Bridging Corporate Claims and Public Perception: Real-Time Validation of ESG Initiatives with Social-Media Analytics

Completed Research Paper

Nidhi Mishra
Karlsruhe Institute of Technology
Karlsruhe, Germany
nidhi.mishra@kit.edu

Moritz Diener
Karlsruhe Institute of Technology
Karlsruhe, Germany
moritz.diener@kit.edu

Daniel Heinz
Karlsruhe Institute of Technology
Karlsruhe, Germany
daniel.heinz@kit.edu

Philipp Spitzer
Karlsruhe Institute of Technology
Karlsruhe, Germany
philipp.spitzer@kit.edu

Philippe Cudré-Mauroux
University of Fribourg
Fribourg, Switzerland
pcm@unifr.ch

Abstract

Stakeholders increasingly question companies' Environmental, Social and Governance (ESG) claims, exposing a gap between self-reported metrics and public perception. We present a design-science artifact – a real-time ESG validation framework that mines social-media discourse to supply organizations with stakeholder feedback. The modular pipeline combines domain-adaptive topic modeling with transformer-based sentiment and emotion classification. A multi-industry case study (IT-consulting, food-and-beverage, tobacco) demonstrates that the system captures sector-specific ESG themes and detects sentiment shifts following announcements, controversies or initiatives. Social media emerges as a sensitive barometer: positivity dominates in sectors aligned with sustainability goals, while tobacco discourse remains negative. Event-based temporal analysis and crowdsourced annotation confirm the framework's accuracy. The study advances information systems research by offering (1) a replicable tool operationalizing social media for ESG accountability, (2) design principles uniting socio-technical relevance and machine learning, and (3) a methodological basis for examining when public discourse legitimizes or penalizes symbolic disclosure.

Keywords: ESG, social media sentiment, design-science research

Introduction

Corporate success is no longer judged solely by financial outcomes. Over the past decade, performance has been reframed through environmental, social, and governance (ESG) responsibility. Whereas earlier research treated responsible behavior as a reputational add-on, current capital-market practice embeds ESG indicators in credit ratings, executive compensation, and investment mandates (Barnett & Salomon, 2006; Friede et al., 2015; Sharma & Dangwal, 2022). This shift has drawn heightened attention from investors, regulators, employees, and civil-society actors, who now demand credible evidence of firms' environmental impact, labor conditions, and governance quality (Saxena et al., 2022). Yet prevailing assurance regimes remain dominated by self-reported metrics and heterogeneous disclosure standards (Johnson et al., 2018), masking under-performance and encourage symbolic "greenwashing" (Delmas & Burbano, 2011) rather than reflecting authentic stakeholder perceptions (Chopra et al., 2024). The resulting information asymmetry erodes stakeholder trust and limits the strategic value of ESG initiatives.

Information systems (IS) research is well positioned to address this credibility gap because it investigates how digital technologies create, transform and disseminate organizational knowledge (Alavi & Leidner, 2001; Bharadwaj et al., 2013; Massa et al., 2023). However, existing ESG-analytics prototypes often repurpose generic text-mining pipelines, privilege data volume over contextual meaning, and rarely progress beyond proof-of-concept in big-data analytics (Abbasi et al., 2016) or ESG platforms (Plugge et al., 2024). Most are evaluated ex post against expert opinion or secondary ratings, neglecting the ambient stakeholder feedback that continuously unfolds on social media (Kietzmann et al., 2011; Treem & Leonardi, 2013). Platforms such as Twitter/X and Reddit provide an unfiltered, real-time chronicle of public reactions to corporate behavior (Rodríguez-Ibánez et al., 2023), yet systematic approaches for converting this noisy stream into decision-ready ESG intelligence remain nascent.

Our study asks *how an IS artifact can systematically capture and interpret public sentiment toward corporate ESG initiatives in real time* so organizations can validate, benchmark and adapt their strategies. We answer through a design science research (DSR) project that follows Peffers et al. (2007) and Tuunanen et al. (2024) refinements. Building on calls to embed sustainability as a high-priority design objective in IS research (Schoormann et al., 2025; vom Brocke et al., 2013), we derive requirements for a context-aware sentiment model. These requirements guide the development of a modular analytics pipeline coupling domain-adaptive topic modeling (Jelodar et al., 2019) with context-aware sentiment and emotion classification based on machine learning and transformer architectures (Bello et al., 2023; Khurana et al., 2023). To validate our approach, we conducted a multi-industry, multi-method evaluation: event-based field validation IT-consulting, food-and-beverage, and tobacco, and crowdsourced ground-truth annotation for an IT-consulting company.

Our research contributes to IS literature in three ways. First, it provides a replicable artifact that operationalizes social media affordances for ESG accountability. Second, it articulates design principles that integrate socio-technical relevance with advanced machine-learning capabilities, enriching DSR knowledge of sustainability analytics. Third, by comparing sentiment dynamics across industries with divergent ESG risk profiles, it enables empirical insights into when public discourse amplifies or undermines symbolic disclosure. This approach addresses the ESG trust deficit by shifting from corporate-controlled reporting to independent stakeholder assessment. Rather than facilitating impression management, our artifact captures authentic public reactions beyond corporate manipulation, serving stakeholders seeking external validation tools. Specifically, investors can apply real-time sentiment analysis for ESG risk screening, regulators for market surveillance and disclosure validation, and civil society organizations for corporate accountability monitoring. The framework thus provides independent, timely insights into public perception, complementing traditional ESG validation methods with external verification.

Literature Review

ESG performance has become the dominant lens for assessing corporate sustainability in strategy and investment decisions (Siegrist et al., 2020), assimilating earlier notions of corporate social responsibility (CSR) and socially responsible investing (SRI). Whereas CSR shifted managerial attention from profit maximization to broader societal obligations (Fatima & Elbanna, 2023), and SRI introduced ethical and sustainability screens into portfolio selection (Martini, 2021; Oehmke & Opp, 2025), contemporary ESG frameworks embed environmental stewardship, labor practices, and governance directly into investment mandates, credit ratings, and regulatory oversight. Despite moves toward harmonized disclosure regimes – exemplified by the Sustainability Accounting Standards Board (SASB) and the European Union's Corporate Sustainability Reporting Directive (CSRD) (European Commission, 2023; García Torea, 2022) – reporting remains largely self-certified and fragmented, inviting accusations of selective disclosure and "greenwashing" (Riyadh et al., 2024). The resulting information asymmetry erodes stakeholder trust and constrains strategic value of sustainability initiatives (Akpan & Oluwagbade, 2023).

Stakeholder theory (Parmar et al., 2010) highlights that firm legitimacy hinges on meeting the expectations of resource-controlling audiences. Favorable ESG perceptions correlate with lower capital costs and stronger customer and employee loyalty (Alsayegh et al., 2020; Choi et al., 2024). Media narratives – and increasingly polyphonic social network discourse –mediate these perceptions by amplifying, reframing, or contesting corporate claims (Zou et al., 2025). For example, consumers and investors are increasingly rewarding businesses perceived as sustainable, and even in emerging markets robust ESG credentials have become critical for attracting investment (Treepongkaruna & Suttipun, 2024). On social media platforms such as Twitter/X and Reddit, high-frequency, user-generated content provides an unfiltered chronicle of public relations that traditional reports such as annual ESG ratings cannot capture (Liu et al., 2023). Social media sentiment also links to financial performance: positive online discussions associate with short-term stock price increases, indicating investors quickly respond to such public narratives (Duz Tan & Tas, 2021).

Advanced natural language processing techniques now make it possible to analyze large social media datasets by processing, analyzing and interpreting streaming information from such platforms, enabling real-time ESG monitoring that can provide real-time feedback to help companies adapt their ESG strategies in response to stakeholder concerns, as opposed to static annual ESG ratings. Park et al. (2023) demonstrate this potential by scraping over 73 million ESG-related posts from Twitter and applying topic modeling and dynamic analysis to reveal evolving discussion themes. However, the methodological challenges are significant: sampling and algorithmic biases, bot manipulation, linguistic nuances, and model drift complicate efforts to extract decision-quality signals from noisy data (Abbasi et al., 2016).

