sustainability

MDPI

Article

From News to Knowledge: Leveraging AI and Knowledge Graphs for Real-Time ESG Insights

Omar Mohmmed Hassan Nassar, Fahimeh Jafari * and Chanchal Jain

Computing and Digital Technology Department, University of East London, London E16 2RD, UK; u2617956@uel.ac.uk (O.M.H.N.); u2240956@uel.ac.uk (C.J.)
* Correspondence: f.jafari@uel.ac.uk

Abstract

Traditional Environmental, Social, and Governance (ESG) assessments rely heavily on corporate disclosures and third-party ratings, which are often delayed, inconsistent, and prone to bias. These limitations leave stakeholders without timely visibility into rapidly evolving ESG events. These assessment frameworks also fail to capture the dynamic nature of ESG issues reflected in public news media. This research addresses these limitations by proposing and implementing an automated framework utilising Artificial Intelligence (AI), specifically Natural Language Processing (NLP) and Knowledge Graphs (KG), to analyse ESG news data for companies listed on major stock indices. The methodology involves several stages: collecting a registry of target companies; retrieving relevant news articles; applying Named Entity Recognition (NER), sentiment analysis, and ESG domain classification; and constructing a linked property knowledge graph to structure the extracted information semantically. The framework culminates in an interactive dashboard for visualising and querying the resulting graph database. The resulting knowledge graph supports comparative inferential analytics across indices and sectors, uncovering divergent ESG sentiment profiles and thematic priorities that traditional reports overlook. The analysis also reveals comparative insights into sentiment trends and ESG focus areas across different exchanges and sectors, offering perspectives often missing from traditional methods. Findings indicate differing ESG sentiment profiles and thematic focuses between the UK (FTSE) and Australian (ASX) indices within the analysed dataset. This study confirms AI/KG's potential for a modular, dynamic, and semantically rich ESG intelligence approach, transforming unstructured news into interconnected insights. Limitations and areas for future work, including model refinement and integration of financial data, are also discussed. This proposed framework augments traditional ESG evaluations with automated, scalable, and context-rich analysis.

Keywords: Environmental, Social, and Governance (ESG); artificial intelligence; natural language processing; knowledge graphs; ESG sentiment analysis; sustainability indicators

check for updates

Academic Editors: Yelin Fu, Zelong Yi and George G.Q. Huang

Received: 17 September 2025
Revised: 17 October 2025
Accepted: 27 November 2025
Published: 12 December 2025

Citation: Hassan Nassar, O.M.; Jafari, F.; Jain, C. From News to Knowledge: Leveraging AI and Knowledge Graphs for Real-Time ESG Insights. *Sustainability* **2025**, *17*, 11128. https://doi.org/10.3390/su172411128

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

1. Introduction

Environmental, Social, and Governance (ESG) factors have become essential for assessing corporate sustainability, financial resilience, and ethical operations. Regulators such as the Financial Conduct Authority (FCA) and the Australian Securities and Investments Commission (ASIC) [1] increasingly demand transparent, real-time ESG intelligence [2]. Yet, traditional ESG assessment methods rely heavily on corporate disclosures and third-party ratings [3], which are often delayed, fragmented, inconsistent [4], and unable to capture

rapidly evolving events [5]. Empirical studies show that correlations between major ESG rating providers can be as low as 0.38, in stark contrast to the 0.92–0.99 correlation observed for credit ratings [6]. Such divergence undermines comparability, investor confidence, and the credibility of ESG evaluations.

Unstructured public data, particularly news articles, offer a valuable, timely counterbalance to corporate reports by surfacing real-time ESG narratives, social discourse, environmental disruptions, and governance failures. Yet, the sheer volume and complexity of this information make manual extraction and analysis infeasible. To ground our contribution against prior work of literature, we add that prior KG-centric sentiment studies in finance demonstrate feasibility of fusing textual analytics with graph structures, however; do not operationalise a unified structure; existing NLP-based ESG studies often focus on sentiment or topic detection in isolation; most Knowledge Graph (KG) implementations rely on structured corporate datasets rather than integrating unstructured, real-time sources. No current framework unifies multi-source near real-time news acquisition, Natural Language Inference (NLI)-based domain classification, sentiment analysis, and semantic linking into a dynamic, queryable ESG intelligence framework; this constitutes the central research gap.

This research addresses the gap by proposing an automated framework that leverages Artificial Intelligence (AI), specifically Natural Language Processing (NLP) techniques and Knowledge Graphs (KG) [1], to address these challenges [2]. However, addressing these limitations necessitates automated methods capable of extracting nuanced insights. This study, therefore, explores an AI-driven framework, critically incorporating Knowledge Graphs (KGs) to semantically link disparate information and uncover hidden ESG patterns from news data.

To demonstrate its capabilities, the framework is applied to firms listed on the FTSE 100, FTSE 250, and ASX 200 indices. These indices were chosen due to their economic significance, sectoral diversity, and regulatory variation, providing a robust comparative testbed for evaluating the framework. The findings reveal distinct ESG sentiment patterns and thematic emphases across markets and sectors, highlighting perspectives that are often overlooked in traditional reports. The contributions of this research are threefold. First, it presents a technical integration: an end-to-end framework that unifies diversified data acquisition, NLI-driven ESG classification, and knowledge graph–based semantic structuring for real-time ESG intelligence. Second, it delivers explainable ESG analytics by constructing a dynamic knowledge graph that exposes relational structures between companies, entities, domains, and sentiment, in contrast to static or opaque metrics. Third, it demonstrates practical impact by showing how the proposed framework can provide investors, regulators, auditors, and corporations with timely, transparent, and scalable ESG insights. Taken together, these contributions advance both the field of computer science through innovative applications of NLP, NLI, and KGs and the practice of sustainable finance by offering an adaptable, real-time solution for ESG data analysis and reporting.

The framework processes news articles about firms from the FTSE 100, FTSE 250, and ASX 200 indices. These indices were chosen for their significant economic impact, broad sector representation, geographic diversity (UK and Australia), enabling comparative analysis, and their established importance in ESG reporting literature. This selection created a solid basis for showcasing the framework's abilities across various market characteristics and regulatory settings, with potential future work to include other global indices. It then applies a series of NLP tasks: sentiment analysis using DistilBERT (distilbert-base-uncased-finetuned-sst-2-english) [3], to classify the tone of articles (Positive, Neutral, Negative) and Named Entity Recognition (NER) to identify key entities like organisations, people, locations, and monetary values; and topic modelling (ESG domain classification) using a hybrid approach of Bag-of-Words and Zero-Shot Classification [4], to map article relevance

to Environmental, Social, or Governance domains. The core contribution lies in constructing a knowledge graph [5] that semantically links these extracted entities and insights, transforming ESG news data from a linear format into a richly connected knowledge space that is ideal for pattern detection and relational analytics. This approach moves beyond isolated metrics to reveal underlying connections and dynamics within ESG narratives. Also, Shahrour [6] proves that higher CSR/ESR engagement is associated with lower market-implied default risk firms, which connects ESG activity to a financial material risk channel

The framework's findings reveal distinct ESG sentiment trends and focus areas across exchanges, sectors, and domains, providing insights often missed by traditional reports. This developed framework offers a modular, scalable, and adaptable way to convert unstructured news into structured, inferential-based ESG knowledge, assisting financial analysts, policymakers, regulators, and impact-driven investors in making better-informed decisions.

The research aims to answer the following questions:

(1) How can AI-driven NLP techniques automate and enrich ESG news data extraction for FTSE 100, FTSE 250, and ASX 200 companies?
(2) Can sentiment and named entity recognition offer timely ESG insights beyond what corporate disclosures provide?
(3) In what ways can knowledge graphs enhance the semantic representation and relationship mapping of ESG information?

