AI-Driven Analysis of Climate-Related Financial Risks in Corporate Annual Reports

Mertin Baker
Department of Computer Science, University of Alabama at Birmingham, Birmingham, AL, USA.
martinwork@uab.edu

Chongchong Tu
Department of Electrical Engineering and Computer Science, University of Kansas, Lawrence, KS, USA.
du1971@ku.edu

Jingwan Tu
Department of Computer Science, Colorado State University, Fort Collins, CO, USA.
jhu@colostate.edu

Quentin A. Edwards
Department of Electrical Engineering and Computer Science, University of Missouri, Columbia, MO, USA.
quentinmail@missouri.edu

Abstract

Climate-related financial risks have become a central concern for corporations, investors, and regulators, yet the information required to assess these risks remains embedded in unstructured narrative disclosures within corporate annual reports. This paper develops a system-level examination of AI-driven approaches for analyzing climate-related financial risk disclosures, with particular attention to the structural, governance, and infrastructural conditions that shape effective deployment. The discussion integrates regulatory disclosure frameworks, corpus construction practices, natural language processing architectures, and financial risk management requirements. It considers how domain-specific language models and semantic anomaly detection can support the identification of material climate risk language, while also highlighting the trade-offs among specificity, interpretability, fairness, robustness, and computational sustainability. The paper argues that AI systems for climate risk disclosure analysis should not be treated as isolated text analysis tools, but as components of broader socio-technical infrastructures that include data governance, auditability, regulatory alignment, and institutional accountability. It further examines the implications of model opacity, distributional shift, and carbon-intensive computation. The analysis draws on current disclosure initiatives, empirical research on climate-related textual analysis, and developments in explainable AI to propose a governance-oriented research agenda. The paper concludes that long-term value from AI-driven climate risk analysis depends less on algorithmic sophistication alone and more on the integration of AI outputs into transparent, contestable, and sustainable institutional processes.

Keywords

climate risk disclosure, natural language processing, corporate annual reports, financial risk management, explainable AI, sustainability governance.

1. Introduction

Climate change is increasingly understood as a source of systemic financial risk rather than a peripheral environmental concern. Physical risks such as extreme weather events, sea level rise, and chronic shifts in temperature interact with transition risks arising from policy changes, technological disruption, and changing market preferences. Corporate annual reports have become a primary site where these risks are described, qualified, and sometimes obscured. The narrative sections of annual reports, including risk factors, management discussion and analysis, and sustainability disclosures, contain dense textual signals that are difficult to aggregate at scale. At the same time, the volume and heterogeneity of corporate filings make manual analysis insufficient for regulators, auditors, and institutional investors seeking to compare climate risk exposures across firms and sectors. This creates a pressing need for machine-readable analytical systems that can process large corpora of unstructured text while remaining robust enough for high-stakes financial interpretation.

The regulatory environment has moved rapidly toward greater standardization of climate risk disclosure. The Task Force on Climate-related Financial Disclosures established widely adopted recommendations around governance, strategy, risk management, and metrics and targets [1]. The Intergovernmental Panel on Climate Change has produced increasingly granular assessments of physical climate impacts and their economic transmission channels [2]. Standard-setting bodies have begun to embed climate considerations within financial reporting frameworks, as reflected in the International Sustainability Standards Board climate disclosure standard [3]. In parallel, the European Union has adopted mandatory corporate sustainability reporting requirements that expand the scope of climate-related narrative information [4]. These developments increase both the opportunity and the obligation to apply advanced analytical methods to corporate disclosures.

Despite regulatory momentum, climate risk disclosure remains characterized by significant heterogeneity in language, granularity, and tone. Firms may use boilerplate language to minimize legal exposure, selectively emphasize favorable adaptation narratives, or describe risks in ways that are difficult to compare across industries. Automated text analysis offers a partial remedy by extracting risk-relevant statements, classifying their topical orientation, and estimating their salience. However, the deployment of AI in this domain creates new challenges related to model explainability, fairness across sectors and firm sizes, data provenance, and the computational costs of large language models. A system-level perspective is therefore necessary to understand how AI-driven climate risk analysis can be designed, governed, and integrated into financial decision-making without reproducing the opacity and inconsistency that the technology is intended to reduce.

