets and removals	Offsetting and carbon-removal language	2	17	carbon offset; carbon credit
Performance progress	Reported progress and emissions changes	3	13	on track to; emissions decreased
Governance and accountability	Oversight and incentive-related language	2	13	ESG committee; climate-linked remuneration
Assurance and validation	Assurance and validation references	2	13	externally assured; ISAE 3000
Finance and materiality	Climate investment, financing, and risk language	2	17	green capex; transition risk
Frameworks	Reporting framework and initiative references	1	25	TCFD; CDP
Zero-weight categories				
Weak context	Generic sustainability-related language	0	15	sustainability; ESG
External-actor flag	External-actor target references	0	8	India's net-zero government target
Total	18 categories	—	273	

Notes: Pattern counts to lexicon entries, not observed matches. Patterns are case-insensitive regular expressions; illustrative terms are plain-language examples rather than complete matching rules. Weights apply to qualifying matches; zero-weight categories do not contribute directly to the weighted score. Section 3.5 specifies the safeguards for weak-context terms, flagged ambiguous acronyms, and bare-number quantification matches. External-actor references are tagged but not excluded from scoring. Matches identify disclosure language, not verified commitments, implementation, or climate outcomes.

Table 2: Contextual Qualification Rules and External-Actor Tagging

Mechanism	Applies to	Qualification rule or tagging action
Gate 1: Context	Weak-context matches	Qualifies only if the same paragraph contains a match with weight ≥ 3 or matches from at least two other distinct substantive categories.
Gate 2: Ambiguous acronyms	Matches on PPA, REC, I-REC, CSO, CDR, VCS, or TPT	Qualifies only if the same paragraph contains a separate match with weight ≥ 2.
Gate 3: Quantification	Quantification-category matches on bare, unlabelled numbers	Qualifies only if the same paragraph contains a match from one of the six supporting categories listed below.
External-actor tag	Paragraphs containing an external-actor match	Flags every match in the paragraph without changing its qualification or contribution to the score.

Notes: Rules operate at the paragraph level. Gate 3 supporting categories are core net-zero, emissions inventory, scope boundary, target commitment, energy actions, and performance progress. Qualification determines whether a match enters the qualifying category count; its contribution to the weighted score depends on the category weight. Weak-context and external-actor categories have weight zero. Contextual qualification does not establish semantic accuracy, company ownership of a commitment, or actual climate performance. Section 3.5 describes the scoring procedure.

Table 3. Descriptive statistics for NZCTD scores and ClimateBERT corroboration

Statistic	NZCTD Score	ClimateBERT corroboration ratio
Available observations	4,048	3,969
Mean	319.6	0.566
Standard deviation	493.7	0.219
Minimum	0.0	0.000
25th percentile	53.8	0.417
Median	140.0	0.609
75th percentile	395.0	0.737
Maximum	8,674.0	1.000
Skewness	4.87	—

Notes: The sample comprises 4,048 company-year observations from 468 companies. The NZCTD score is the weighted sum of qualifying category-level matches (Section 3.5). The ClimateBERT corroboration ratio is the proportion of lexicon-qualifying paragraphs that ClimateBERT classifies as climate-related (Section 3.6); it represents model-based corroboration, not classification accuracy or disclosure validity. Ratio statistics summarize available company-year ratios, excluding 79 observations that were unavailable and retaining valid ratios of zero. Score statistics include zero-score observations. The skewness is not reported for the corroboration ratio.

Table 4. Annual Sample Coverage, NZCTD Scores, and ClimateBERT Corroboration

Fiscal-year start	Company-years (n)	Mean NZCTD score	Median NZCTD score	ClimateBERT (n)	Mean corroboration ratio	Zero-score observations (n)
2013	5	32.2	26.0	5	0.614	0
2014	146	53.3	25.0	142	0.536	4
2015	285	75.8	41.0	283	0.441	2
2016	313	99.3	54.0	309	0.477	4
2017	338	113.1	58.0	337	0.479	1
2018	346	136.5	70.5	345	0.474	1
2019	355	146.3	81.0	353	0.496	2
2020	375	206.8	100.0	375	0.537	0
2021	387	292.6	151.0	386	0.581	1
2022	405	468.6	318.0	404	0.676	1
2023	432	579.8	414.0	427	0.671	5
2024	458	611.3	463.5	403	0.678	55
2025	203	909.7	683.0	200	0.684	3
Total	**4,048**	—	—	**3,969**	—	**79**