It seeks to answer the questions through the following objectives:

(1) To implement pre-trained NLP model utilisation for sentiment analysis, named entity recognition, and topic modelling of ESG-related news articles, terms, and actors within the fetched articles.
(2) To construct and visualise a dynamic knowledge graph representing ESG event relationships among FTSE 350 and ASX 200 companies.
(3) To further dive into the semantic representation of interrelationships between ESG narratives that are not captured through traditional methods of analytics.
(4) Exploration of inferential analytics that may be derived from the output of the proposed framework to be further used for augmenting the decision-making process.

The primary contributions of this research are as follows:

(1) The development and implementation of a novel, automated framework that integrates NLP and KG technologies for ESG news analysis. This included custom-built scripts for entity collection, news aggregation, and hybrid approaches for ESG domain classification combining a Bag-of-Words and Zero-Shot Classification.
(2) The construction of a dynamic, property knowledge graph using Neo4j version 5.26.6 to structure semantically and link extracted ESG entities, sentiments, and thematic domain classifications from unstructured news data.
(3) A scalable and adaptable framework for transforming unstructured public news into interconnected, queryable knowledge, offering an alternative to traditional, often static, ESG reporting inferential analysis.
(4) An interactive dashboard for visualising and exploring the KG, enabling stakeholders to identify sentiment trends, ESG focus areas of market directionalities, and relational dynamics across different companies and sectors within the FTSE 100, FTSE 250, and ASX 200 indices.
(5) Demonstration of how AI-driven analysis of news data can provide timely, sentiment-aware, and context-rich insights often missing from corporate disclosures, aiding financial analysts, policymakers, and researchers.

2. Relevant Literature Review

Traditional Environmental, Social, and Governance (ESG) data sources, such as corporate self-reporting and third-party rating agencies, face considerable scrutiny regarding their opacity, subjectivity, and lack of timeliness. The dynamic nature of ESG issues is often inadequately captured by these methods, leading to inconsistencies in ratings. For instance, a recent study [7] highlights that ESG ratings from major providers can exhibit correlations as low as 0.38, a stark contrast to the high correlation (0.92 to 0.99) observed for credit ratings, underscoring the challenge of achieving standardised and comparable ESG assessments.

This discrepancy points to differing methodologies and regulatory landscapes that contribute to inconsistent investor insights and potentially misleading ratings. Ref. [8] further argues that much of the existing ESG data appears incomplete and unreliable, potentially reflecting "noise" rather than clear signals, or worse, facilitating "greenwashing" where sustainability claims lack substantive backing [9]. The challenge is compounded by the sheer volume of metrics involved, with some research indicating over 1000 metrics in play [10], leading to incomparable ratings. In response to these limitations, unstructured data from public news media has emerged as a rich, alternative source of real-time ESG indicators. However, the manual extraction and analysis of this voluminous and complex information are infeasible. This necessitates the development of automated, data-driven frameworks that can deliver timely, transparent, and meaningful interpretations.

Natural Language Processing (NLP) techniques are increasingly being applied to ESG analysis to address these challenges. Sentiment analysis, for example, can gauge the tone of ESG narratives. While lexicon-based tools like TextBlob or VADER have been used, their efficacy in capturing the nuances of complex texts, such as news articles or press releases, is limited. Dia [11] notes that these tools struggle with contextual understanding compared to statistical models like BERT (Bidirectional Encoder Representations from Transformers) and its variants, which offer superior contextual understanding for lengthy and domain-specific texts [12]. Studies like [13] confirm the value of bidirectional encoders for analysing ESG narratives. Parallel evidence from price-prediction tasks indicates deep learning's edge over traditional ML in financial contexts, motivating our use of transformer-based text models [14]. Araci [15] specifically addresses FinBERT and demonstrated its effectiveness for sentiment analysis in financial texts, indicating that domain-specific fine-tuning enhances performance. The DistilBERT model, a distilled version of BERT, has been shown to retain 97% of BERT's capabilities while being significantly faster and lighter, making it suitable for resource-constrained studies [3]. Named Entity Recognition (NER) plays a crucial role in identifying key actors and elements (e.g., organisations, people, locations, monetary values) within ESG texts. While general-purpose NER tools like SpaCy provide foundational capabilities, their application to ESG-specific terms (e.g., "carbon offsetting initiative," "gender pay audit") remains an area for development, as generic models may require domain adaptation.

ESG topic modelling and classification aim to categorise articles into relevant domains (Environmental, Social, Governance). While BoW offers a rapid baseline, it lacks semantic depth [16]. Supervised classifiers can outperform basic tagging but require substantial labelled training data, which is scarce in the ESG domain. Studies like [4] demonstrate the effectiveness of zero-shot classification using Natural Language Inference (NLI) models for text categorisation without task-specific training data, a valuable approach when such data is unavailable, as is often the case for nuanced ESG narratives. This is particularly relevant as ESG themes frequently overlap, suggesting that multi-label classification models might be more effective, though they introduce computational complexity. Popularised by Google's Knowledge Graph, KGs enable entity linkage, disambiguation,

relationship inference, and contextual querying. While KGs have found applications in finance for areas like fraud detection and risk analytics, their use for ESG analysis, particularly leveraging unstructured news data, is a relatively less explored frontier. Most existing ESG-related KG implementations tend to rely on structured data (e.g., board compositions, shareholder structures) or internal sources, without integrating external, unstructured content like news. For instance, Driller [17] investigated the use of KGs for ESG metrics but presented a theoretical prototype without a working example. Conversely, Angioni [18] demonstrated the feasibility of constructing a KG from news articles (in a post-COVID context) with reasonable accuracy, highlighting the potential for tracking complex socio-environmental events.

Relative to sentiment-only frameworks, we contribute structure and queryability; relative to finance-tuned models, we emphasise ESG domaining and entity-level linkage; relative to KG prototypes, we integrate multi-source retrieval with NLI-assisted tagging classifications to produce cross-index, graph native analytics. That situates our contribution across ESG news NLP, knowledge graphs, and event-driven finance. To our knowledge, no prior study delivers the same end-to-end, real-time ESG news-to-KG with inferential indicators, making like-for-like benchmarking infeasible at present; future work will enable comparability by releasing task definitions, datasets, and evaluation protocols to establish community baselines. On the notes of analytical synthesis and gap, news-centric ESG NLP studies often adopt general transformers or FinBERT for sentiment but stop at document-level scores without semantic structuring of actors and relations; KG-oriented work shows the feasibility of graphing news but typically lacks end-to-end, real-time frameworks with investable indicators that are built on empirical underlying movement of sentiment shifts. Collectively, prior literature provides techniques in isolation (e.g., FinBERT sentiment; zero-shot ESG classification; KG exemplars) yet does not deliver a unified system that (a) ingests real-time news, (b) entity-resolves and time-stamps events, (c) materialises them as queryable knowledge graphs, and (d) outputs inferential indicators suitable for downstream finance studies. Our framework's design explicitly targets this gap, integrating zero-shot classification and transformer-based sentiment with dynamic KG construction over FTSE/ASX universes, and specifying the pathway to empirical financial tests outlined later.