This paper contributes a structured interdisciplinary analysis of AI-driven climate risk disclosure analysis. It examines the disclosure ecosystem as a socio-technical system, discusses data architecture and corpus construction, reviews AI and natural language processing approaches, and analyzes the measurement of materiality and temporal orientation. It then addresses governance, fairness, explainability, integration with financial risk infrastructure, structural trade-offs, and sustainability. The overarching argument is that AI systems in this field must be evaluated not only by accuracy metrics but also by their robustness, auditability, and alignment with institutional accountability mechanisms.

2. Climate-Related Financial Risk Disclosure as a Socio-Technical System

Climate-related financial risk disclosure operates at the intersection of corporate reporting, regulatory supervision, investor information processing, and environmental science. It is best understood as a socio-technical system in which textual artifacts, institutional rules, analytical tools, and human judgment interact. Corporate annual reports are legal documents shaped by securities regulation, accounting standards, and litigation concerns. Their climate risk narratives are not neutral descriptions of physical or transition exposures, but strategically constructed disclosures influenced by managerial incentives, reputational considerations, and uncertainty about future conditions. The United States Securities and Exchange Commission has sought to standardize climate risk disclosure through formal rulemaking, reflecting a shift from voluntary toward mandatory reporting expectations [5]. Institutional investors have also placed increasing importance on climate risk information, using it to assess portfolio vulnerabilities and engagement priorities [6].

The effectiveness of AI-driven analysis depends on recognizing the institutional context in which disclosures are produced. A firm may discuss climate risk in general terms while avoiding quantified scenario analysis because of uncertainty about future regulatory pathways. Another firm may provide extensive discussion of physical asset locations but omit transition risk exposure linked to carbon-intensive revenue streams. These choices are shaped by managerial discretion, legal constraints, and the absence of fully harmonized disclosure taxonomies. AI systems that simply search for climate-related keywords may conflate substantive disclosure with boilerplate language. More sophisticated systems must distinguish between risk identification, risk assessment, and risk management language, and must account for the fact that similar phrases may carry different meanings across sectors and jurisdictions.

From a system design perspective, climate risk disclosure analysis requires more than a model trained on labeled text. It requires a data pipeline that can ingest filings, preserve document structure, map textual passages to reporting categories, and track changes over time. It also requires governance mechanisms that define how model outputs are used in risk management and regulatory processes. The socio-technical lens highlights that model errors do not remain limited to technical systems; they can propagate into investment decisions, supervisory judgments, and public perceptions of corporate resilience. Consequently, architectural choices in data ingestion, model selection, and output presentation are inseparable from questions of institutional responsibility.

The heterogeneity of climate disclosures also poses a fundamental measurement challenge. Unlike financial statement line items, climate risk narratives lack a common unit of account. Some disclosures are qualitative, some are quantitative, and many are forward-looking statements with inherent uncertainty. The lack of standardized semantic structures means that AI models must learn to identify risk-relevant signals from context, negation, hedging, and temporal qualification. A robust system must therefore combine domain knowledge with statistical learning, and it must be designed to accommodate ongoing regulatory change. This requires modular architectures in which the disclosure taxonomy, annotation schema, and model components can evolve as reporting standards mature.

3. Data Architecture and Corpus Construction

The construction of a high-quality corpus is foundational for AI-driven climate risk analysis. Corporate annual reports are distributed across multiple jurisdictions, formats, and filing systems. In the United States, annual reports on Form 10-K contain risk factor sections that are particularly relevant for climate risk extraction, though these documents are often provided as HTML or PDF and may contain tables, exhibits, and embedded graphics. Textual analysis in finance has long recognized the importance of 10-K disclosures for understanding risk communication, and financial text researchers have developed methods to parse and analyze such documents at scale [7]. A similar logic applies to climate-related sections of European annual reports, integrated reports, and sustainability statements, which vary significantly in structure and language.