Study	Real-time	Social-media data	Company-specific	Beyond-financial
El Barachi et al. (2021)	✓	✓	✗	✓
Jaiswal et al. (2024)	✗	✓	✓	✗
Kvam et al. (2024)	✗	✓	✓	✓
Fischbach et al. (2023)	✓	✗ (news)	✓	✓
Bariz et al., 2021	✓	✗ (reports)	✓	✗
Li et al. **(2024)**	✗	✗ (reviews)	✓	✓
Leung et al. (2024)	✗	✗ (news)	✓	✓
Present study	✓	✓	✓	✓

Table 1. Comparison of ESG analysis approaches

Recent studies have begun to operationalize computational ESG analytics, but the state of the art remains fragmented, with most failing to provide company-specific insights. Early work relied on traditional media (Fischbach et al., 2023) and annual reports (Bariz et al., 2021), which offered company-level perspectives but lacked the real-time responsiveness needed to capture rapidly evolving sentiment. A case study analyzed social media posts from prominent figures and their followers to examine public reaction to matters such as natural disasters, political developments, and sports events in real time (El Barachi et al., 2021). Li & Zhao (2024) and Leung et al. (2024) assessed non-financial aspects such as employee reviews on Glassdoor but used non-real-time methods. Jaiswal et al. (2024) incorporated real-time analysis of millions of ESG-related social media posts, establishing a framework for sentiment analysis and topic modeling on alternative data, yet without company-specific granularity. Overall, most efforts privilege either speed or depth, rarely both; context-aware firm-level resolution across ESG subdimensions is uncommon, and classifier performance is often unreported. Table 1 maps representative studies against four criteria derived from our review: real-time capability, use of social-media sources, company-specificity, and coverage of non-financial ESG aspects. No extant artifact meets all four simultaneously.

In parallel with analytical advances, the IS community has used DSR to create artifacts addressing ESG and stakeholder engagement challenges. One notable direction is digital platforms for ESG reporting and transparency. For example, Plugge et al. (2024) document a co-design process for an ESG platform ecosystem that reconciles fragmented stakeholders requirements; Hueller et al. (2024) built a platform that automates multi-tier supply chain disclosure from public data; and Vieru & Plugge (2025) proposed an explainable-AI platform for ESG scoring. These prototypes illustrate the synergy between stakeholder engagement and artifact refinement but also reveal limits: evaluations rarely extend beyond expert walkthroughs; sustainability is treated as an application domain rather than a boundary condition shaping design objectives; and few IS studies exploit the affordances of social media datasets that make online discourse a powerful governance resource (Kietzmann et al., 2011; Treem & Leonardi, 2013).

Methodology

This study reports findings from a DSR project that instantiates and evaluates an artifact that transforms noisy, high-frequency Twitter/X discourse into decision-relevant ESG intelligence. Addressing the gap identified in the literature review – rigorously tested, context-aware social media analytics artifacts remain absent – we design, implement, and evaluate a scalable analytics pipeline that classifies public sentiment across all three ESG dimensions and visualizes temporal shifts to alert firms or the public to emerging legitimacy risks or opportunities. Our approach exploits four key social media affordances (Treem & Leonardi, 2013) for ESG validation: (1) *Visibility* – capturing authentic stakeholder reactions in public ESG discussions, unlike filtered corporate surveys; (2) *Persistence* – leveraging permanent digital traces for longitudinal sentiment analysis and reputation tracking; (3) *Editability* – analyzing nuanced expressions of ESG concerns, beyond binary rating systems; and (4) *Association* – examining how ESG narratives spread through networked communities. These affordances enable real-time monitoring, early detection of controversies before escalation, and identification of influential voices shaping ESG perceptions.

The research process follows the echeloned DSR model proposed by Tuunanen et al. (2024), which refines the five-step scheme of Peffers et al. (2007) into five interlocking echelons: problem analysis, objectives and requirements definition, design and development, demonstration, and evaluation. Figure 1 depicts the iterative flow among echelons and how insights in one phase informed adjustments in the next. In problem analysis we conducted expert interviews to surface pain points in current ESG validation; objectives and requirements definition translated these into design objectives. Design and development yielded a modular pipeline coupling domain-adaptive topic modeling with transformer-based sentiment and emotion classifiers, instantiated in Python with a Twitter/X web crawler and a PostgreSQL back-end for streaming storage. Demonstration involved three industry cases – IT consulting, food and beverage, and tobacco company – selected for their contrasting ESG risk profiles. Evaluation combined event-study analysis with a crowdsourced ground-truth benchmark to assess both managerial usefulness (interpretability and timeliness of the alerts) and technical performance.

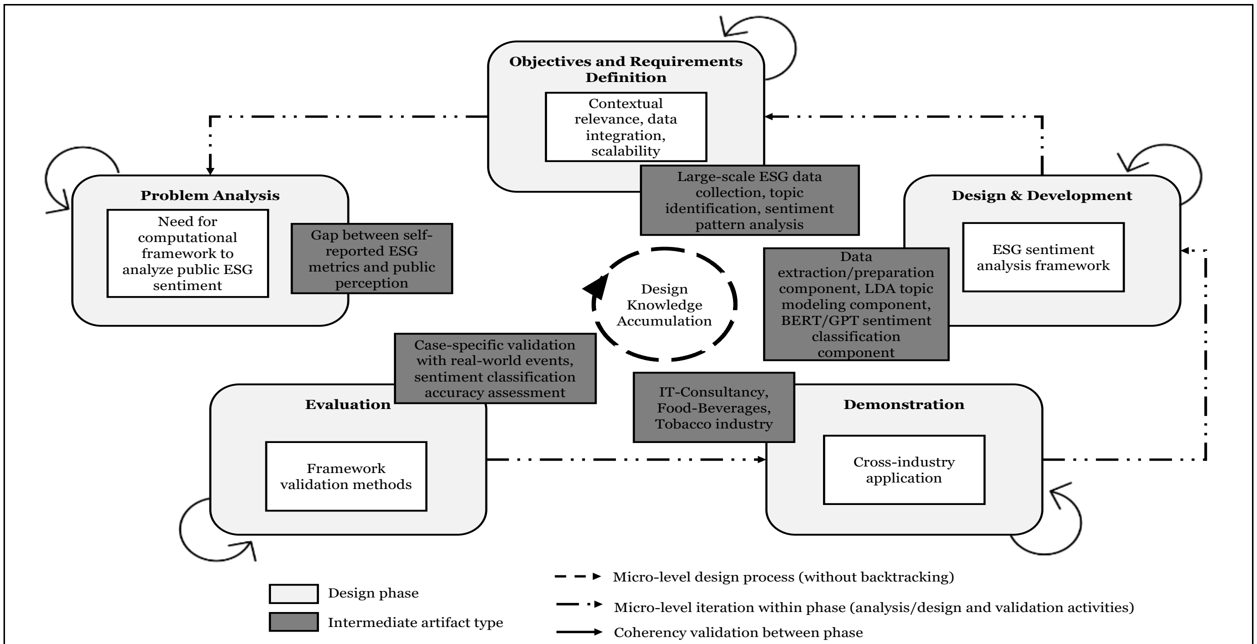

Figure 1. Design science research metamodel for ESG sentiment analysis framework

Integrating advanced natural language processing with DSR's problem-solving paradigm extends analytics literature while providing sustainability-oriented design knowledge that is transferable to other IS contexts. The following section presents the respective methodological rationale and details for each phase, directly linking them to interim results.

Designing a Real-Time ESG Validation Pipeline

Problem Analysis

The problem identification phase of our DSR approach was primarily informed by structured interviews (Table 2) with five ESG experts selected based on their professional experience (minimum five years), diverse industry backgrounds, and direct involvement in ESG implementation or research. These interviews, lasting between 48-60 minutes (average 52 minutes), were recorded and transcribed with participant consent to ensure accurate representation of expert perspectives. Our expert panel included two ESG consultants, one sustainability director from a multinational corporation, one ESG-focused investment analyst, and one academic researcher specializing in corporate sustainability reporting. This diverse composition ensured a comprehensive understanding of ESG validation challenges across different contexts and stakeholder perspectives. The interviews followed a structured format consisting of ten predetermined questions covering the evolution of ESG reporting, current validation practices, and potential improvements.