A critical challenge in constructing KGs from news data is entity normalisation and disambiguation, ensuring that textual mentions are correctly mapped to real-world concepts. This is particularly complex in ESG contexts where aliases, abbreviations, and context-dependent terms are common, and centralised ESG entity registries are lacking. The current landscape reveals a significant research gap in the synergistic combination of NLP-driven news analysis with dynamic KG construction for comprehensive ESG intelligence. While some studies, for example [5], argue for using robust NLP on textual data for ESG analysis and suggest investigating beyond mere correlations, and [1] have developed ESG-specific corpora such as ESG-FTSE, there is a clear need for frameworks that can transform unstructured, real-time news into interconnected, semantic insights. This research addresses this gap by developing an AI framework that bridges unstructured ESG news content with structured semantic analysis using KGs, aiming to provide a more holistic, timely, and nuanced understanding of the ESG landscape. Finally, the literature also cautions that news-driven indicators reflect editorial and regional discourse norms; we therefore acknowledge bias transmission and include mitigation in our limitations and future-work plan (fairness audits, provenance-aware aggregation, and sensitivity checks) consistent with the recent calls for transparency in NLP-for-ESG measurement [11–13].

3. Methodology

The framework within this study employs a data-driven pipeline architecture, as illustrated in Figure 1, that leverages AI techniques, primarily NLP, to automate the collection, analysis, and structuring of ESG-related news content for companies in the FTSE 100, FTSE 250, and ASX 200 indices.

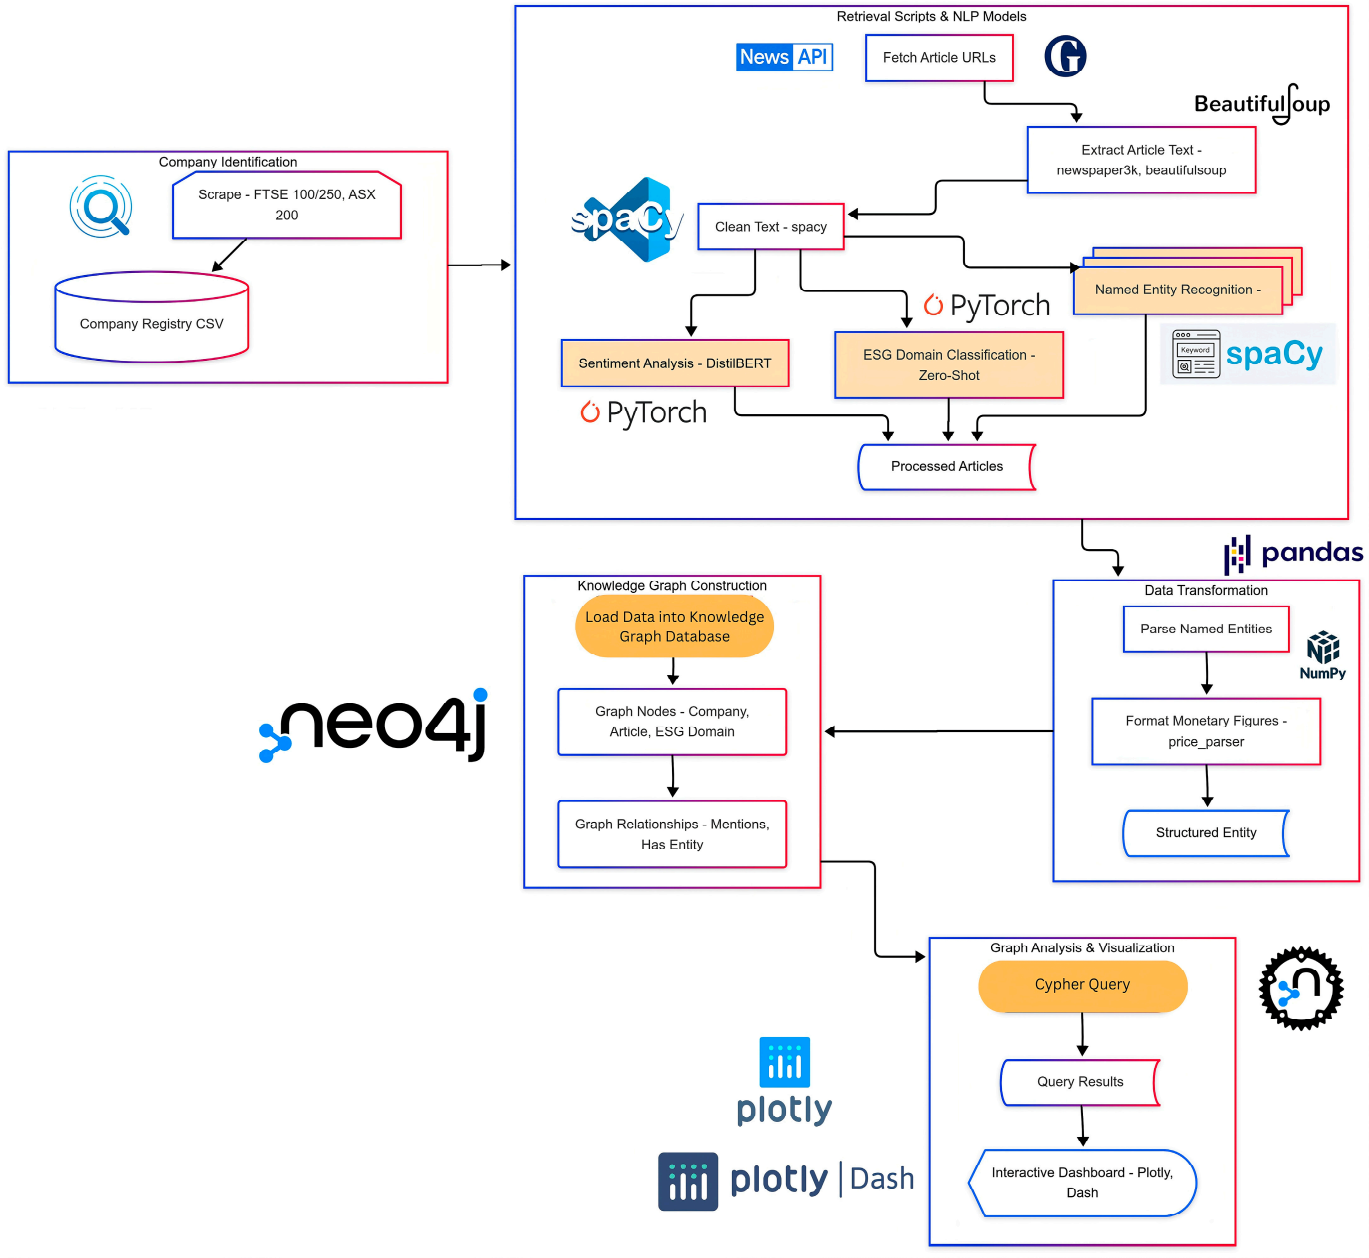


Figure 1. Framework pipeline flowchart overview.

The process involves several key stages that can be broken down as follows:

A. Entity Collection:

The process begins by constructing a "Company Registry" master dataset by scraping publicly available Wikipedia index pages for the FTSE 100, FTSE 250, and ASX 200, yielding a dataset of 550 firms. Each entry includes the company name, stock ticker, and exchange; a sample is presented in Table 1.

Table 1. Company registry sample.

Company Name	Ticker	Sector	Exchange
3i	III	Financial Services	FTSE 100
Admiral Group	ADM	Insurance	FTSE 100
Allianz Technology Trust	ATT	Investment Trusts	FTSE 250
Alpha Group International	ALPH	Financial Services	FTSE 250
Pro Medicus	PME	Healthcare	ASX 200
Tyro Payments	TYR	Information Technology	ASX 200

B. ESG News Article Collection

News articles were collected using NewsAPI and The Guardian's developer API, queried by company name, ticker, and ESG-related keywords (e.g., "ESG," "Sustainability Reports"). Fallback web-scraping methods for sources like BBC and CNN were also implemented using libraries such as BeautifulSoup and Newspaper. Strategies to manage API rate limits included request throttling and API key rotation. Filtering mechanisms were applied to distinguish ESG-relevant articles from purely financial news. Statistics of the acquired data are detailed in Tables 2 and 3.

