

Automated ESG Compliance Auditing using Retrieval-Augmented Generation and Semantic Vector Search: A Computational Economics Approach

Aditi Shastri¹, Aditya Karumbaiah G U², Aaditya S Rao³, Arshia Sirohi⁴, Angela Varghese⁵ and Gangadhar Angadi⁶

¹⁻⁶Dept. of ISE, RV College of Engineering, Bengaluru, India

Email: aditishastri.is23@rvce.edu.in, adityakgu.is23@rvce.edu.in, aadityasrao.is23@rvce.edu.in, arshiasirohi.is23@rvce.edu.in, angelavarghese.is23@rvce.edu.in, gangadharangadi@rvce.edu.in

Abstract— ESG (Environmental, Social, Governance) is a framework used by investors and companies to assess an organization's sustainability, ethical impact, and long-term performance along with financial metrics. Given that only larger companies are required to report ESG metrics, corporate annual reports, which are mandatory for market listed companies can be used to estimate the ESG scores of a company. However, such reports are often long, fraught with technical jargon, and inconsistent narrative language. To address this challenge, this paper proposes an artificial intelligence based automated auditing methodology using the RAG (Retrieval-Augmented Generation) framework. Here, both disclosures and compliance queries are encoded as semantic vector embeddings and information are retrieved based on semantic similarity rather than keyword matching. A large language model then processes the collected evidence into a set of relevant disclosure categories, finally culminating in a proposed ESG score to assess firms and compare them with their peers. The proposed framework is expected to improve the scalability of ESG disclosure analysis and also enable cost effective automated verification of ESG reports using other reports and material published by the firm under scrutiny.

Index Terms— ESG Analytics, Retrieval-Augmented Generation, Semantic Search, Natural Language Processing and Sustainability Compliance.

I. INTRODUCTION

Environmental, Social, and Governance (ESG) considerations have now become a central element of corporate performance evaluation. ESG has moved well beyond its earlier role as a supplementary, non-financial indicator [1]. Corporate evaluation and review is the last stage in the Corporate Performance Management, it is necessary to ensure consistency between implementation and planned strategic actions, with a continuous measurement of the circumstances both inside and outside the company [2]. ESG has emerged as a means for companies to be evaluated on non-financial factors. Investors, regulators and policymakers are starting to recognize that issues regarding environmental degradation, workforce welfare, unethical practices and governance failures pose a direct material financial risk [3].

The significance of ESG is acknowledged, (International Sustainability Standards Board (ISSB), 2023) [4], in the way this report introduces and formalizes the notion of ESG within the world of accounting standards, emphasizing the integral part such disclosures play for the enterprise value calculation required by the investor community. In his letter to the CEOs, as the author of [5] puts it, "climate risk is investment risk, and we are seeing a fundamental shift of capital," as the world is moving toward an era in which "sustainability was increasingly part of the fabric."

Moreover, the substitution of ESGs is also visible in emerging economies such as India, where the government has developed a better sustainability assist the auditing process and enhance the transparency in reporting ESGs, has recently introduced the Business Responsibility and Sustainability Reporting (BRSR) form [6]. This mandate for reporting now requires the companies to share their information in the domains of environmental, social, and governance reporting. Though the BRSR has succeeded in increasing the quantum of ESG reporting, it has still not been able to overcome the problem of utilizing these reports to set new goals and gain insights in the domain of corporate reporting [7].

In current practices, the case of ESG-related information being dispersed in small amounts in vast and lengthy annual reports with page numbers exceeding the hundreds is a challenging one. Any relevant disclosure is stuck within management jargon, governance and legal terms, and various other vernacular parlance, making it difficult to decipher the true state of a company's promises to sustainability, social and governance good. Disclosures are frequently written using contrasting or unrelated terminology. Considering this great lack of standardization in the terminologies, language used and format of disclosures, it is understandable that there are multiple subjective interpretations for the same piece of information [8]. Stakeholders often face asymmetry, where corporate management holds an advantage over investors and external evaluators in the aspects of essential information and the firm's true sustainability posture. While reports like those of Bajaj Finance [9] and Infosys [10] have started to include ESG related information in separate chapters, an automated framework would allow for better contextual support.