A well-designed data architecture for climate risk analysis must address ingestion, normalization, and linkage. Ingestion involves retrieving documents from regulatory repositories or corporate websites while preserving metadata such as filing date, company identifier, industry classification, and reporting period. Normalization involves converting heterogeneous file formats into machine-readable text while retaining section boundaries and paragraph structure. This is essential because climate risk language may appear in risk factors, management discussion, sustainability reports, or notes to financial statements, and the location of a statement can affect its interpretive weight. Linkage involves associating textual passages with firm-level attributes such as emissions data, geographic asset exposure, and financial statement variables so that text-derived risk measures can be connected to economic outcomes.

Beyond technical plumbing, corpus construction raises important questions about representativeness and bias. A corpus built primarily from large listed firms may underrepresent small and medium-sized enterprises, whose disclosures are often less standardized and less detailed. Language coverage is another concern, as global climate risk analysis requires multilingual processing across reporting regimes with different legal traditions and discursive norms. Historical coverage is equally important because climate risk language changes over time in response to regulatory developments, scientific assessments, and market events. Corpora that do not span multiple years and multiple regulatory environments may encode temporal and jurisdictional biases that limit generalizability.

Data governance is therefore a central component of corpus construction. Versioning, provenance tracking, and access controls are necessary to ensure that model training data can be audited and reproduced. In high-stakes financial applications, undocumented changes to a corpus can undermine the validity of downstream analyses and expose institutions to legal or regulatory challenge. A mature data architecture must also incorporate mechanisms for handling missing data, inconsistent labels, and evolving taxonomies. Rather than assuming that a fixed corpus can support all future analyses, system designers should build feedback loops that allow domain experts to correct annotation errors and refine category definitions over time.

4. AI and NLP Architectures for Climate Risk Extraction

Early approaches to financial text analysis relied heavily on dictionary-based methods, in which predefined word lists were used to classify documents according to sentiment or risk orientation. While such methods are transparent and computationally efficient, they often struggle with the contextual nuance of climate risk language. Modern natural language processing has moved toward contextual representation learning, with transformer-based encoders such as BERT providing the foundation for many financial and sustainability text applications [8]. These models learn contextual word representations that can capture differences between phrases like transition risk and transition opportunities, reducing the need for handcrafted keyword lists.

Domain-specific pretraining has further improved the performance of climate text models. ClimateBert, for example, was developed to handle climate-related vocabulary and reporting contexts that general language models may not represent well [9]. Research using such models has shown that corporate climate disclosures often exhibit patterns of cherry-picking and cheap talk, in which firms discuss climate initiatives without providing substantive quantitative commitments [10]. This line of work demonstrates the value of combining domain-specific language modeling with careful measurement of disclosure quality. Firm-level climate exposure measures derived from earnings call transcripts and annual reports have also been constructed using supervised machine learning, enabling researchers and practitioners to quantify climate risk discussions in a more systematic way [11].

Financial applications of natural language processing have extended beyond classification to the construction of risk factors and hedging portfolios. Studies have shown that text-based measures of climate news can be used to build portfolios that hedge climate risk, indicating that textual signals contain economically relevant information [12]. These methods rely on the attention architecture that underlies modern transformer models, which allows a model to weight relevant parts of a document when making predictions [13]. Attention-based representations support the extraction of complex semantic relationships across long passages and are particularly useful for climate risk narratives that span multiple sentences and regulatory sections.

Recent work has extended these ideas to the detection of structural deviation in risk factor disclosures. One approach models semantic anomalies in 10-K risk factor sections by combining representation learning with explainable artificial intelligence, allowing analysts to identify when a disclosure departs from expected structural patterns [14]. Such anomaly detection is valuable because material changes in climate risk disclosure may not always be signaled by the introduction of new keywords, but rather by shifts in the semantic organization of risk narratives. By treating disclosure structure as an object of analysis, these methods move beyond bag-of-words approaches and align more closely with the interpretive practices of financial analysts and auditors.