ID	Role	Industry	Background	Experience	ESG Expertise
I1	Sustainability Director	Technology	Corporate	15+ years	ESG strategy, reporting, and stakeholder engagement
I2	ESG Researcher	Multiple	Academia	8 years	ESG standard development, CSR measurement
I3	ESG Consultant	Finance	Consulting	12 years	Sustainability investing, ESG data validation
I4	CSR Manager	F&B	Corporate	7 years	Supply chain sustainability, environmental initiatives
I5	ESG Analyst	Multiple	Investment	9 years	ESG performance evaluation

Table 2. Overview of interviewees

The experts unanimously identified several critical challenges in current ESG validation approaches. First, they highlighted the inherent *limitations of self-reported data,* with one expert noting: "*The fundamental challenge with ESG reporting today is the lack of independent verification. Companies control both the narrative and the metrics, creating an unavoidable conflict of interest.*" Second, all five experts emphasized the *time lag between ESG reporting and validation,* creating a disconnect between corporate claims and public perception. An ESG investment analyst observed: "*By the time third-party verification is completed, the market has already formed its own opinion based on social media discourse and news coverage. This temporal gap significantly reduces the strategic value of traditional validation.*" Third, the experts consistently pointed to the *fragmentation of ESG standards and frameworks* as a major obstacle to effective validation. A sustainability director with experience across multiple reporting frameworks explained: "*Companies are navigating a complex maze of ESG standards, each with its own metrics and validation approaches. This makes it nearly impossible for stakeholders to make meaningful comparisons across organizations.*"

When asked about potential solutions, all five experts expressed strong interest in leveraging social media data as a complementary validation source. They viewed social media platforms as valuable channels for capturing unfiltered stakeholder sentiment in real-time. One academic researcher specifically noted: "*Social media represents the authentic voice of stakeholders. It contains valuable signals about how ESG initiatives are actually perceived, rather than how companies wish them to be perceived.*" The experts highlighted several key requirements for an effective social media-based validation approach. These included the need for context-specific analysis that considers industry differences, the ability to distinguish between different ESG dimensions (environmental, social, and governance), and integration capabilities with existing validation frameworks. They also emphasized the importance of temporal analysis to track how public sentiment evolves in response to corporate initiatives and external events.

Based on these expert insights, we identified the need for a computational framework that could systematically extract, analyze, and interpret public sentiment towards corporate ESG initiatives using social media data. This framework would need to address the subjectivity, standardization, and temporal limitations of current validation approaches while providing actionable insights for corporate stakeholders. The problem identification thus established both the practical relevance of our research question and the specific requirements that would guide the design of our computational artifact. By grounding our problem understanding in expert perspectives, we ensured that our DSR approach would address genuine challenges faced by ESG practitioners while contributing to the broader discourse on ESG validation methodologies.

Objectives and Requirements Definition

Following Tuunanen et al. (2024) DSR methodology, the objectives definition phase transformed the identified problems into concrete goals for our artifact. This phase established the specific capabilities our computational framework should possess to effectively address the challenges in ESG validation. Our objectives were informed by both the problems identified in the previous phase and the potential of advanced computational techniques to offer innovative solutions. Based on our problem identification and structured interviews with ESG experts, we defined three primary design objectives (DOs) for our computational framework.

First, the framework should enable the *systematic collection and processing of large volumes of ESG-related social media data (DO1)*, addressing the need for more diverse and timely information sources in ESG validation. The experts emphasized that social media represents a rich, largely untapped source of stakeholder perspectives that could complement traditional ESG metrics. One expert particularly noted that "capturing the unfiltered voice of stakeholders through social media could reveal blind spots in corporate ESG assessments that might otherwise remain hidden." Second, our framework should *identify and categorize relevant topics within public ESG discourse (DO2)*, providing structured insights into the specific sustainability issues that resonate with stakeholders. The experts highlighted the importance of understanding not just general sentiment, but the particular ESG dimensions driving public reactions. This objective addressed the problem of limited context in traditional validation methods, which often fail to capture the nuanced, multifaceted nature of public responses to corporate sustainability efforts. Third, the framework should *analyze sentiment patterns within identified ESG topics (DO3)* to provide actionable insights for corporate decision-makers. This objective aimed to address the gap between corporate ESG reporting and public perception, offering organizations a mechanism to gauge the effectiveness of their sustainability initiatives in real-time. The experts emphasized that timely feedback on public sentiment could enable more responsive and effective ESG strategies.

Through this objectives and requirements definition phase, we established a clear vision for what our artifact should accomplish and concrete criteria against which its effectiveness could be evaluated. These objectives directly informed the subsequent design and development phase, ensuring our computational framework would address the specific challenges identified in ESG validation while leveraging the opportunities presented by social media analytics.

Design & Development

Building on the DOs, our design and development phase created a framework to effectively analyze ESG-related social media sentiment. Our implementation resulted in a computational framework with three interconnected components forming a cohesive analytical pipeline as illustrated in Figure 2. This cyclical process begins with data collection through a Twitter web crawler that identifies mentions of company names and ESG keywords. The data then flows through preprocessing steps including text normalization, stop word removal, tokenization, and duplicate removal. Finally, as shown in Figure 2, the framework performs simultaneous sentiment analysis and topic modeling, using LDA topic modeling to identify key ESG thematic themes while applying BERT embeddings with ML models for sentiment classification.

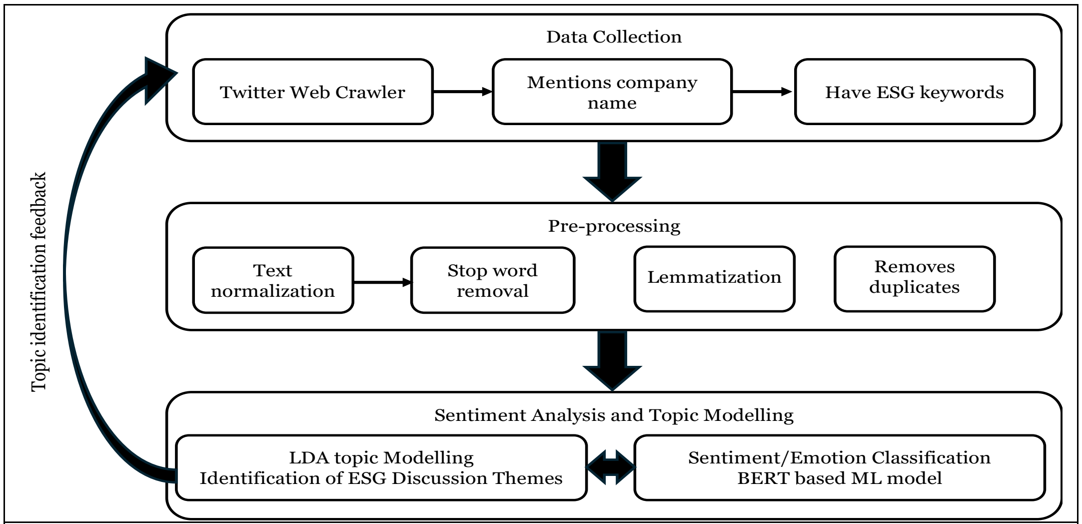

Figure 2. ESG sentiment analysis framework

The *data extraction and preparation* component is the foundation of our analysis process. Our collection method gathers ESG-related tweets based on relevant hashtags, keywords, and company mentions, then filters for content with significant public engagement or specific ESG discussions. The *preprocessing stage*, depicted in the middle section of Figure 2, ensures analytical reliability through systematic data cleansing by addressing missing data in timestamps and content, normalizing text by removing special characters, hashtags, mentions, and URLs, standardizing text to lowercase, removing duplicate tweets, applying tokenization to segment text into individual words, filtering common stop words while preserving ESG-relevant terminology, and implementing lemmatization to reduce words to their root forms.

Our data comes primarily from Twitter/X, chosen for its extensive reach, real-time nature, and diverse content. Twitter's public design enables comprehensive ESG discourse analysis, unlike Facebook's privacy restrictions that limit data accessibility. It also serves as a key venue for immediate reactions to corporate announcements and ESG events – making it well suited for temporal sentiment analysis. The platform hosts active discussions among sustainability professionals, investors, NGOs, and policymakers, providing access to influential stakeholder voices. Its character-limited, text-focused format is also more amendable to sentiment analysis than image-heavy or video-based platforms (e.g., Instagram, TikTok). Prior ESG sentiment research further validates Twitter as a representative source of stakeholder views on corporate sustainability (Jaiswal et al., 2024). While multi-platform analysis could broaden coverage, Twitter's unique characteristics make it the most suitable for rigorous, real-time ESG sentiment validation.

Specifically, we exploit Twitter's affordances: visibility of public discourse enables authentic stakeholder voice capture; persistence supports longitudinal tracking; editability yields nuanced expressions; and association patterns reveal how narratives propagate. Using the focused web crawler shown in the top section of Figure 2, we collected ESG-related tweets through targeted keywords (e.g., "ESG", "sustainability", "corporate responsibility", company mentions). To avoid corporate greenwashing, the data collection model was designed to filter out posts from companies themselves, ensuring authentic external stakeholder and public perspectives. The dataset spans 2021-2023, providing a longitudinal perspective across industries, with each tweet including structured metadata (timestamps, engagement metrics) for quantitative and qualitative analysis. Initial data gathering occurred from January 2021 to December 2023, company-specific filtering in January-February 2024, and crowdsourced validation in March 2024.