Table 2. Acquired data for graph construction.

Total News Articles Distribution:	**News Article Distribution by Exchange:**	**News Article Distribution by ESG Domain:**
- **450** Companies out of **550** - **1098** Articles, (**17** Neutral) - **93** Unique Sectors; **29** Grouped	- **FTSE 100;** 214 Articles - **FTSE 250;** 485 Articles - **ASX 200;** 399 Articles	- **Environmental;** 403 - **Governance;** 211 - **Social;** 484
Total Character Count: 1,961,454	Total Characters **with Spaces**; 1,055,018	Total Characters **without Spaces**; 906,436
Corpus: 436 Pages of Content	**Corpus Total Lines:** 20,866	**Corpus Total Words:** 131,161

Table 3. Article distribution (positive/negative).

Stock Index	ESG Domain	Total Count	Positive Articles Count	Negative Articles Count
FTSE 250	Environmental	159	122	37
FTSE 250	Governance	102	44	58
FTSE 250	Social	224	146	78
ASX 200	Environmental	156	100	56
ASX 200	Governance	77	49	28
ASX 200	Social	166	111	55
FTSE 100	Environmental	88	46	42
FTSE 100	Governance	32	18	14
FTSE 100	Social	94	64	30

C. NLP: NER, Sentiment Analysis, and ESG Domain Classification

(a) Named Entity Recognition:

Named Entity Recognition (NER): SpaCy's 'en_core_web_lg' model was used to identify and categorise entities (persons, organisations, locations, monetary values, etc.) within each article. Monetary values were further parsed into numerical formats using the price_parser library for consistency. While not fine-tuned on a domain-specific ESG

corpus due to resource constraints in this study, their generalisability was leveraged. The potential for improved performance with ESG-specific fine-tuning is acknowledged in the limitations and future work.

(b) Sentiment Analysis:

The DistilBERT model was used to assign a polarity score to each article. For individual articles, sentiment was then broadly categorised as 'POSITIVE', 'NEGATIVE', or 'NEUTRAL'. In aggregated analyses presented in this study, such as the average sentiment per exchange, specific thresholds were applied to the average scores; an average sentiment score greater than 0.1 was classified as 'Positive Leaning,' a score less than −0.1 as 'Negative Leaning,' and scores between −0.1 and 0.1 (inclusive) as 'Neutral/Mixed'. These specific thresholds were chosen pragmatically for this study to provide a precise tripartite classification based on the observed distribution of sentiment scores from the DistilBERT model on the initial dataset, allowing for differentiation beyond a simple zero-crossing. While not benchmarked against external studies for these specific values, they served to operationalise sentiment trends for comparative analysis within this framework. Future work could involve sensitivity analysis or calibration against manually labelled datasets to refine these thresholds.

(c) ESG Domain Classification:

A hybrid approach was used. Initially, a Bag-of-Words (BoW) technique with custom keyword lists tailored to ESG themes was developed for the Bag-of-Words component to ensure initial relevance to the domains of Environmental, Social, and Governance categories to provide a baseline classification. This was augmented by a Zero-Shot Classification model ('facebook/bart-large-mnli') leveraging Natural Language Inference (NLI) to predict the ESG domain without task-specific training data. A sample of the processed data is shown in Table 4 below.

Table 4. Graph database data sample.

Company	Sector	Content of Article	Sentiment of Article	ESG Domain	Named Entities Extracted	Exchange
Admiral Group	Insurance	Content Deprecated 500+ Words	Negative	Governance	Admiral Group Plc (ORG); 31 March 2025 (DATE); Justine Roberts (PERSON)	FTSE 100
ASOS	Retailers	Content Deprecated 500+ Words	Negative	Social	LONDON (GPE); Tuesday 16th November 2024 (DATE); Centre Sustainable Fashion CSF (ORG)	FTSE 250
BHP	Materials	Content Deprecated 500+ Words	Positive	Social	Chile (GPE); BHP (ORG); USD 797 million (MONEY)	ASX 200

D. Knowledge Graph Construction and Linkage

The structured outputs from the NLP phase were used to construct a property knowledge graph using Neo4j. Nodes were defined for entities like Company, Article, ESG Domain, Named Entity, Exchange, Sector, and Grouped Sector, with attributes storing relevant information (e.g., ticker, sentiment score). The node schema is detailed in Table 5, shown below.

The relationship schema, detailed in Table 6, connects nodes (from Table 5) using relationships like MENTIONS, HAS_ESG_DOMAIN, LOCATED_ON, and BELONGS_TO, creating a rich, queryable ESG information network using Cypher. These specific properties were chosen to meet core ESG inferential analysis requirements. For instance, linking an "Article" to a "Company" via "Mentions," and to an "ESG Domain" with a sentiment score,

enables queries about how companies are portrayed regarding specific ESG themes and sentiment. Relationships such as "Belongs To" and "Located On" allow for comparative analyses, like examining sentiment trends or ESG focus across industries or markets. This relational structure turns isolated data into an interconnected network, aiding pattern discovery and deeper insights.

Table 5. Graph database node schema.

Element Type	Label	Attributes	Description
Node	Company	Ticker (string, unique), Company (string), Sector (string), Exchange (string),	Represents a company from the input company registry file.
Node	Article	URL (string, unique), Source (string), Content (string), Sentiment (string), ESG tags (string),	Represents a processed news article that is retrieved in the article retrieval stage.
Node	ESG Domain	Domain(string)	Represents a primary ESG category for each article (Environmental, Social, Governance), derived from the topic modelling procedure.
Node	Named Entity	Entity (string, unique), Entity Label(string), Context(string), Numerical Value (float)	Represents an extracted entity from each article; a single article may contain multiple entities (person, org, money, etc.) Numerical Value only holds the money entities.
Node	Exchange	Exchange (string, unique)	Represents a stock exchange or index out of the pre-defined ones (FTSE 100, FTSE 250, ASX 200). Can be customised to accept other exchanges.
Node	Grouped Sector	Normalised Sector (string, unique)	Represents a grouping of sectors to ease relating the sectors that closely align with a broader one, to summarise how major market sectors relate to the other nodes.
Node	Sector	Sector (string, unique)	Represents an industry section that businesses deal in (Finance, Technology, Investments, etc.) Sectors are derived from the company registry file.

Table 6. Graph database relationship schema definition.

Element Type	Label	Description
Relationship	Has Named Entity	Links (Article) → (Named Entity)
Relationship	Has ESG Domain	Links (Article) → (ESG Domain) → (Company)
Relationship	Mentions	Links (Article) → (Company) → (Named Entity)
Relationship	Located On	Links (Company) → (Exchange)
Relationship	Belongs To	Links (Company) → (Sector)
Relationship	Belongs To Group Sector	Links (Sector) → (Grouped Sector) → (Company)

An example Cypher query is shown in Figure 2, and the textual output can be visualised as shown in Figure 3 or can also be visualised as a network graph of interconnected nodes.

```
// Find the first 5 companies listed on the FTSE 100 exchange. – This is a comment in cypher query //
MATCH (c:Company)
WHERE c.exchange = 'FTSE 100'
RETURN c.company AS CompanyName, c.ticker AS Ticker, c.exchange AS Exchange
LIMIT 5
```

Figure 2. Cypher query code sample.