Keyword searches or term frequency make up the foundation for early automated approaches in attempting factual interpretation of ESG text analysis [11]. They however, do not account for words implying the same idea but with differing labels across different firms. This yields systematic misidentification of governance practices and misrepresent adherence to regulations, leading to inaccurate and downright wrong ESG ratings and benchmarks. This approach causes distrust in the use of automation to interpret ESG text and it shows the need for approaches that employ more than just rudimentary text matching.

Through this work, an attempt to address these difficulties has been made. The work presents an automated ESG compliance auditing framework based on Retrieval-Augmented Generation (RAG), where both disclosures and compliance queries are encoded as semantic vector embeddings, and the information is retrieved based on semantic similarity rather than keyword matching [11], [12]. A lightweight large language model then processes the collected evidence into a set of disclosure categories relevant to the assessment of compliance, providing a consistent and interpretable outcome. The proposed framework is expected to improve the accuracy and scalability of ESG disclosure analysis in large and unstructured corporate reports.

II. THEORETICAL FRAMEWORK

This paper proposes an ESG auditing framework that uses semantic retrieval and generative classification. The theoretical basis for such an approach can be found in Agency Theory [13], Signaling Theory [14], and Legitimacy Theory [15] as explained below.

A. Agency Theory and ESG as a monitoring technique

The firm is visualised as a contract between principals, or shareholders and agents, or managers in Agency Theory [13]. As players have different and conflicting incentives, suboptimal outcomes may occur. For instance, managers may prioritize, say, short term financial performance, personal compensation, or reputational concerns, thus negatively impacting long term sustainability and compliance. Thus, ESG disclosures can serve as a monitoring mechanism in order to align managerial behaviour with shareholder interests.[4], [15]. While ESG disclosures offer enhanced transparency, there are compliance risks of the regulatory, conduct and reputational nature [16] This work proposes the use of corporate reports instead to identify the ESG performance of firms. Corporate reports, however, are voluminous, unstructured, and linguistically complex. Adding to this, they can often exceed several hundred pages, making it impractical and inefficient to manually extract and verify ESG-relevant information [9], [8]. The proposed Retrieval-Augmented Generation framework [12] addresses this limitation by automating the identification and extraction of ESG compliance evidence. This framework aims to

thus reduce cost of computing ESG compliance metrics in the absence of ESG reports. Even when such reports are present, it can serve as a tool to compare and contrast with the evidence extracted from annual corporate reports.

B. Signaling Theory and the Difference Between Symbolic and Substantive Disclosure

Signaling Theory deals with information asymmetry between actors. Firms with higher transparency may thus be perceived as more high quality [14]. A large portion of sustainability reporting consists of aspirational or narrative statements, that depict commitment to ethical practices and environmental responsibility. Since it is inexpensive to create such reports, there exists an equilibrium where both high- and low-quality firms seem similar in nature on the basis of their reporting alone, which is allowed to exist due to the prevailing information asymmetry. Thus, vague and abstract disclosures can be interpreted as 'weak' signals and quantitative ones with formal mechanisms as 'strong'. Reporting measurable outcomes, numerical targets, or structural oversight bodies implies operational investment and exposes firms to risks, which further adds to the credibility of classifying such signals as 'strong'. The proposed framework thus leverages Signaling Theory by classifying disclosures into commitment only statements, formal policies, and quantitative evidence.