We employ Latent Dirichlet Allocation (LDA) topic modeling to uncover thematic structures in ESG-related social media discourse. Using a Bayesian probabilistic approach, each tweet is modeled as a mixture of latent topics, with hyperparameters $\alpha = 0.1$ and $\beta = 0.01$ optimized for topic coherence (NPMI > 0.3). A low α enforces sparse topic mixtures, assuming that individual tweets typically focus on specific ESG aspects rather than multiple themes, while a low β ensures topics are defined by distinctive vocabulary, yielding interpretable and coherent thematic clusters. These parameter choices are designed to produce clearly distinguishable ESG themes for interpretable analysis (Jelodar, et al., 2019). The NPMI coherence threshold of 0.3 provides a standard benchmark for semantic meaningfulness in topic modeling. For implementation, we utilize collapsed Gibbs sampling to infer the posterior distribution of topic assignments. The optimal number of topics is selected by combining perplexity measurements and coherence scores, ensuring interpretable clusters. To visualize and refine the topic models, we use LDAvis, which represents topics as circles sized by prevalence and positioned by semantic similarity. Word relevance is weighted to balance term frequency within a topic against distinctiveness across the corpus. This systematic approach extracts meaningful ESG-related themes from unstructured social media data, forming the basis for subsequent sentiment and emotion analysis.

Our sentiment analysis uses BERT (Bidirectional Encoder Representations from Transformers), a contextual language model capturing bidirectional word relationships. We implemented a fine-tuned BERT-base model (12 transformer encoders, 12 attention heads, 110M parameters) specifically adapted for sentiment classification in ESG-related discourse. Tweets were processed through sequential layers of self-attention mechanisms, enabling it to capture semantic nuances critical for accurate sentiment determination. For preprocessing, we implemented specialized tokenization handling for Twitter-specific elements such as hashtags and mentions, followed by the addition of special tokens ([CLS] and [SEP]) and attention masking to accommodate varied text lengths. The model was fine-tuned on a balanced dataset of manually labeled ESG tweets with cross-validation to ensure robustness. For emotion analysis, we extended beyond the positive-negative-neutral scheme to encompass a finer-grained emotional taxonomy, detecting four distinct emotional states: optimism, joy, anger, and sadness. This granular emotional analysis was implemented through a transfer learning approach, adapting the BERT architecture with a modified classification head trained on emotion-labeled data. Both sentiment and emotion models were evaluated using precision, recall, and F1 scores, with close attention to class-specific performance to address potential imbalances in the distribution of sentiments across ESG topics.

This framework provides organizations with a systematic approach to monitor and analyze public sentiment regarding their ESG initiatives. By combining advanced data collection methods, rigorous preprocessing, sophisticated topic modeling, and nuanced sentiment analysis, the framework enables more informed strategic decision-making regarding ESG policies, communications, and stakeholder engagement.

Demonstration: Results of Case-specific Analysis

To demonstrate the artifact's capabilities, we applied our framework to a dataset of 90,000 tweets reflecting public sentiment on ESG over two years, enabling comprehensive longitudinal sentiment analysis. We focused on three industrial contexts – IT-Consultancy, Food-Beverages, and Tobacco – selecting one major prominent company per industry. From the dataset, we filtered company-specific mentions to create case samples: 2,600 tweets for the IT consultancy company, 9,800 for the food & beverage company, and 590 for the tobacco company. These figures reflect the actual volume of ESG-related discourse available for each firm during our analysis period. The variation in tweet volumes illustrates different levels of public engagement with ESG topics across industries, with the tobacco company attracting fewer but more intense reactions, while the food & beverage company generated the highest volume of ESG-related discussions. This case selection strategy was designed to test the framework's adaptability across varied ESG landscapes with distinct sustainability challenges and public expectations. Our dataset consists of tweets from the general public discussing company ESG practices and behavior. We analyzed external stakeholder tweets about the companies but excluded tweets posted by the companies themselves to ensure our analysis captured authentic public perception rather than corporate messaging.

The data collection process presented several methodological challenges. Managing the high daily volume of tweets required efficient filtering mechanisms to isolate relevant ESG content. We implemented keyword filters based on industry-specific ESG terminology, company names, and hashtags to ensure data relevance. Processing occurred in sequential stages, incrementally refining the dataset to focus on substantive ESG discussions while minimizing noise. For each industry, we applied topic modeling to identify key themes, followed by sentiment classification and temporal analysis to track shifts in public perception over time. This analytical approach tested the framework's capacity to detect industry-specific ESG discourse patterns and sentiment dynamics. The following section presents case findings (see also Figure 3), highlighting variations in sentiment patterns, emotional responses, and thematic emphases across sectors.

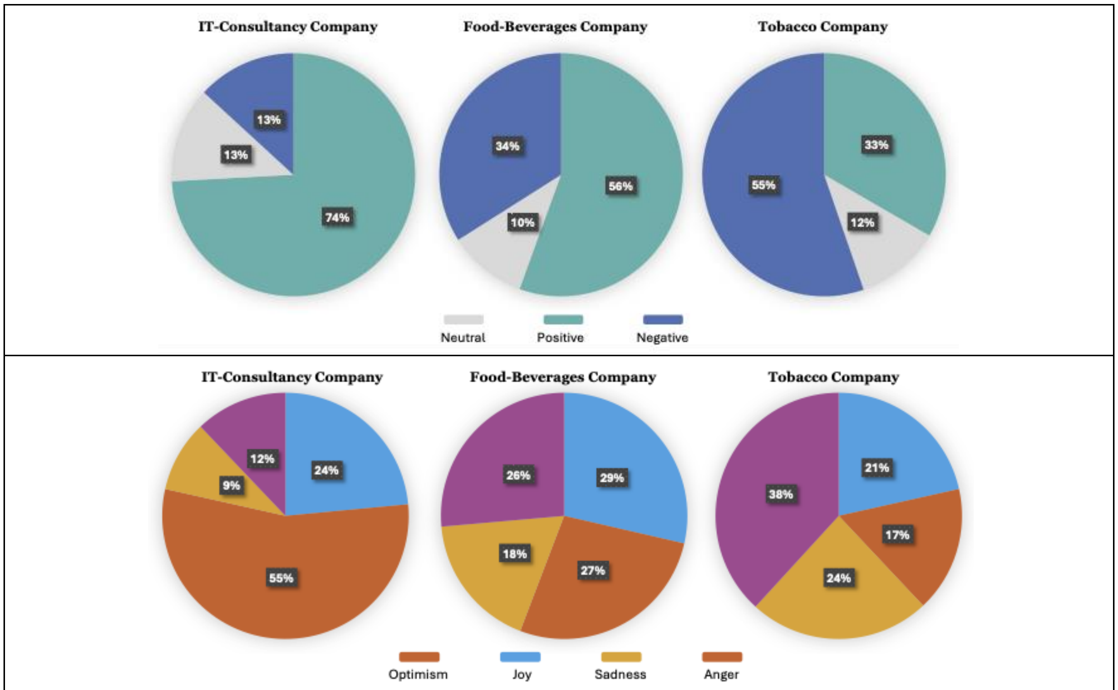

Figure 3. Sentiment and emotions identified in all three industries

Case 1: IT-Consulting. Analysis of 2,600 tweets about this company revealed strong engagement with technology-enabled ESG progress. *Topic modeling* surfaced the following five most prominent ESG themes: Climate action and net-zero roadmaps (e.g., "net zero", "2040", "renewables", "emissions-cut") generated predominantly positive sentiment. Responsible supply chain management (e.g., "ecovadis", "RBA", "labor standards") received mixed reactions. Inclusion, diversity, and equality initiatives ("inclusion networks", "ID&E", "belonging") formed a largely positive sentiment cluster. Workforce development ("upskill", "cloud certs", "AI-training") reflecting favorable views of talent development efforts. By contrast, workforce reductions ("head-count", "restructuring", "rolling-layoffs") produced a distinct negative cluster. These topics, selected from the broader modeling output, capture the most coherent and representative ESG discourse in this industry.

Sentiment analysis (Figure 3) revealed 74% positive sentiment, mainly tied to progress in climate action roadmaps, inclusion, diversity & equality, and AI-driven upskilling solutions. Neutral sentiment (13%) correlated with general discussions about partnerships and corporate updates, while negative sentiment (13%) reflected critiques of layoffs. *Emotion analysis* showed optimism (55%) and joy (23.5%) as predominant emotional responses, with sadness (9.4%) and anger (12.1%) emerging around unmet sustainability targets, specifically regarding carbon neutrality commitments. This emotional landscape underscores the industry's complex relationship with sustainability expectations, balancing technological innovation against implementation challenges. This sector's pivotal role in facilitating digital sustainability innovations was reflected in its predominantly positive sentiment patterns.