5. Measuring Materiality, Salience, and Temporal Orientation

A central challenge for AI-driven climate risk analysis is the measurement of materiality. Not every mention of climate change is financially material, and the same phrase may be material for one firm but immaterial for another. Machine learning models can be trained to classify passages according to financial materiality, but the definition of materiality itself varies across jurisdictions and user groups. In some regulatory contexts, materiality is judged from the perspective of a reasonable investor, while in others it includes broader environmental and social impacts. AI systems must therefore be able to encode multiple materiality frameworks or be configurable to the reporting regime in which they are deployed. This configuration requirement has implications for model architecture, annotation schema, and evaluation metrics, because a single model may not be appropriate across all regulatory settings.

Salience is related to but distinct from materiality. Salience refers to the prominence and specificity of climate risk language within a document. A highly salient disclosure might include quantitative scenario analysis, named physical locations, or explicit links between climate risks and financial statement line items. Low-salience disclosures may use generic language that acknowledges climate risk without identifying transmission channels or financial consequences. Measurement of salience often requires fine-grained sequence labeling and relation extraction, because the relevant information may be spread across multiple clauses. AI systems can learn to distinguish between substantive and symbolic disclosure by attending to linguistic markers of specificity, commitment, and causal reasoning.

Temporal orientation is another important dimension. Climate risk disclosures may refer to risks that have already materialized, risks expected within a short horizon, or long-term uncertainties spanning decades. Forward-looking statements are often accompanied by hedging language and safe harbor disclaimers, which complicates their interpretation. A model that does not distinguish between past observations and future projections may misclassify the nature of a disclosure. Robust temporal extraction therefore requires not only recognition of tense and temporal expressions but also an understanding of scenario-based language, in which firms describe possible future pathways under different climate and policy assumptions. The architecture of such systems must support the representation of time as a structured dimension rather than a simple lexical feature.

The measurement of these dimensions is constrained by data availability and annotation cost. Supervised learning requires high-quality labeled examples, but climate risk materiality and salience labels are difficult to produce consistently. Expert disagreement is common because materiality judgments require contextual knowledge of a firm's business model, geographic footprint, and capital structure. Crowd-sourced labels may be insufficient for high-stakes financial applications. Hybrid approaches that combine limited expert labels with weak supervision or active learning can help reduce annotation burden, but they introduce additional design complexity. System designers must balance the desire for rich measurement against the practical constraints of label quality, reproducibility, and auditability.

6. Governance, Fairness, and Explainability in AI Deployment

The deployment of AI systems for climate risk disclosure analysis raises significant governance challenges. Because these systems may influence investment decisions, credit assessments, and regulatory review, their outputs must be interpretable and justifiable to human stakeholders. Explainability is particularly important when models flag a firm as having abnormal climate risk disclosure or when model outputs are used to compare firms within a sector. Interpretable models or post hoc explanations can help analysts understand which textual features drove a particular classification. In high-stakes settings, there is a strong argument for using interpretable models whenever possible rather than relying solely on opaque black-box systems [15]. Documentation practices such as model cards can also support transparency by recording intended use, performance characteristics, and known limitations [16].

Fairness in climate risk analysis is not reducible to demographic parity. Instead, the relevant fairness concerns involve differential performance across sectors, firm sizes, regulatory jurisdictions, and disclosure practices. A model trained predominantly on large carbon-intensive firms may misclassify the disclosures of small service companies or firms in emerging markets. Similarly, a model optimized for English-language 10-K filings may fail when applied to European sustainability statements written in French or German. Fairness assessments must therefore be domain-specific and should include disaggregated performance reporting across relevant subgroups. This requires evaluation datasets that represent diverse disclosure styles and a governance process for identifying and mitigating performance disparities.