C. Legitimacy Theory and the Credibility of a Disclosure

Legitimacy Theory proposes that organizations seek to maintain social acceptance by aligning their activities with prevailing norms and values[15]. ESG reporting is one such technique through which firms can show compliance with such standards[5]. If enforcement is weak, however, such disclosures may prioritise appearance over factual reporting[17]. Thus, firms are incentivised to strategically emphasize positive narratives while downplaying negative events. Superficial analysis then is a major limitation for automated systems that rely solely on keyword frequency or document length, since they may fail to capture the magnitude of negative disclosures [18]. The proposed system uses an aggregate score across the ESG domains. In order to compensate for the potential downplaying of negative disclosures in reports, the framework adopts negatively weighted scores for adverse indicators.

III. METHODOLOGY

The proposed framework consists of three phases, namely, data ingestion and text chunking, semantic embedding and retrieval and generative classification with scoring. It is a typical Retrieval-augmented generation pipeline with necessary changes as needed to benefit the purpose of ESG scoring.

A. Data Ingestion and Text Chunking

This is the first phase in the framework, which aims to clean the provided reports and create chunks of the reports. These chunks are then vectorized and analyzed in subsequent steps. This phase is important because of a few reasons. The first being that corporate reports are typically lengthy PDF documents that lack a standardized structure, and directly embedding these documents is infeasible, both due to context window limitations and also because of computational cost. Directly vectorizing entire documents would also dilute the semantic relevance of particular disclosures, and thus prevent explainability[19]. Thus, the following preprocessing is undertaken. The raw text obtained from the document is cleaned by removing headers, footers, and page numbers, and a fixed length sliding window is used to create overlapping chunks. The specific sliding window size used in the experiments below was 1000 characters, with a 200-character overlap.

B. Semantic Embedding and Vector Retrieval

Unlike generic document summarization pipelines, the system operates on a predefined set of ESG compliance queries derived directly from the : This query-driven design ensures that chunking and retrieval remain aligned with regulatory intent rather than narrative prominence. The text chunks so obtained in the previous steps are transformed into dense vector representations using a pretrained sentence-level embedding model. The one used in the experiments undertaken in this study is all-MiniLM-L6-v2, available in the Python implementation of the sentence-transformers library [20]. Such dense embeddings capture context and meaning, thus allowing for semantic relationships between phrases which can be measured using similarity metrics[21]. The specific metric used in this study was cosine similarity. These obtained embeddings are compared against compliance queries which are derived from the Business Responsibility and Sustainability Reporting (BRSR) mandate as outlined in Table IV-A. In addition to the original set of BRSR guideline questions, a few queries pertaining to negative indicators are also used in the framework. The BRSR mandate is a compulsory framework (effective from FY 2022-23) formulated by the SEBI (Securities and Exchange Board of India) for the 1000 largest listed entities

(by market capitalization) to disclose their ESG performance. [7]. These compliance queries are vectorized into using the same embedding model as the document chunks, and cosine similarity is used to compute the relevance between a query and all the document chunks. The top k similar document chunks are selected as high signal [15] evidence for a particular query. The choice of k can influence how lenient or strict the mechanism is, with higher k being predisposed to more false positives and lower k susceptible to missing out on evidence. In this study, k=5 produced consistent results and was thus chosen [22], [21].

C. Generative Classification and Scoring

This is the final stage of the pipeline. The k most similar document chunks are passed to a large language model (LLM) with a prompt to classify it as one of four predefined categories: quantitative evidence, formal policy disclosure, commitment-only statement, or absence of disclosure. Gemma:2B [23] was chosen as the LLM due to its smaller size, so that the framework could be scaled with lesser computational cost and latency. Each of these 4 categories carries a different weight in the ESG scoring. The highest weight is allotted to quantitative evidence, since they represent strong and metric based signals. Formal policy disclosure and commitment-only statements receive lower scores.

Absence of disclosure receives the lowest, albeit positive weight. Document chunks for negative indicator queries are allotted negative penalty scores. These include questions on the presence of regulatory fines, litigations and other such compliance breaches. The final ESG score is then computed by aggregating weighted results across all the compliance queries.