Case 2: Food-Beverages Industry. From 9800 tweets about this company, *topic modeling* surfaced numerous ESG themes, with five standing out for prominence and sentiment significance. Net-zero targets and regenerative agriculture (e.g., "net zero", "2050", "regenerative", "cocoa-plan") drew largely positive sentiment. Plastic pollution formed the second major topic, dominated by terms such as "top plastic polluter", "sachets", and "BreakFreeFromPlastic". Water rights controversies (e.g., "water", "san bernardino forest", "drought") generated predominantly negative reactions. Child labor and human rights concerns in supply chains created a fourth cluster with mixed but mostly negative sentiment, while health and nutrition reformulation ("sugar-reduction", "portion-control", "healthier-portfolio") completed the top five themes. These themes represented the most influential and distinct clusters for this sector.

Compared to the IT-Consultancy sector, the *sentiment analysis* showed a more polarized sentiment distribution (Figure 3). Positive sentiment (55.9%) correlated with regenerative agriculture initiatives and nutrition reformulation. Negative sentiment (35.0%) centered on child labor, human-rights, water extraction practices and their community impact, and perceived delays in sustainability commitment fulfillment. Neutral sentiment (9.1%) reflected informational content such as press releases. *Emotional analysis* (Figure 3) revealed optimism and joy in response to sustainability campaigns, particularly packaging innovations, whereas anger and sadness dominated reactions to environmental criticisms, especially allegations of resource exploitation in water-scarce regions. Public discourse conveyed a distinct sense of urgency, calling for faster transition to sustainable operational practices.

Case 3: Tobacco Industry. From 590 tweets, *topic modeling* revealed multiple ESG themes amid predominantly critical public discourse, with five standing out. The industry's "smoke-free" transformation strategy (e.g., "iqos", "zyn", "unsmoke") drew mixed sentiment with some cautious optimism. ESG ratings and DJSI inclusion formed a second topic cluster with rare positive sentiment. Health and youth-vaping regulation ("e-cigarette", "youth", "fda", "harm-reduction") attracted strongly negative reactions. Climate and environmental pledges ("net zero", "cdp", "a-list") generated mixed but somewhat positive sentiment. Finally, supply-chain labor and greenwashing allegations, ("child labor", "malawi", "human-rights", "greenwashing") dominated critical discussions.

Sentiment analysis (Figure 3) revealed predominantly negative sentiment (55%), with discourse focused on health risks, environmental impacts, and greenwashing allegations. Positive sentiment (33%) correlated primarily with corporate philanthropy, such as public health program funding and carbon emission reduction commitments, while neutral sentiment (12%) reflected general discussions of corporate ESG statements and initiatives. *Emotion analysis* (Figure 3) showed anger and sadness as dominant emotional responses, highlighting widespread criticism of the industry's perceived inability to reconcile its fundamental business model with ESG principles. Joy and optimism, although less frequent, appeared around specific environmental initiatives like renewable energy adoption and supply chain waste reduction.

Evaluation

We used two complementary approaches to evaluate the artifact: real-world event validation and human crowdsourcing validation allowed us to assess both the framework's temporal sensitivity to external ESG events and its alignment with human judgment.

Real-World Event Validation

Our *event-based validation* focused on case-specific validation through real-world events, examining the correlation between sentiment trends identified by our framework and documented ESG-related events within each industry. By mapping significant sentiment fluctuations against known corporate announcements, controversies, and initiatives, we assessed the framework's sensitivity to external factors that influence public opinion. For each industry, we identified key ESG-related events during our analysis period and examined sentiment patterns before, during, and after these events.

In the IT-Consultancy sector, we tracked sentiment shifts following major sustainability announcements, technology launches with environmental benefits, and corporate governance changes. For the Food-Beverages industry, we analyzed sentiment responses to packaging sustainability initiatives, supply chain modifications, and community engagement programs. In the Tobacco sector, we monitored public reaction to regulatory developments, harm reduction product launches, and environmental sustainability commitments. This event-based validation approach enabled us to assess the framework's temporal accuracy, contextual sensitivity, and industry-specific relevance. We examined various performance metrics including topic classification accuracy, sentiment distribution alignment with known events, and the framework's ability to detect subtle shifts in public opinion. The evaluation also considered technical performance aspects such as processing efficiency, scalability with increasing data volumes, and the system's adaptability to evolving ESG terminology and emerging topics. Through this, we established robust evidence of the artifact's provides actionable insights for organizations seeking to monitor and respond to stakeholder perceptions.

Crowdsourcing Validation

Our second evaluation approach incorporated human judgment through crowdsourcing to establish a standard for sentiment classification. After automated classification using the sentiment analysis, we created a human-labeled set for our IT-Consultancy corpus of 450 tweets. Three annotators were recruited specifically for their prior exposure to ESG reporting and social-media analytics: an ESG Consultant with MBA and 6 years of experience in ESG strategy development, a Manager with PhD and 8 years of experience in ESG standard development and CSR measurement, and a researcher in Computer Science with 5 years of experience in Natural Language Processing and ESG research focusing on computational text analysis. Before labeling, they completed a 10-minute micro-training module that (i) introduced ESG constructs, (ii) explained our three-point sentiment scale (0 = negative, 1 = neutral, 2 = positive), and (iii) walked through edge-case examples such as sarcasm and mixed-sentiment clauses. Annotation was performed in a internal web interface that displayed one tweet at a time; raters were required to select a sentiment. Once all three labels were collected, a simple majority-vote script resolved the final class; in the 8% of cases with a three-way split, the item was routed to a internal ESG researcher for adjudication within the same interface. The resulting adjudicated dataset became the reference benchmark for evaluating our BERT model. Through this comprehensive evaluation strategy combining automated analysis with human validation, we established robust evidence of the artifact's effectiveness in analyzing ESG-related public sentiment across diverse industrial contexts, providing actionable intelligence for organizations seeking to monitor and respond to stakeholder perceptions.

The crowdsourced evaluation contributed to our findings in several key ways. First, it revealed systematic patterns in classification disagreements, particularly for tweets containing technical ESG terminology, sarcasm, or mixed sentiments. We reported agreement levels specifically for negative tweets (83% agreement) because negative sentiment classification showed the highest accuracy and reliability, making it a robust benchmark for our model's performance. Positive and neutral classifications showed lower agreement rates (62% and 54% respectively) due to the subjective nature of interpreting corporate ESG claims. These crowdsourcing results directly informed our sentiment classification algorithm design through: (1) Implementation of confidence thresholding to flag uncertain classifications, (2) Development of industry-specific lexicons based on disagreement patterns, (3) Addition of context-aware preprocessing steps to handle ESG-specific terminology, and (4) Integration of ensemble methods that combine multiple classification approaches for ambiguous cases identified through human evaluation.

EE1: Real time ESG announcement analysis. The first evaluation episode focused on assessing our framework's ability to detect and analyze public sentiment shifts in response to significant ESG-related corporate announcements. This evaluation approach tested the temporal sensitivity and contextual accuracy of our *sentiment analysis framework* by examining how effectively it captured real-time public reactions to ESG initiatives across the three industry sectors. For our real-time ESG announcement analysis, we focused again on the same major companies per industry as before (IT consultancy, food and beverage, and tobacco). For each company, we identified major ESG-related announcements during our analysis period and examined sentiment patterns before, during, and after these events, as depicted in Figure 4. This single-company-per-industry approach allowed for in-depth analysis of sentiment dynamics while ensuring sufficient tweet volume for meaningful temporal analysis.

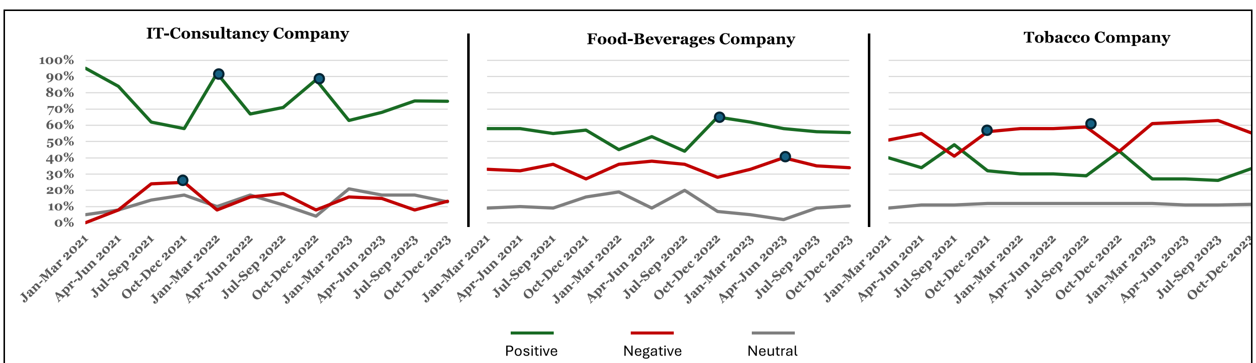

Figure 4. Sentiments timeline across all three industries

For IT-Consultancy company, we analyzed sentiment shifts (Figure 4) around three significant peaks of Twitter activity that occurred in March 2022, and November 2022. Our timeline analysis revealed that these peaks were predominantly associated with positive sentiment, particularly related to digital sustainability initiatives and green IT announcements. On the social front, the company launched over 70 inclusion networks, and this peak of positive sentiment could be seen in December 2022. One area of mild criticism is related to business performance and restructuring – for instance, employee layoffs in January 2022 as part of cost-cutting drew some negative tweets. The framework successfully detected a notable increase in optimism-related emotions following companies' sustainability announcements, demonstrating its ability to capture nuanced emotional responses beyond basic sentiment polarity.