	CompanyName	Ticker	Exchange
1	"Admiral Group"	"ADM"	"FTSE 100"
2	"Airtel Africa"	"AAF"	"FTSE 100"
3	"Anglo American plc"	"AAL"	"FTSE 100"
4	"Antofagasta plc"	"ANTO"	"FTSE 100"
5	"Ashtead Group"	"AHT"	"FTSE 100"

Figure 3. Cypher query text representation.

E. Data Normalisation and Pre-Processing

Following the Natural Language Processing (NLP) extraction, a crucial pre-processing stage ensured data normalisation, numerical consistency, and entity resolution, all vital for reliable knowledge graph construction. This involved several steps:

- Company Entity Normalisation: News mentions of companies were mapped to the canonical names and tickers stored within the "Company Registry." This process ensured that varied references to the same company in different articles were resolved to a single, consistent entity in the knowledge graph.
- Parsing Complex NER Outputs: The outputs from Named Entity Recognition (NER) were parsed to extract and structure relevant information.
- Standardising Data Types: Specific data types were standardised to ensure uniformity across the dataset.
- Conversion of Monetary Values into Numerical: Textual monetary values, such as 'USD 5 million' or '7 cents', were converted into uniform numerical formats. This was achieved using data manipulation libraries like pandas, numpy, and price_parser.
- Normalising Sector Names: Sector names were normalised across the dataset to ensure consistency.

This comprehensive normalisation stage produced a cleaned, entity-centric dataset ready for ingestion into the graph database. While advanced co-reference resolution for all entity types was not part of this initial implementation, the described normalisation steps were key to building an accurate knowledge graph.

F. Visualisation and Analysis

The KG was visualised using Neo4j Bloom for network graph exploration and Plotly Dash for creating an interactive dashboard to display aggregated insights and query results. Standard charting libraries like Matplotlib and Plotly were also used for supplementary

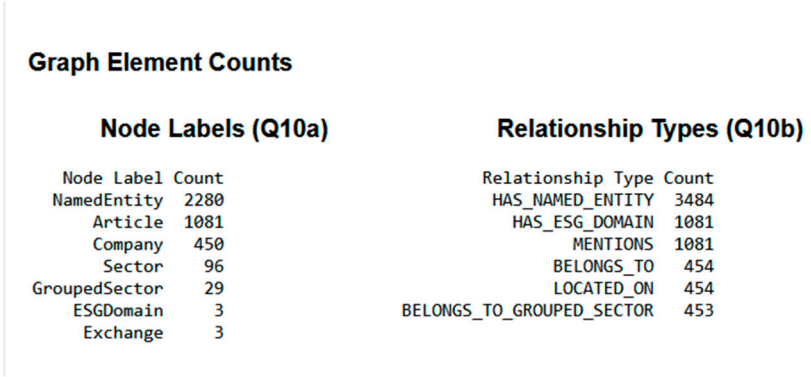
visualisations, with the pre-defined query results passed to Plotly Dash to build an interactive, filterable dashboard for inferential analytics, network graph visual as shown in Figure A1, Appendix A).

4. Results and Discussion

This implementation of the proposed framework, culminating in a queryable knowledge graph and an interactive dashboard, provides substantial evidence of its capacity to automate and enrich ESG news data analysis. The findings, derived from ESG-related news articles concerning FTSE 100, FTSE 250, and ASX 200 companies, address the core research questions by demonstrating the utility of AI-driven NLP techniques and knowledge graphs in generating timely and interconnected ESG insights.

4.1. Environmental, Social, and Governance Insights Across Indices

The automated framework successfully transformed unstructured news articles into a structured graph database, enabling sophisticated queries and visualisations. As illustrated in Figure 4, shown below, the graph comprises various nodes (e.g., Company, Article, ESGDomain, NamedEntity) and relationships (e.g., HAS_NAMED_ENTITY, MENTIONS, HAS_ESG_DOMAIN), which are the direct products of NER and relationship extraction processes. This structure allows for the exploration of complex interdependencies. For instance, Table 7, shown below, showcases an overall sentiment analysis for each stock exchange across different ESG domains. This analysis revealed that the FTSE 250 leaned positively in the environmental domain (0.53 average sentiment), while exhibiting negative leaning in governance (−0.14) and positive leaning in social aspects (0.30). Comparatively, the ASX 200 demonstrated positive leaning across all domains, whereas the FTSE 100 showed neutral/mixed sentiment for environmental issues (0.05) but positive leaning for governance (0.15) and social issues (0.36). These insights, derived from the aggregation of sentiment scores from news articles, offer a perspective often missing from traditional, static corporate disclosures. This is significant as it suggests differing media narratives or public perceptions surrounding ESG aspects for these indices, potentially influencing investor sentiment or highlighting areas where companies on certain exchanges face greater scrutiny or praise for specific ESG domains. For instance, the FTSE 100's neutral/mixed environmental sentiment despite positive social/governance could indicate a complex interplay of large multinational operations facing global environmental criticisms alongside strong domestic governance reporting.



Graph Element Counts

Node Labels (Q10a)

Node Label	Count
NamedEntity	2280
Article	1081
Company	450
Sector	96
GroupedSector	29
ESGDomain	3
Exchange	3

Relationship Types (Q10b)

Relationship Type	Count
HAS_NAMED_ENTITY	3484
HAS_ESG_DOMAIN	1081
MENTIONS	1081
BELONGS_TO	454
LOCATED_ON	454
BELONGS_TO_GROUPED_SECTOR	453

Figure 4. Graph Database Nodes & Relationships Data.

4.2. Average Sentiment Across Indices for Technology & Media Grouped Sector

The inferential analytics achieved through this framework extend to deep market analysis through the mapping of semantics and comparability of indices. The ability to

relate these analytics to various figures and graphs within the dissertation provides a multi-faceted view of ESG performance and perception.

For example, the observation that the FTSE 250's Technology & Media grouped sector had a potentially weaker sentiment profile (neutral/mixed with a score of 0.07 in, as shown in Figure 5 below) compared to the FTSE 100 (0.24) and ASX 200 (0.40) in the same sector, coupled with a strong indicative focus on the Social (S) domain, can be contextualised, this can be further explored to slice into the FTSE 250 index sector by filtration and cross check the inferential analytics, as illustrated in Figure 6, shown below. This aligns with documented tendencies for FTSE 250 companies to emphasise social factors in their ESG narratives and executive remuneration schemes.

Table 7. Overall index sentiment analysis.