Human oversight is essential. AI outputs should be treated as decision support, not as automatic determinations of climate risk performance. Analysts must be able to trace model outputs to source passages, review the confidence of model predictions, and override model classifications when contextual knowledge suggests an alternative interpretation. Auditability also depends on maintaining a complete record of model versions, training data, and preprocessing steps. Without such records, it becomes difficult to reconstruct why a model produced a particular output, which undermines accountability in regulatory and legal contexts. Governance mechanisms should therefore include model versioning, data provenance logs, and regular independent review.

Regulatory alignment is another dimension of governance. AI systems used by financial institutions may be subject to existing risk management and model risk management requirements. Supervisors may expect that model validation includes conceptual soundness, data quality, and outcome analysis, all of which apply to climate text models. The integration of climate risk disclosure analysis into regulatory reporting pipelines must therefore anticipate supervisory expectations around documentation, stress testing, and internal controls. This is not merely a technical requirement but a structural condition for the legitimacy of automated climate risk analytics.

7. Integration with Financial Risk Management and Regulatory Infrastructure

The value of AI-driven climate risk disclosure analysis is realized only when its outputs are integrated into broader financial risk management processes. Banks, insurers, asset managers, and credit rating agencies increasingly need to incorporate climate risk into credit analysis, underwriting, portfolio construction, and stress testing. Text-derived climate risk measures can complement quantitative metrics such as emissions intensity, energy mix, and asset location. For example, a bank assessing a corporate borrower may use AI-extracted evidence of transition risk disclosure to adjust its internal risk scoring or to trigger additional due diligence. However, integration requires mapping AI outputs onto existing risk taxonomies and decision workflows. A text classification indicating high transition risk exposure is not directly usable unless it can be linked to credit risk parameters or portfolio exposure models.

System integration also involves temporal alignment. Annual reports are published periodically, but risk management processes operate on different cycles, including quarterly reviews, continuous monitoring, and event-driven assessments. AI systems must therefore be capable of incremental updating and backtesting, allowing risk managers to observe how climate risk language changes between reporting periods and how such changes correlate with financial outcomes. This requires a data architecture that supports versioned document collections and reproducible feature extraction. It also requires model monitoring, because the distribution of climate disclosure language shifts over time due to regulatory reforms, scientific developments, and market shocks. Without ongoing monitoring, model performance may degrade silently, producing misleading risk signals.

Regulatory infrastructure is another integration domain. Supervisors may use AI-driven analysis to screen filings for climate risk disclosure deficiencies, identify outlier firms, or prioritize examinations. Automated screening tools can increase the scale and consistency of supervisory review, but they must be designed to minimize false positives that could impose unnecessary compliance burdens. Supervisory use of AI also requires transparency about model objectives and limitations, because regulated firms may challenge findings based on opaque algorithms. The credibility of such tools depends on their alignment with legal definitions of materiality and their ability to generate evidence that can be reviewed using established administrative procedures.

Interoperability between AI systems and reporting standards is essential. Climate disclosure frameworks such as the TCFD recommendations and IFRS S2 provide conceptual categories that AI models can use as classification targets [1][3]. Mapping textual passages to these categories allows analysts to compare disclosures across firms and jurisdictions more systematically. However, interoperability also requires careful handling of overlapping and evolving taxonomies. A governance layer must manage taxonomy versions, mapping rules, and validation against regulatory texts. This layer ensures that AI-driven disclosure analysis remains consistent with the institutional language used by regulators and standard setters, thereby increasing the legitimacy and utility of automated outputs.

8. Structural Trade-offs, Robustness, and Sustainability

Every design choice in AI-driven climate risk analysis involves structural trade-offs. A model that is highly specialized for climate risk language may perform well on climate-specific corpora but generalize poorly to other risk domains. Conversely, a general financial text model may miss climate-specific semantics that are not represented in its training distribution. Similar trade-offs arise between precision and recall in risk extraction. A high-precision system may correctly identify a small set of clearly material climate disclosures while missing less explicit but economically relevant statements. A high-recall system may capture more potential risk language but increase the burden of false positives for analysts. These trade-offs cannot be resolved solely by adjusting model thresholds; they must be managed through application-specific evaluation criteria and human review workflows.