In the case of Food and Breverage Industry (Figure 4), our framework identified a major shift in Twitter activity, with two distinct sentiment patterns. The first peak showed a shift toward more positive sentiment in November 2022, corresponding with the company's announcements of sustainable packaging initiatives. The second peak was in June 2023 characterized by predominantly negative sentiment, coinciding with public criticism of the company's water usage practices. People mentioned "100% recyclable" claims in August 2023 further painted company's messaging as deceptive. This temporal pattern demonstrated our framework's ability to detect sentiment shifts in response to different types of ESG-related events.

For Tobacco (Figure 4), our analysis tracked sentiment around a significant peak in Twitter activity between November 2021 to August 2022, which was predominantly characterized by negative sentiment. This aligned with public criticism of the company's harm reduction claims, highlighting the framework's capability to identify industry-specific ESG challenges. Interestingly, our analysis also detected that in 2022, a WHO report explicitly accused major tobacco companies (our targeted company included) of "greenwashing": showcasing environmental projects and ESG accolades to rehabilitate their image, while continuing to profit from a product that kills millions. This longitudinal perspective demonstrated the framework's ability to capture sentiment evolution over extended periods.

Across all three cases, our evaluation confirmed the framework's effectiveness in detecting sentiment shifts aligned with real-world ESG announcements and events. The sentiment patterns identified by our computational analysis showed strong temporal correlation with documented corporate ESG initiatives and public controversies. For instance, in the case of Food and Breverage Industry, the shift from anger to joy in the two 2023 peaks accurately reflected the company's response to water usage criticism through the announcement of new sustainability commitments. This evaluation episode also provided insights into the varying public responses to ESG announcements across different industries. The predominantly positive response to IT sector sustainability initiatives contrasted sharply with the more skeptical reactions to similar announcements in the tobacco industry, highlighting the importance of industry context in ESG sentiment analysis. Overall, this evaluation episode demonstrated that our framework successfully captures real-time public sentiment shifts in response to ESG announcements, providing valuable temporal insights that can help organizations, investors public can gauge the impact of their ESG initiatives.

EE 2: crowd sourcing. The second evaluation episode involved crowdsourcing to validate our sentiment analysis model's performance. The crowdsourced classifications revealed a sentiment distribution of 59% positive, 24% negative, and 17% neutral tweets in the IT consultancy industry dataset. When compared against our automated classification methods, we found that the BERT model achieved 62% agreement with human classifications. The BERT method demonstrated greater accuracy in identifying negative sentiment (83% agreement with human classifications) but tended to overclassify neutral content.

Analysis of disagreement patterns between human evaluators and automated methods revealed that Tweets containing technical terminology, sarcasm, or mixed sentiments within the same message presented the greatest classification difficulties. Additionally, referencing multiple ESG dimensions simultaneously (such as environmental initiatives alongside governance concerns) often received different classifications depending on which aspect the evaluator or model prioritized. This crowdsourcing evaluation episode provided valuable insights into the strengths and limitations of our approaches. It highlighted the importance of incorporating human judgment in the evaluation process, particularly for content where contextual understanding and domain knowledge significantly impact sentiment interpretation.

Discussion

The purpose of this study was to investigate whether social media discourse can complement – or even partially replace – traditional, firm-controlled ESG assurance mechanisms. By designing an artifact that captures Twitter/X conversations, distills domain-specific topics, and assigns sentiment labels, we sought to operationalize the call for "ambient" stakeholder feedback loops regarding ESG performance, which is gaining momentum in both governance (European Commission, 2023) and IS research (Schoormann et al., 2025). The following sections outline the implications of our findings, points out under which conditions real-time sentiment analysis adds value, and outlines avenues for future research.

Real-Time Sentiment as a Legitimacy Barometer

The artifact's cross-industry demonstration and evaluation confirms that Twitter/X acts as a sensitive – though not uniformly receptive – barometer of stakeholder legitimacy judgements. Our temporal analysis of sentiment revealed that public perception is highly responsive to corporate announcements and controversies, suggesting that social media sentiment serves as a real-time feedback mechanism for ESG initiatives. Specific value of real-time analysis includes: (1) Early warning detection – in our food & beverage case, negative sentiment about water usage practices emerged weeks before formal legal challenges, providing early controversy signals that traditional annual ESG reports would miss; (2) Immediate campaign effectiveness measurement – the IT consultancy's sustainability announcements showed sentiment improvements within hours, enabling rapid assessment of communication strategy success; (3) Optimal timing identification – real-time monitoring revealed that ESG announcements during high-engagement periods generated substantially more positive sentiment than routine disclosure timing; (4) Competitive benchmarking – simultaneous tracking across industries showed that tobacco industry ESG initiatives consistently faced predominantly negative sentiment regardless of content, while IT sector initiatives maintained strong positive reception. This aligns with Liu et al. (2023), who noted that social media offers an "unfiltered and immediate view of stakeholder sentiment that traditional reports cannot capture." The peaks in Twitter/X activity for each company corresponded with significant corporate announcements or controversies, demonstrating the platform's value as a barometer for public reaction.

For the IT company that positions itself as an enabler of digital sustainability, 78.5% of tweets were positive, and emotion analysis was dominated by optimism and joy. Peaks in those emotions coincided with announcements of climate-action roadmaps and inclusion programs, mirroring the event-based validation that mapped three sentiment spikes (May 2021, December 2021, and April 2022) to green-IT launches and governance reforms. By contrast, the food-and-beverage company faced a polarized landscape: 55.8% positive tweets on regenerative agriculture and nutrition reformulation were offset by 35.0% negative tweets focused on water extraction and child-labor allegations. Tobacco discourse was more severe, with roughly 62% negative sentiment and anger/sadness dominating despite repeated "smoke-free" pledges.

Taken together, these patterns resonate with legitimacy theory and particularly legitimacy-threshold arguments (Soublière & Gehman, 2020; Suchman, 1995). Stakeholders appear willing to reward firms whose core business model is compatible with sustainability narratives, yet remain skeptical when the underlying product (e.g., combustible tobacco) conflicts with ESG goals. For IS scholars, the evidence implies that digital trace data (here: social media data) can reveal sector-specific credibility ceilings that conventional ESG ratings, which aggregate scores across industries, may obscure. More granular emotional analysis deepens this insight. In food-and-beverage, anger clusters around hashtags such as #BreakFreeFromPlastic or #WaterRights, whereas joy emerges when packaging innovations are announced. Emotion therefore signals whether discourse is moving toward moral condemnation or constructive engagement and thus provides additional contextual understanding of stakeholder engagement with corporate sustainability efforts, supporting Zou et al. (2025) assertion that "media narratives serve as a powerful intermediary that can amplify or downplay a company's ESG actions."

While directional sentiment changes may appear intuitive, our analysis revealed several unexpected findings that challenge conventional ESG assumptions. First, we discovered industry-specific legitimacy thresholds where tobacco companies face mostly negative sentiment regardless of ESG initiative quality, while IT consultancy maintains mostly positive sentiment, revealing sector-specific credibility ceilings that traditional ESG ratings obscure. Second, our temporal analysis showed that ESG sentiment shifts occur within hours rather than the weeks assumed in traditional validation cycles. Third, emotional granularity provided nuanced insights beyond basic sentiment – anger clustered around specific issues like plastic pollution (#BreakFreeFromPlastic) while joy emerged from packaging innovations, enabling targeted strategic responses. Finally, cross-industry comparison revealed that identical ESG initiatives (net-zero pledges) generate completely different stakeholder reactions depending on industry context, challenging assumptions about universal ESG strategy transferability.