Exchange	Environmental Domain	Governance Domain	Social Domain
FTSE 250	Positive Leaning (0.53%)	Negative Leaning (0.14%)	Positive Leaning (0.30%)
FTSE 100	Neutral / Mixed (0.05%)	Positive Leaning (0.15%)	Positive Leaning (0.36%)
ASX 200	Positive Leaning (0.28%)	Positive Leaning (0.27%)	Positive Leaning (0.34%)

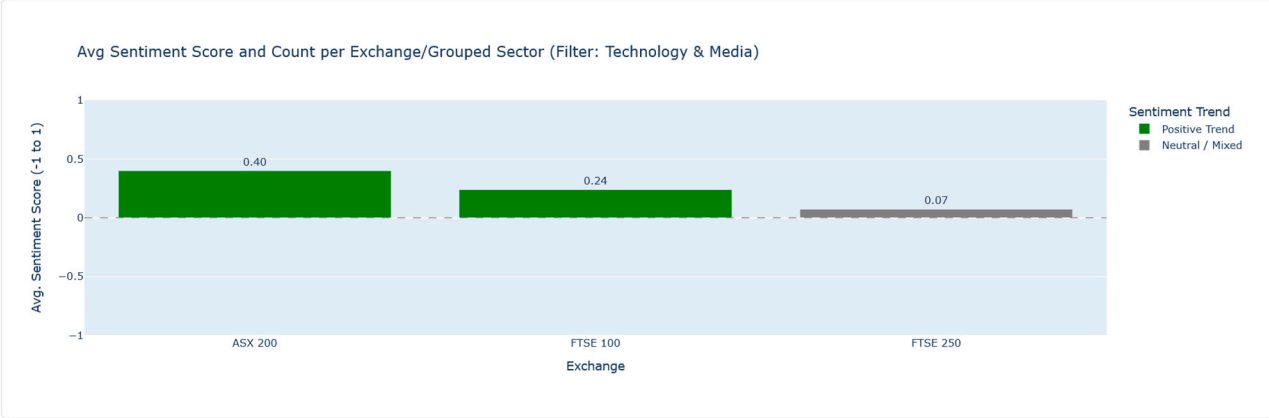


Figure 5. Technology & media (grouped sector) sentiment across indices.

E vs S vs G Focus Direction per Exchange/Grouped Sector (1 entries for filters: Ex=FTSE 250, Sec=Technology & Media)

Exchange	Grouped Sector	Focus Direction	Env. Focus Index	Soc. Focus Index	Gov. Focus Index
FTSE 250	Technology & Media	Focus: Social	0.62	1.15	0.38

Figure 6. Filtration for FTSE 250 (technology and media, grouped sector).

The news sentiment analysis, therefore, captures this prioritisation, which might, in turn, overshadow signals related to environmental and governance performance within that specific sector and index. Furthermore, the framework allows for granular exploration of sentiment within specific sectors across different indices. Figure 7, below, presents the average sentiment for the "Consumer Staples" sector, showing the FTSE 100 (−0.60) significantly underperforming compared to the FTSE 250 (0.67) and ASX 200 (0.55). Such visualisations highlight how external perceptions, reflected in news media, can diverge from a company's or an index's self-assessment or broad market trends.

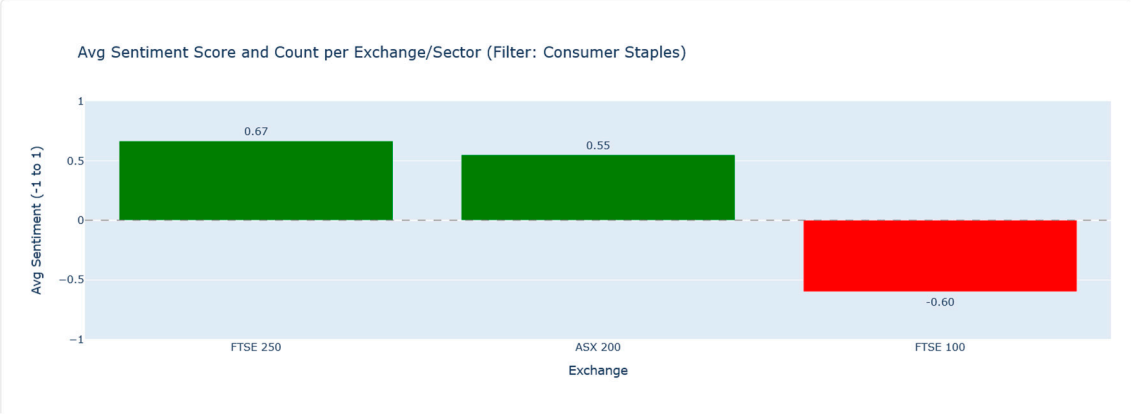


Figure 7. Consumer staples (unique sector) average sentiment across all indices.

4.3. Regulatory Context & Implications

Observed sentiment differences between FTSE and ASX companies may reflect institutional context. UK firms operate under the FCA's evolving ESG disclosure expectations, while Australian firms engage with ASIC's guidance on green claims; both shape media narratives and stakeholder scrutiny. For investors, these sentiment skews indicate where governance versus environmental themes attract heightened external attention, informing stewardship priorities and screening thresholds.

The knowledge graph structure is pivotal in enhancing semantic representation and relationship mapping. It moves beyond simple tabular data by explicitly modelling how entities like companies, articles, ESG domains, and named entities (including monetary values, as seen in Figure 8, illustrated below, linking Barclays to 'USD1 trillion' in the social domain for FTSE 100) are interconnected.

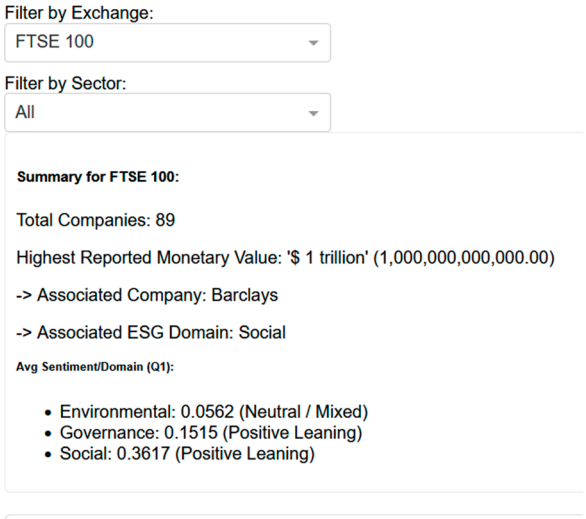


Figure 8. NER by topic, company, and ESG domain (index overview).

This enables the traversal of multi-step relationships, such as mapping the connection between an exchange and the average sentiment score for a specific ESG domain, as explored in Table 7, illustrated on page 12, or determining a grouped sector's relative focus (E vs. S vs. G) by aggregating across linked companies and articles, this is derived from the underlying analysis that can also be viewed in Figure 9, shown below.

E vs S vs G Focus Direction per Exchange/Grouped Sector (57 entries for filters: Ex=All, Sec=All)

Exchange	Grouped Sector	Focus Direction	Env. Focus Index	Soc. Focus Index	Gov. Focus Index
ASX 200	Retail & Consumer Goods	Focus: Environmental & Social	0.74	0.74	0.32
ASX 200	Technology & Media	Focus: Governance	0.83	0.67	1.00
ASX 200	Telecommunications	Focus: Social	0.91	1.45	0.27
ASX 200	Utilities	Focus: Environmental & Social	0.67	0.67	0.33
FTSE 100	Aerospace & Defence	Focus: Social	0.50	1.00	0.50
FTSE 100	Chemicals	Focus: Environmental	2.00	1.00	0.00

Figure 9. ESG focus directionality per exchange over grouped sectors.

Figure 10, shown below, provides a visual example of this semantic mapping, illustrating article sentiment within the Technology sector linked to specific ESG domains and companies, suggesting, for instance, a potentially lower focus on governance aspects for the depicted companies compared to environmental or social issues. A crucial feature of the knowledge graph is colour-encoding for ease of representation, as shown in Figure 10. Negative articles are coloured in a purple colour, while the positive ones are in yellow, with company nodes in light blue, and ultimately ESG domains in orange.