IV. RESULTS AND DISCUSSION

The evaluation of the proposed framework was designed to assess its technical performance and the analytical usefulness. The experiments conducted tested whether semantic retrieval could extract relevant statements from corporate reports, and also whether quality of disclosures was evaluated (strong versus weak signals).

A. Qualitative Results

The system successfully identified relevant disclosure evidence for all mandatory Environmental, Social, and Governance queries across the evaluated companies, including Bajaj Finance and Infosys. Firms which employ structured and metric-driven reporting were seen to exhibit consistently higher scores than those relying mainly on narrative descriptions. Companies that reported explicit numerical values, such as emissions data, workforce statistics, or governance structures, achieved enhanced classification confidence due to stronger signal strength. Firms with limited quantitative evidence received lower scores despite extensive narrative sections. These results indicate that the framework prioritizes disclosure quality rather than document length, formatting or style of presentation.

The framework also was seen to have demonstrated resilience in handling terminological variation. Conceptually similar but linguistically different disclosures were correctly classified. Thus, semantic embeddings can be used to search for meanings and treat all synonyms similarly.

TABLE I: THE BRSR COMPLIANCE QUERY SET

Pillar	Audit Queries
Environment	Does the company have a stated environmental policy?
	Does the company report on GHG emissions or carbon footprint?
	Are renewable energy initiatives disclosed
	Is water usage or conservation discussed?
	Is waste management or recycling reported?
Social	Are employee health and safety measures disclosed?
	Is workforce diversity or inclusion reported?
	Are training and skill development initiatives described?
	Are community or CSR initiatives discussed?
Governance	Are human rights or labour practices disclosed?
	Are the board composition and independence disclosed?
	Are ethics or anti-corruption policies described?
	Is a whistleblower mechanism reported?
	Is executive remuneration disclosed?
Adverse Indicators	Is risk management or internal control discussed?
	Has the company paid fines for non-compliance?
	Are there pending litigations against the company?
	Has the company reported any cyber incidents?

V. CONCLUSIONS

This study presents an ESG compliance auditing framework that utilizes semantic vector search with retrieval-augmented generation. This framework focuses on contextual interpretation in contrast to traditional ESG analysis methods that depend primarily on keyword-based matching. It could thus be used to enhance comparability of compliance assessments across firms that use diverse formatting and styles.

The experimental results demonstrate that the system consistently detects relevant ESG related content in extensive annual reports. It also is able to differentiate between strong and weak signals, i.e. quantitative and qualitative metrics. The use of penalty mechanisms for adverse disclosures ensures balanced outcomes. The major advantage associated with this framework is its per metric explainability as to how the proposed ESG score was calculated.

The proposed framework could reduce the cost and effort associated with ESG monitoring. It thus could reduce information asymmetry in markets and increase transparency. These characteristics could be leveraged in large scale screening of textual reports, particularly in the domain of regulatory compliance verification.

A major limitation with this framework is that it can be applied only on publicly available reports and disclosures. Unreported operational practices cannot be captured by text-based analysis, regardless of the techniques or methods employed [8]. Comprehensive ESG evaluation thus depends on third party audits. External datasets, trend analysis, media coverage analysis could be explored in future work to improve the auditing and analysis of sustainability compliance.

ACKNOWLEDGMENT

The authors wish to thank the staff and faculty at RV College of Engineering for their guidance, their friends and family for the support provided, and the research community for their contributions