Methodological Additions to ESG Text Analytics and Design Science

Our project advances ESG text analytics by combining domain-adaptive topic discovery with transformer-based sentiment classification – an architecture that ensures both contextual richness and computational tractability. Topic modeling first surfaced themes such as “net-zero roadmaps” in IT or “plastic-pollution backlash” in food-and-beverage, after which sentiment models were constrained to those themes. Although we did not compute full precision-recall matrices, the crowdsourcing benchmark on 450 IT-industry tweets demonstrates the pipeline’s overall practical effectiveness. Disagreement concentrated in tweets containing sarcasm, mixed valence, or references to multiple ESG dimensions, echoing Van Atteveldt et al. (2021) conclusions about the difficulties of computational approaches in sentiment analysis when dealing with co-occurring concepts and their critique of simplified polarity scales. Future iterations could implement probabilistic sentiment scores or multi-label outputs to flag classification uncertainty in real time.

DSR literature benefits from sustainability-oriented design knowledge in the form of a modular analytics pipeline that couples domain-adaptive topic modeling (Jelodar et al., 2019) with context-aware sentiment and emotion classification based on machine learning and transformer architectures (Bello et al., 2023; Khurana et al., 2023) that is transferable to other IS contexts in the sustainability domain. For example, similar analytical frameworks could be deployed for real-time monitoring of supply chain sustainability disclosures, tracking stakeholder responses to climate adaptation measures, or evaluating the perceived authenticity of corporate diversity and inclusion initiatives – all areas where ambiguous terminology and complex sentiment patterns similarly challenge conventional measurement approaches.

Implications for Managers, Auditors, and Regulators

For *managers*, our research offers a methodological blueprint for incorporating social media sentiment analysis into ESG reporting and validation processes. As noted by Van Atteveldt et al. (2021), traditional methods of collecting and validating ESG data have been criticized for being time-consuming, costly, and subject to accuracy and credibility issues. Our artifact addresses these limitations through automated, real-time analysis, providing firms with a systematic approach to understanding public perception of their ESG initiatives, acting as what Ignatov (2023) describes as an “early warning system” or potential controversy. In food-and-beverage, negative sentiment about water extraction spiked weeks before a lawsuit surfaced, giving management a window to engage community stakeholders. In IT consulting, a drop in sentiment around layoffs in January 2022 signaled reputational damage that dissemination of positive sustainability content alone could not neutralize. These use cases parallel Wade & Hulland (2004) resource-based view: analytical capabilities deliver value only when matched with absorptive capacity, i.e., teams capable of interpreting and acting on those insights. Interviewed sustainability staff highlighted a skills gap in data literacy, suggesting that training investments may be a prerequisite for harvesting social media intelligence. *Regulators* and *auditors* can likewise leverage the artifact. The U.S. Securities and Exchange Commission’s proposed climate-disclosure rule (SEC, 2024) foresees third-party assurance. Sentiment heat maps could guide auditors toward contested disclosure points, increasing the efficiency of limited assurance budgets. Within the European Union, the draft CSRD (European Commision, 2023) might integrate social-media indicators into its risk-screening toolkit, bridging reported compliance and perceived legitimacy.

Limitations & Future Work

The study has several limitations that open avenues for future research. First, Twitter/X’s demographic skews toward younger, urban, wealthier users may bias sentiment estimates (Blank & Lutz, 2017); additional platforms such as LinkedIn for professional voices or reddit for Generation Z perspectives would eventually improve representativeness. Second, the artifact was evaluated on language tweets; multi-modality analysis such as incorporating images would give more contextual knowledge. Third, the evaluation covered three industries; extending to low-salience sectors (e.g., utilities) could test whether sentiment dynamics differ under weaker public scrutiny. Fourth, our human-in-the-loop crowdsourcing evaluation relied on three annotators; future work should employ larger coder pools to derive robust inter-rater reliability estimates and potentially adopt active-learning protocols that prioritize uncertain cases (Settles, 2009). Finally, while temporal alignment with events was demonstrated qualitatively, causal directionality remains to be established through quasi-experimental designs, such as difference-in-differences natural experiments (Angrist & Pischke, 2014).

A second set of limitations concerns model performance. Our BERT-based approach tended to classify content as neutral, suggesting several directions for improvement. Finally, our BERT methodology presents several limitations requiring future attention. The model displayed a tendency to classify content as neutral and showed declining performance when applied across different industries (IT-Consultancy: F1=0.72, Food-Beverages: F1=0.64, Tobacco: F1=0.58) due to domain-specific sentiment patterns, confirming cross-industry generalizability challenges from training primarily on IT-consultancy data. Fine-tuning BERT specifically on ESG terminology could enhance performance by better recognizing industry-specific sentiment indicators. Incorporating domain experts in the validation process would provide richer ground truth data, while adapting the analysis framework to account for industry-specific ESG priorities would enable more precise sentiment evaluation tailored to company-specific contexts.

Ethical risks should also be foregrounded. Although we stripped personal identifiers and complied with platform policies, sentiment monitoring can entrench surveillance norms and reinforce algorithmic biases documented in natural language processing models (Bender et al., 2021). Participatory audits and differential-privacy techniques offer routes for mitigation. The study also prompts reflection on the scope of ESG assurance. Traditional frameworks, such as the Global Reporting Initiative, emphasize backward-looking metrics audited annually. Our findings suggest that forward-looking, sentiment-based indicators can serve as leading signals of latent reputational risk. Integrating such indicators into balanced-scorecard systems (Kaplan & Norton, 1996) could enable dynamic reallocation of sustainability budgets toward issues that matter most to vocal stakeholders, thereby enhancing strategic agility (Teece, 2018).

Future work should broaden the scope to include a wider range of companies, industries, and social media platforms to assess the robustness of sentiment dynamics across contexts, as suggested by Bouadjenek et al. (2023). Longitudinal studies that track the co-evolution of ESG reporting and public sentiment would provide valuable insights into how disclosure practices and stakeholder reactions interact over time, advancing our understanding of corporate sustainability communication as a dynamic process.

Acknowledgment of AI Use

The authors acknowledge the use of generative AI tools, specifically GPT-4, 4o, 5 (OpenAI) and Claude 3.5 Sonnet (Anthropic), for editing and writing assistance during manuscript preparation. All research design, methodology, data analysis, findings, and intellectual contributions represent the authors' original work.

References

Abbasi, A., Sarker, S., & Chiang, R. (2016). Big data research in information systems: Toward an inclusive research agenda. *Journal of the association for information systems*, 17(2), i-xxxii. https://doi.org/10.17705/1jais.00423

Akpan, J. U., & Oluwagbade, O. (2023). Social and environmental responsibility in accounting: Beyond financial metrics. *International Journal of social sciences and management research*, 9(9), 163-188. https://doi.org/10.56201/ijssmr.v9.no9.2023.pg163.188

Alavi, M., & Leidner, D. E. (2001). Knowledge management and knowledge management systems: Conceptual foundations and research issues. *MIS quarterly*, 25(1), 107-136. https://doi.org/10.2307/3250961

Alsayegh, M. F., Abdul Rahman, R., & Homayoun, S. (2020). Corporate economic, environmental, and social sustainability performance transformation through ESG disclosure. *Sustainability*, 12(9), 3910. https://doi.org/10.3390/su12093910

Bariz, H., Kiesel, F., & Schreiber, N. (2021). The Information content of ESG disclosures and the public perception of credit risk. *ICIS 2021*. https://aisel.aisnet.org/icis2021/is_sustain/is_sustain/10

Bello, A., Ng, S. C., & Leung, M. F. (2023). A BERT framework to sentiment analysis of tweets. *Sensors*, 23(1), 506. https://doi.org/10.3390/s23010506

Bharadwaj, A. E. (2013). Digital business strategy: Toward a next generation of insights. *MIS quarterly*, 37(2), 471-482. https://doi.org/10.25300/MISQ/2013/37:2.3

Bouadjenek, M. R., Sanner, S., & Wu, G. (2023). A user-centric analysis of social media for stock market prediction. *ACM transactions on the web*, 1-22. https://doi.org/10.1145/3532856

Choi, S., Jeong, K. S., & Park, S. R. (2024). ESG activity recognition enhances organizational commitment and service-oriented organizational citizenship behavior among insurance call center staff. *Heliyon*, 10(11). https://doi.org/10.1016/j.heliyon.2024.e31999

Chopra, S. S., Senadheera, S. S., Dissanayake, P. D., Withana, P. A., Chib, R., Rhee, J. H., & Ok, Y. S. (2024). Navigating the challenges of environmental, social, and governance (ESG) reporting: The path to broader sustainable development. *Sustainability*, 16(2), 606. https://doi.org/10.3390/su16020606

Delmas, M. A., & Burbano, V. C. (2011). The drivers of greenwashing. *California management review*, 54(3), 64-87. https://doi.org/10.1525/cmr.2011.54.1.64

Duz Tan, S., & Tas, O. (2021). Social media sentiment in international stock returns and trading activity. *Journal of behavioral finance*, 22(3), 221-234. https://doi.org/10.1080/15427560.2020.1772261