Robustness to distributional shift is a major concern. Climate disclosure language changes as regulatory expectations evolve and as firms respond to new climate-related events. A model trained on historical filings may not recognize emerging terms such as climate litigation, nature-related financial risks, or transition finance. Adversarial actors may also deliberately modify disclosure language to evade automated detection or to project a favorable image without substantive changes in risk exposure. Robustness therefore requires continuous model retraining, anomaly detection, and adversarial evaluation. System designers should incorporate monitoring mechanisms that detect drift in input distributions and alert governance bodies when model behavior deviates from validated performance.

Sustainability is an increasingly important but often overlooked dimension of AI system design. Large language models require substantial computational resources for training and inference, and their carbon footprint can be significant. Research has documented the energy and policy considerations associated with deep learning for natural language processing [17]. The trend toward ever larger models has raised concerns about environmental impact, data governance, and the concentration of AI capacity among a small number of well-resourced organizations [18]. The carbon footprint of large-scale language models must be measured and disclosed if climate risk analytics are to avoid undermining their own normative objectives [19]. System-level sustainability therefore involves evaluating the full lifecycle energy cost of model development, including data collection, pretraining, fine-tuning, and continuous deployment.

Sustainable AI deployment also requires attention to institutional capacity. Many smaller organizations, including public sector supervisors and non-governmental organizations, may lack the computational infrastructure to train or even fine-tune large language models. This creates a risk that advanced climate risk analytics become concentrated among large financial institutions and technology providers. To promote broader access, system architectures can include shared model repositories, open annotation schemas, and modular components that can be deployed in resource-constrained environments. Sustainability in this sense extends beyond environmental impact to include the resilience and equity of the analytical ecosystem.

9. Future Directions and Policy Implications

Future research in AI-driven climate risk disclosure analysis should move toward deeper integration of textual, financial, and physical data. Linking narrative disclosures to asset-level geolocation data, supply chain information, and climate model outputs would enable more precise estimates of physical risk exposure. Transition risk analysis could be strengthened by connecting disclosure language to policy tracking databases and technology cost curves. Such integration requires advances in entity resolution, knowledge graphs, and multimodal representation learning. It also requires careful attention to data privacy, intellectual property, and the legal status of derived risk signals.

Another promising direction is the use of causal inference methods to understand whether disclosure changes precede or follow changes in real climate risk exposure. Observational studies have shown that market attention to global warming varies over time, suggesting that investor processing of climate information is not stable [20]. AI systems that simply measure textual change may confuse shifts in attention with shifts in underlying risk. Causal frameworks, combined with careful identification strategies, could help separate substantive disclosure from strategic communication. This would improve the usefulness of AI-generated metrics for both investors and regulators.

Policy implications are substantial. Regulators should consider establishing minimum standards for the auditability and transparency of AI systems used in climate risk disclosure analysis. These standards could include requirements for model documentation, data provenance, validation against expert judgments, and public disclosure of known limitations. Standard setters may also encourage the development of open benchmark datasets that reflect diverse jurisdictions, languages, and firm sizes. Without such public infrastructure, it will be difficult to compare AI systems or to hold vendors accountable for performance claims.

International coordination is necessary because climate risk is a global phenomenon while disclosure regimes remain fragmented. AI systems trained in one regulatory context may not transfer cleanly to another. Policy frameworks should therefore support cross-border research collaboration, shared vocabularies, and interoperability between reporting taxonomies. The goal should not be to impose a single technical architecture, but to ensure that AI-driven climate risk analysis supports comparable, transparent, and decision-useful information. This requires sustained investment in public data infrastructure and institutional capacity building, especially for smaller regulators and civil society users.

10. Conclusion

AI-driven analysis of climate-related financial risks in corporate annual reports offers substantial potential to improve the scale, consistency, and depth of climate risk assessment. At the same time, it introduces complex system-level challenges that cannot be addressed through model accuracy alone. The architecture of data pipelines, the choice of language models, the measurement of materiality and temporal orientation, and the integration of model outputs into risk management and regulatory processes all shape the ultimate usefulness and legitimacy of these systems. Governance, fairness, explainability, robustness, and sustainability are not secondary concerns but constitutive elements of a credible analytical infrastructure.