Notes: Fiscal years are identified by their starting year; for example, 2013 denotes FY2013–14. The NZCTD score is the weighted sum of qualifying category-level matches. The ClimateBERT corroboration ratio is the proportion of lexicon-qualifying paragraphs that ClimateBERT classifies as climate-related. ClimateBERT(n) denotes the company-years with an available ratio. Mean corroboration ratios are arithmetic means of available company-year ratios, not pooled paragraph-level proportions; unavailable ratios are excluded, while valid ratios of zero are retained. The 79 unavailable ratios are reported to coincide with the 79 zero-score observations. Annual score summaries include zero-score observations. The panel is unbalanced, and annual comparisons reflect changing company coverage rather than necessarily representing within-company changes.

Table 5. Category-level prevalence and contribution to the aggregate NZCTD score

Category	Weight	Prevalence (% of company-years)	Share of aggregate NZCTD score (%)
Positive-weight categories			
Emissions inventory	4	75.94	28.03
Quantification	2	76.48	13.75
Scope boundary	3	43.80	10.68
Energy actions	1	94.84	10.28
Assurance and validation	2	97.16	9.51
Time and baseline	2	86.36	7.33
Core net-zero	4	38.93	5.89
Frameworks	1	52.82	4.07
Transition plan	3	68.35	2.94
Operational actions	1	76.80	1.90
Target commitment	3	46.76	1.67
Governance and accountability	2	34.29	1.30
Offsets and removals	2	33.47	1.24
Finance and materiality	2	16.67	0.64
Performance progress	3	27.94	0.51
Value-chain actions	1	21.62	0.25
Zero-weight categories			
Weak context	0	77.50	0.00
External-actor flag	0	2.89	0.00

Notes: The sample comprises 4,048 company-year observations. Prevalence is the percentage of company years with at least one qualifying match in the relevant category. Categories can occur within the same company-year; therefore, prevalence percentages need not sum to 100%. Each category's score share equals its qualifying-match count summed across company-years and multiplied by its weight, divided by the aggregate NZCTD score across the sample (1,293,664), expressed as a percentage. These shares are not averages of company-year contribution percentages. Positive-weight categories are ordered by descending score share. Displayed shares sum to 99.99% because of rounding. Weak context and the external-actor flag contribute zero directly because their weights are zero.

Table 6. Score Sensitivity, Inter-Category Correlations, and Principal Component Analysis

Diagnostic	Result	Comparison
Weighted–unweighted score comparison		
Pearson correlation	0.9855	NZCTD score versus the unweighted sum of qualifying counts across all eighteen categories
Spearman's rank correlation	0.9919	Same comparison, based on company-year ranks
Inter-category correlations		
Mean pairwise Pearson correlation	0.401	Across all 153 distinct category pairs
Highest Pearson correlation	0.882	Emissions inventory and scope boundary
Lowest pairwise Pearson correlation	0.031	Quantification and external-actor flag
Principal component analysis		
PC1 explained variance	46.9%	First principal component
PC2 explained variance	8.1%	Second principal component
Cumulative explained variance, PC1–PC3	60.6%	First three principal components
Pearson correlation between PC1 scores and NZCTD score	0.8655	Component scores versus the weighted disclosure score

Notes: N = 4,048 company-year observations. Inter-category correlations use category-level qualifying-match counts. Principal component analysis (PCA) uses the eighteen standardized (z-scored) category-count variables, including weak context and the external-actor flag. PC denotes principal component. The unweighted comparator sums qualifying counts across all eighteen categories, whereas weak context and the external-actor flag contribute zero to the weighted NZCTD score. Thus, the comparison changes both the categories' weights and their inclusion in the score. The PC1–NZCTD correlation is an internal association between summaries derived from the same underlying counts. Explained-variance percentages describe the observed component structure.