REFERENCES

- [1] Robert G. Eccles, Ioannis Ioannou, and George Serafeim. The impact of corporate sustainability on organizational processes and performance. *Management Science*, 60(11):2835–2857, 2014.
- [2] Ilori David Babafemi. Corporate Strategy, Planning and Performance Evaluation: A Survey of Literature. *Journal of Management Policies and Practices*, 3(1), 2015.
- [3] Gunnar Friede, Timo Busch, and Alexander Bassen. Esg and financial performance: Aggregated evidence from more than 2000 empirical studies. *Journal of Sustainable Finance & Investment*, 5(4):210–233, 2015.
- [4] International Sustainability Standards Board (ISSB). Overview of IFRS S1 and IFRS S2. tech report, IFRS Foundation, 6 2023. Released 26 June 2023; effective for periods beginning on or after 1 January 2024.
- [5] Larry Fink. Larry Fink CEO Letter | BlackRock. Technical report, 2020. In his annual letter to CEOs, BlackRock’s Larry Fink discusses how the events of 2020 are helping to accelerate an economic transformation.
- [6] Kundan Pandey. Sebi expands ESG focus to supply chains, brings green credits into disclosure norms. *Mongabay*, apr 1 2025. India’s environmental science and conservation news.
- [7] BRSR.in. Brsr - business responsibility and sustainability reporting. <https://brsr.in/>, 2025. Accessed: 2025-02-07.
- [8] Florian Berg, Julian Kolbel, and Roberto Rigobon. Aggregate confusion: The divergence of esg ratings. *Review of Finance*, 26(6):1315–1344, 2022.
- [9] BAJAJ FINANCE LIMITED. 38th Annual Report 2024-2025. tech report, Bajaj Finance Corporate Office, Viman Nagar, Pune, Pune, 2024.
- [10] ESG REPORT. Insight TO IMPACT. tech report 2024-25, null, null, 2024.
- [11] Philipp Baier, Marc Berninger, and Florian Kiesel. Environmental, social and governance reporting in annual reports: A textual analysis. *Financial Markets, Institutions & Instruments*, 29(3):93–118, jun 26 2020.
- [12] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Kuttler, Mike Lewis, Wen-tau Yih, Tim Rocktaschel, Sebastian Riedel, and Douwe Kiela. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 9459–9474. Curran Associates, Inc., 2020.
- [13] Michael C. Jensen and William H. Meckling. Theory of the firm: Managerial behavior, agency costs and ownership structure. *Journal of Financial Economics*, 3(4):305–360, 10 1976.
- [14] Michael Spence. Job Market Signaling. *The Quarterly Journal of Economics*, 87(3):355, 8 1973.
- [15] Mark C. Suchman. Managing Legitimacy: Strategic and Institutional Approaches. *The Academy of Management Review*, 20(3):571, 7 1995.
- [16] Association for Financial Markets in Europe (AFME) & Ernst & Young LLP. Esg and the role of compliance. Technical report, Association for Financial Markets in Europe, November 2022. Discussion paper on ESG compliance roles and practices.

- [17] Magali A. Delmas and Vanessa Cuerel Burbano. The Drivers of Greenwashing. *California Management Review*, 54(1):64–87, 10 2011.
- [18] TIM LOUGHRAN and BILL MCDONALD. When Is a Liability Not a Liability? Textual Analysis, Dictionaries, and 10ks. *The Journal of Finance*, 66(1):35–65, jan 6 2011.
- [19] Nicolas Webersinke, Mathias Kraus, Julia Anna Bingler, and Markus Leippold. Climatebert: A Pretrained Language Model for Climate-Related Text. Technical report, 2022. arXiv:2110.12010v3 [cs.CL], 17 Dec 2022.
- [20] Nils Reimers and Iryna Gurevych. Sentence-transformers: Multilingual sentence, paragraph and image embeddings with bert, 2020.
- [21] Nils Reimers and Iryna Gurevych. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. Technical report, 2019. arXiv:1908.10084v1 [cs.CL] 27 Aug 2019.
- [22] Christopher D. Manning, Prabhakar Raghavan, and Hinrich Schütze. " Introduction to Information Retrieval. Cambridge University Press, 2008.
- [23] Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Riviere, Mihir Sanjay Kale, Juliette Love, Pouya Tafti, et al. Gemma: Open models based on gemini research and technology. arXiv preprint arXiv:2403.08295, 2024. Accessed: 2025-02-07