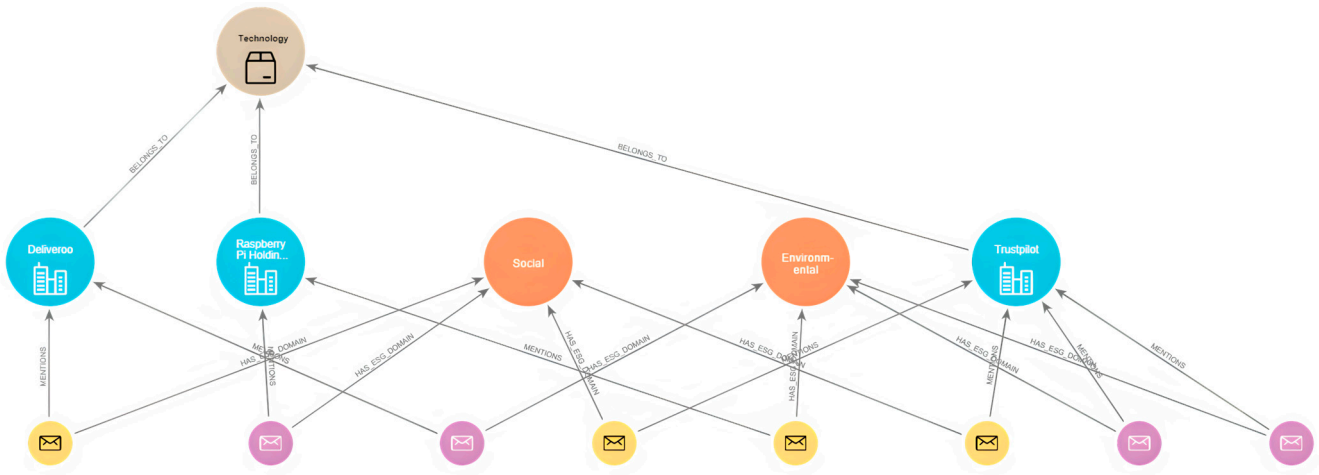


Figure 10. Visual network graph, technology sector, article sentiment sample.

Similarly, the dashboard offers a comparative analysis of mean sentiment scores across various aggregated industry classifications, with the "Food & Tobacco" sector showing the most pronounced negative average sentiment (−0.35), while "Energy & Resources" was closest to neutral (0.07), as shown in Figure A2 in Appendix A.

4.4. Inferential Analytics Discussion Through Dashboards

The disparity where the FTSE 250 showed more positive environmental sentiment than the FTSE 100 (which was neutral/mixed in Figure A2, Appendix A) can be partly attributed to structural differences. The FTSE 100's significant international exposure makes its constituents' sentiment susceptible to global negative news (e.g., trade tariffs) [19], which can obscure company-specific ESG efforts. The FTSE 250's predominantly domestic

focus might insulate it from such global headwinds [20], and research also indicates these companies may inherently cause less environmental damage and are subject to frameworks encouraging positive environmental performance [21,22]. The ASX 200's positive environmental leaning, while potentially reflecting genuine reporting improvements, might also be influenced by a media environment that is possibly less critically focused on ESG issues compared to the UK. In synthesis, the achieved inferential analytics, supported by the visual and relational capabilities of the knowledge graph, provide a dynamic and context-rich lens into the ESG performance and perception of companies. This approach effectively complements traditional corporate disclosures by offering timely, externally focused insights derived from the continuous flow of public news data, highlighting how AI and KG technologies can transform unstructured information into actionable ESG intelligence.

To provide an initial comparative benchmark within this primarily unsupervised study, a subset of the news data was manually annotated with sentiment and domain labels. The accuracy figures derived from this self-labelled sample offer a preliminary understanding of the pre-trained models' performance on the specific dataset.

In evaluating the performance of the framework on a manually labelled subset of raw ESG news data, see Table 8 for performance breakdown, the pre-trained DistilBERT model achieved an overall accuracy of 66.00% for sentiment analysis. In comparison, the Zero-shot model recorded 60.00% accuracy for ESG domain classification. These figures serve as a practical benchmark, considering the inherent complexities of the varied and raw nature of real-world news data. For sentiment analysis, the DistilBERT model demonstrated reasonable performance in identifying 'Positive' sentiment, achieving an F1-score of 0.81 for this majority class. However, it showed lower precision for 'Negative' instances despite high recall. In domain classification, the Zero-shot model performed adequately for core ESG categories, with F1-scores around 0.68–0.69 for 'Environmental' and 'Social' domains and 0.67 for 'Governance'. While these results represent a sound starting point, they also underscore the potential for further refinement through avenues detailed in the "Limitations and Future Work" Chapter V of this paper, such as employing domain-specific models, multi-label classification approaches, and more granular sub-theme modelling.

Table 8. Tabulated quantitative figures of the models.

Prediction Class	Model Used (Pre-Trained)	Metric	Value of Metric
Sentiment Analysis	DistilBERT	Overall Accuracy	66.00%
		F1-Score (Positive)	0.81%
		F1-Score (Negative)	0.30% (Recall: 0.75; Precision: 0.19)
ESG Domain	Zero-Shot Classification	Overall Accuracy	60.00%
		F1-Score (Environmental)	0.68%
		F1-Score (Social)	0.69%
		F1-Score (Governance)	0.67%

5. Limitations and Future Work

- Entity Collection: Reliance on Wikipedia for company registries is pragmatic due to its accessibility and frequent updates, but it could be augmented with official exchange data feeds in the future for enhanced canonical accuracy and to mitigate lag in reflecting index composition changes. The initial scope, limited to FTSE 350 (FTSE 100 and FTSE 250 combined) and ASX 200 for their economic influence and geographic diversity, could be expanded to other global indices or sectors to provide richer cross-market ESG comparisons and improve the generalizability of insights.
- Article Retrieval and Relevance: Public APIs present limitations such as rate restrictions, potentially impacting data volume and recency, and variable search result

relevance from general keyword queries, which can introduce noise. Refining ESG relevance filtering beyond current keyword-based approaches (e.g., "ESG," "Sustainability Reports") using fine-tuned classification models, trained on a dedicated ESG news corpus, would significantly improve precision in identifying truly pertinent articles and reduce the manual effort needed for verification.

- News Coverage & External Validity: Our corpus (~1098 articles retrieved across 450/550 firms) is sufficient for prototyping but not for exhaustive coverage. Future work should explore different language representativeness for better context on regional differences; commit to extending the time window for article collection, also to enable time-series analysis; incorporate additional news APIs and local outlets; and add other indices (e.g., STOXX, S&P 500 /ASX extensions) to strengthen cross-market generalisability.
- NLP—Sentiment Analysis: DistilBERT provides a good balance of performance and computational efficiency, crucial for processing large article volumes; however, larger or domain-specific models like FinBERT could capture more ESG-specific sentiment nuances and contextual subtleties often missed by general-purpose transformers. Addressing mixed sentiments within articles, perhaps through aspect-based sentiment analysis, and mitigating the impacts of text truncation on capturing comprehensive sentiment from lengthy reports remain key areas for future improvement and model refinement.
- NLP—Topic Modelling: The hybrid BoW and zero-shot classification (using models like facebook/bart-large-mnli) was effective for initial ESG domain categorisation without requiring task-specific training data, but this approach could be enhanced for greater depth. Future work should explore multi-label classification to accurately capture articles spanning multiple ESG domains simultaneously and develop more granular ESG sub-theme models (e.g., distinguishing "water scarcity" from "carbon emissions"), likely requiring dedicated, expertly annotated ESG news datasets and defined taxonomies.
- NLP—Named Entity Recognition: SpaCy's en_core_web_lg model is robust for familiar entities, but fine-tuning with an ESG-specific annotated corpus would improve precision for specialised terms (e.g., "carbon offsetting initiative," "green bonds") not well-represented in general training data; this includes enhancing the parsing of complex monetary values beyond the current price_parser capabilities if needed. Advanced entity disambiguation, linked to canonical databases like Wikidata or bespoke ESG entity repositories, is a complex but valuable future addition for enhancing knowledge graph accuracy and resolving entity ambiguities.
- Knowledge Graph Construction: The current property graph schema effectively models core relationships between companies, articles, and ESG domains, but it could be significantly enriched by integrating external structured datasets, such as official company-reported emissions figures or third-party ESG ratings, for cross-validation. Linking entities and concepts to established ontologies like DBpedia or financial ontologies (e.g., FIBO) would also provide greater semantic depth, facilitate more sophisticated reasoning, and improve interoperability with other ESG data systems.
- Knowledge Graph Enhancements & Explainable AI: While the KG provides a rich semantic structure, the complex web of relationships can sometimes make it challenging to trace the exact reasoning behind emergent patterns or inferential conclusions. To enhance user trust and facilitate easier validation of insights, future work should explore the integration of explainable AI (XAI) techniques. Methods such as generating textual explanations for paths within the graph or adapting principles from tools like SHAP or LIME to graph-based queries could help clarify how specific ESG as-