El Barachi, M., AlKhatib, M., Mathew, S., & Oroumchian, F. (2021). A novel sentiment analysis framework for monitoring the evolving public opinion in real-time: Case study on climate change. *Journal of cleaner production*, 312, 127820. https://doi.org/10.1016/j.jclepro.2021.127820

European Commission (2023). Corporate sustainability reporting directive. *http://data.europa.eu/eli/reg_del/2023/2772/oj*

Fatima, T., & Elbanna, S. (2023). Corporate social responsibility (CSR) implementation: A review and a research agenda towards an integrative framework. *Journal of business ethics*, 184(1), 105-121. https://doi.org/10.1007/s10551-022-05047-8

Fischbach, J., Adam, M., Dzhagatspanyan, V., Mendez, D., Frattini, J., Kosenkov, O., & Elahidoost, P. (2023). Automatic eSG assessment of companies by mining and evaluating media coverage data: NLP approach and tool. *IEEE Big Data 2025*. https://doi.org/10.1109/BigData59044.2023.10386488

García Torea, N. (2022). Sustainability accounting standards board (SASB). *In Encyclopedia of sustainable management*, 1-3. https://doi.org/10.1007/978-3-030-02006-4_841-1

Hueller, L., Kuffner, T., Schneider, M., Schuhmann, L., Cauderay, V., Buz, T., ..., Uebernickel, F. (2024). Designing a collaborative platform for advancing supply chain transparency. *arXiv preprint*, https://doi.org/10.48550/arXiv.2409.08104

Ignatov, K. (2023). When ESG talks: ESG tone of 10-K reports and its significance to stock markets. *International review of financial analysis*, 89, 102745. https://doi.org/10.1016/j.irfa.2023.102745

Jaiswal, R., Gupta, S., & Tiwari, A. K. (2024). Decoding mood of the Twitterverse on ESG investing: Opinion mining and key themes using machine learning. *Management research review*, 47(5), 1221-1252. https://doi.org/10.1108/MRR-07-2023-0526

Jawahar, G., Sagot, B., & Seddah, D. (2019). What does BERT learn about the structure of language? *ACL 2019*. https://inria.hal.science/hal-02131630v1

Jelodar, H., Wang, Y., Yuan, C., Feng, X., Jiang, X., Li, Y., & Zhao, L. (2019). Latent dirichlet allocation (LDA) and topic modeling: Models, applications, a survey. *Multimedia tools and applications*, 78(11), 15169-15211. https://doi.org/10.1007/s11042-018-6894-4

Johnson, Z., Ashoori, M. T., & Lee, Y. J. (2018). Self-reporting CSR activities: When your company harms, do you self-disclose? *Corporate reputation review*, 21, 153-164. https://doi.org/10.1057/s41299-018-0051-x

Khurana, D., Koli, A., Khatter, K., & Singh, S. (2023). Natural language processing: State of the art, current trends and challenges. *Multimedia tools and applications*, 82, 3713-3744. https://doi.org/10.1007/s11042-022-13428-4

Kietzmann, J. H., Hermkens, K., McCarthy, I. P., & Silvestre, B. S. (2011). Social media? Get serious! Understanding the functional building blocks of social media. *Business horizons*, 54(3), 241-251. https://doi.org/10.1016/j.bushor.2011.01.005

Kvam, E., Molnár, P., Wankel, I., & Ødegaard, B. A. (2024). Do stock prices react to ESG sentiment? *SSRN preprint*, 4057988. https://dx.doi.org/10.2139/ssrn.4057988

Leung, M., Liang, C., Lourie, B., & Zhu, C. (2024). Voting with their feet: Employee responses to ESG news. *SSRN preprint*, 4937880. https://dx.doi.org/10.2139/ssrn.4937880

Li, Z., & Zhao, X. (2024). Text mining analysis of ESG practices in employee online reviews for organizational insight. *PACIS 2024*. https://aisel.aisnet.org/pacis2024/track17_socmedia/track17_socmedia/5/

Liu, M., Luo, X., & Lu, W. Z. (2023). Public perceptions of environmental, social, and governance (ESG) based on social media data: Evidence from China. *Journal of cleaner production*, 387, 135840. https://doi.org/10.1016/j.jclepro.2022.135840

Martini, A. (2021). Socially responsible investing: From the ethical origins to the sustainable development framework of the European Union. *Environment, development and sustainability*, 23, 16874-16890. https://doi.org/10.1007/s10668-021-01375-3

Massa, S., Annosi, M. C., Marchegiani, L., & Petruzzelli, A. M. (2023). Digital technologies and knowledge processes: New emerging strategies in international business. A systematic literature review. *Journal of knowledge management*, 27 (11), 330–387. https://doi.org/10.1108/JKM-12-2022-0993

Oehmke, M., & Opp, M. M. (2025). A theory of socially responsible investment. *Review of Economic Studies*, 92(2), 1193-1225. https://doi.org/10.1093/restud/rdae048

Park, J., Choi, W., & Jung, S. U. (2022). Exploring trends in environmental, social, and governance themes and their sentimental value over time. *Frontiers in psychology*, 13, 890435. https://doi.org/10.3389/fpsyg.2022.890435

Parmar, B. L., Freeman, R. E., Harrison, J. S., Wicks, A. C., Purnell, L., & De Colle, S. (2010). Stakeholder theory: The state of the art. *The academy of management annals*, 4(1), 403-445. https://doi.org/10.5465/19416520.2010.495581

Peffers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of management information systems*, 24(3), 45-77. https://doi.org/10.2753/MIS0742-1222240302

Plugge, A., De Reuver, M., & Vieru, D. (2024). Co-designing ESG platform functionality: A digital platform ecosystem approach. *HICSS 2024*. https://hdl.handle.net/10125/107190

Riyadh, H. A., Garad, A., & Al-Shmam, M. A. (2024). A critical overview of “the good,”“the bad,” and “the ugly” in CSR contemporary business landscapes. *Impact of corporate social responsibility on employee wellbeing*. https://doi.org/10.4018/979-8-3693-3470-6.ch001

Rodríguez-Ibánez, M., Casánez-Ventura, A., Castejón-Mateos, F., & Cuenca-Jiménez, P. M. (2023). A review on sentiment analysis from social media platforms. *Expert systems with applications*, 223, 119862. https://doi.org/10.1016/j.eswa.2023.119862

Schoormann, T., Möller, F., Hoppe, C., & vom Brocke, J. (2025). Digital sustainability. *Business & information systems engineering*, 67(3), 429–438. https://doi.org/10.1007/s12599-025-00937-3

SEC (2024). SEC adopts rules to enhance and standardize climate-related disclosures for investors. https://www.sec.gov/newsroom/press-releases/2024-31

Siegrist, M., Bowman, G., Mervine, E., & Southam, C. (2020). Embedding environment and sustainability into corporate financial decision-making. *Accounting & finance*, 60(1), 129-147. https://doi.org/10.1111/acfi.12533

Treem, J. W., & Leonardi, P. M. (2013). Social media use in organizations: Exploring the affordances of visibility, editability, persistence, and association. *Annals of the international communication association*, 36(1), 143-189. https://doi.org/10.1080/23808985.2013.11679130

Treepongkaruna, S., & Suttipun, M. (2024). The impact of environmental, social and governance (ESG) reporting on corporate profitability: Evidence from Thailand. *Journal of financial reporting and accounting*, https://doi.org/10.1108/JFRA-09-2023-0555

Tuunanen, T., Winter, R., & vom Brocke, J. (2024). Dealing with complexity in design science research: A methodology using design echelons. *MIS quarterly*, 48(2), 427–458. https://doi.org/10.25300/MISQ/2023/16700

van Atteveldt, W., van der Velden, M. A., & Boukes, M. (2021). The validity of sentiment analysis: Comparing manual annotation, crowd-coding, dictionary approaches, and machine learning algorithms. *Communication methods and measures*, 15(2), 121-140. https://doi.org/10.1080/19312458.2020.1869198

Vieru, D., & Plugge, A. (2025). Overcoming the AI opacity in ESG reporting: A digital platform-based knowledge boundary-spanning perspective. *HICSS 2025*. https://r-libre.teluq.ca/3438/

vom Brocke, J., Watson, R. T., Dwyer, C., Elliot, S., & Melville, N. (2013). Green information systems: Directives for the IS discipline. *Communications of the association for information systems*, 30, 77-102. https://doi.org/10.17705/1CAIS.03330

Zou, X., Jiang, J., Zhang, H., & He, H. (2025). ESG performance, media coverage and brand value. *Asia Pacific journal of marketing and logistics*, 37(1), 171-188. https://doi.org/10.1108/APJML-04-2024-0441