This paper has argued for a socio-technical perspective in which AI systems are understood as embedded within institutional and regulatory environments. Corporate climate disclosures are strategically produced narratives, not neutral data, and automated systems must account for this reality through careful corpus design, contextual representation learning, and continuous monitoring. The use of domain-specific language models and semantic anomaly detection can help identify substantive disclosure patterns, but only when complemented by human oversight and traceable model governance.

Policy and research efforts should focus on building shared data infrastructure, promoting transparency, and ensuring that advanced climate risk analytics are accessible beyond a small set of large institutions. Future work should also integrate textual analysis with physical and financial data, develop causal approaches to disclosure interpretation, and evaluate the environmental footprint of the AI systems themselves. Ultimately, the long-term contribution of AI to climate risk disclosure analysis will depend on its ability to support accountable, fair, and sustainable decision-making in the face of deep uncertainty.

References

1. Task Force on Climate-related Financial Disclosures. (2017). Recommendations of the Task Force on Climate-related Financial Disclosures. Financial Stability Board.

2. Intergovernmental Panel on Climate Change. (2022). Climate change 2022: Impacts, adaptation and vulnerability. Contribution of Working Group II to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge University Press.

3. IFRS Foundation. (2023). IFRS S2 climate-related disclosures. International Sustainability Standards Board.

4. European Commission. (2022). Directive (EU) 2022/2464 of the European Parliament and of the Council of 14 December 2022 amending Regulation (EU) No 537/2014, Directive 2004/109/EC, Directive 2006/43/EC and Directive 2013/34/EU, as regards corporate sustainability reporting. Official Journal of the European Union, L 322, 15-80.

5. U.S. Securities and Exchange Commission. (2024). The enhancement and standardization of climate-related disclosures for investors. SEC Release No. 33-11275.

6. Krueger, P., Sautner, Z., & Starks, L. T. (2020). The importance of climate risks for institutional investors. Review of Financial Studies, 33(3), 1067-1114.

7. Loughran, T., & McDonald, B. (2011). When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. Journal of Finance, 66(1), 35-65.

8. Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. Proceedings of NAACL-HLT 2019, 4171-4186.

9. Webersinke, N., Kraus, M., Bingler, J. A., & Leippold, M. (2022). ClimateBert: A pretrained language model for climate-related text. arXiv.

10. Bingler, J. A., Kraus, M., & Leippold, M. (2022). Cheap talk and cherry-picking: What ClimateBert has to say on corporate climate risk disclosures. Finance Research Letters, 47, 102776.

11. Sautner, Z., van Lent, L., Vilkov, G., & Zhang, R. (2023). Firm-level climate change exposure. Journal of Finance, 78(3), 1449-1498.

12. Engle, R. F., Giglio, S., Kelly, B., Lee, H., & Stroebel, J. (2020). Hedging climate change news. Review of Financial Studies, 33(3), 1184-1216.

13. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. Advances in Neural Information Processing Systems, 30, 5998-6008.

14. Sun, F., He, S., Wang, R., Ke, L., Shen, H., & Liao, Q. (2026). Modeling Structural Deviation in 10-K Risk Factors: A Semantic Anomaly Detection and Explainable AI Approach. Risks, 14(4), 87.

15. Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. Nature Machine Intelligence, 1(5), 206-215.

16. Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I. D., & Gebru, T. (2019). Model cards for model reporting. Proceedings of the Conference on Fairness, Accountability, and Transparency, 220-229.

17. Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, 3645-3650.

18. Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, 610-623.

19. Luccioni, A. S., Viguier, S., & Ligozat, A.-L. (2023). Estimating the carbon footprint of BLOOM, a 176B parameter language model. Journal of Machine Learning Research, 24(253), 1-15.

20. Choi, D., Gao, Z., & Jiang, W. (2020). Attention to global warming. Review of Financial Studies, 33(3), 1112-1145.