sessments or sentiment trends are derived from the underlying data and relationships within the KG.

- Inclusion of Financial Data: Integrating quantitative financial metrics such as company revenue, market capitalisation, stock price volatility, credit ratings, or specific reported ESG performance indicators alongside news-derived insights could enable robust correlational analyses. This would bridge qualitative ESG news narratives with quantitative market realities, offering a more holistic assessment and allowing for the investigation of potential links between ESG news sentiment, subsequent stock movements, or overall corporate financial health.
- Ethical Considerations and Model Bias: The insights derived from this framework are contingent on the news data processed and the inherent characteristics of the AI models used. News media itself can reflect various biases, including market sentiment, political leanings, or selective reporting, which can influence the textual data fed into the framework itself. Consequently, the sentiment and topics identified by the NLP models may, to some extent, mirror these existing biases from the source material. Future work should incorporate systematic approaches for bias detection and mitigation within the NLP framework. This could involve using fairness toolkits, carefully curating and balancing training data if models are fine-tuned for ESG tasks, and conducting fairness audits to understand how the framework performs across different types of ESG narratives and company profiles. Furthermore, as AI-driven ESG analysis becomes more prevalent, establishing clear lines of accountability for automated outputs is crucial. Overreliance on these tools without understanding their limitations (including inherent biases from data and models) could lead to flawed decision-making by investors, regulators, or companies themselves. Future frameworks should aim for greater transparency in how conclusions are drawn.
- NLP Model Bias: Because news sources embed editorial and regional norms as mentioned before, modelled sentiment and topics can partially inherit those biases. Going forward, we will integrate a formal bias-detection and mitigation pipeline: data curation with stratified sampling, reweighting, and counterfactual data augmentation (e.g., entity/region swaps; negation injection) to balance ESG narratives; fairness auditing using slice-based performance by sector, region, firm size, calibration metrics, and reporting via model cards. To reduce NLP misinterpretations, we will deploy negation and hedging detectors, co-reference and entity disambiguation, contradiction and stance checks (NLI), sarcasm and modality cues, topic-to-entity disentanglement, and KG consistency rules; add uncertainty quantification with abstention or human-in-the-loop review for low-confidence, high-materiality events; and perform domain adaptation and multilingual normalisation to stabilise performance across regions. These measures aim to prevent overreliance on purely automated outputs, foreground model limitations, and make downstream decisions more robust and auditable.
- News Media Bias: As researchers, we need to concur that news-driven ESG signals inherit editorial agendas and regional journalistic norms, producing systematic selection and framing effects. Classic agenda-setting research shows that media emphasis shapes public salience, implying coverage $\neq$ ground truth; thus, indicator intensity can reflect editorial priorities as much as event magnitude. We mitigate by aggregating across diversified outlets/APIs, de-duplicating, recording outlet/region provenance in the KG, normalising for time-varying source coverage, and enabling region-weighted analyses; nevertheless, residual selection effects persist and are explicitly acknowledged and stress-tested via sensitivity analyses. Further research should be expanded thoroughly.

- Financial Linkage to Market Outcomes: While this study primarily focuses on developing and implementing the framework structure, a full econometric validation remains beyond the current scope. Future work will explicitly extend the analysis toward quantitative benchmarking and validation against third-party ESG rating datasets (e.g., MSCI, Refinitiv, or Sustainalytics). This will involve comparing company-level graph-based indicators with existing ESG scores using rank concordance, correlation, and sector-level benchmarking metrics. Additionally, we plan to integrate financial datasets to enable a comprehensive econometric evaluation—transforming time-stamped ESG sentiment intensity, topic salience, and controversy signals into investable indicators. These will be examined through event studies, time-fixed panel regressions, and Granger or VAR-based lead–lag analyses, with appropriate confounder controls, multiple-testing corrections, and robustness checks. This next stage will provide a systematic pathway to assess how the graph-derived ESG indicators correspond to financial performance and market behaviour.
- Evaluation and KG Governance: The project primarily focused on implementation success, demonstrating a functional framework and qualitative insight generation via the dashboard. Future work should incorporate formal evaluation metrics for each NLP task and the KG itself; this would necessitate the development or availability of suitable, expertly annotated ESG benchmark datasets to establish ground truth for entities, sentiments, topics, and relations, allowing for rigorous quantitative performance assessment. As entities and relations evolve, we will maintain versioned graph snapshots with timestamped provenance, periodic schema reviews, and scheduled model refreshes; we will log all ingestion and inference settings to ensure longitudinal comparability.

6. Conclusions & Real-World Application Retrospectives

This research successfully demonstrated the feasibility and potential of an AI-driven framework, integrating NLP and KG technologies, to enhance the collection, assessment, and semantic structuring of ESG data derived from public news sources. The developed end-to-end framework automates the transformation of unstructured news into an interconnected, queryable knowledge graph, representing companies, ESG events, sentiment, and other relevant entities.

The application of knowledge graphs offers a path to model complex ESG relationships, providing a richer, more contextualised view than traditional static data formats allow. The framework offers a valuable proof-of-concept for generating timely, dynamic, and interconnected insights that complement traditional ESG reporting methods. While operating within practical constraints, the identified limitations also present clear avenues for future refinement and expansion, underscoring the robustness and adaptability of the core methodology.

Potential Real-World Applications of This Framework

- Investors and financial analysts can utilise the framework to gain timely insights into ESG-related risks and opportunities potentially not yet reflected in formal company disclosures or lagging third-party ratings. Tracking ESG domain focus (e.g., Figure 9) and sentiment shifts (e.g., Figures 5 and 7) for specific companies or sectors can inform investment decisions, due diligence, and portfolio management, particularly for Socially Responsible Investing (SRI) strategies. Further, hedge funds and asset managers could employ the framework as an early-warning mechanism for reputational or regulatory shocks, incorporating ESG sentiment indicators into trading algorithms and risk assessment models.

- Regulators and policymakers could leverage the framework to monitor corporate ESG discourse in near real-time, identify emerging areas of concern (e.g., widespread negative sentiment in a sector), or detect patterns indicative of potential 'greenwashing' across industries or indices. This can support evidence-based policy making and market oversight activities. For example, securities commissions could use ESG sentiment graphs to detect discrepancies between corporate sustainability disclosures and external narratives, informing targeted audits or compliance investigations.
- ESG auditors and consultants might employ the tool to complement their traditional audit processes by providing an external, media-based perspective on a company's ESG performance and stakeholder perception, helping to identify discrepancies or areas requiring deeper investigation beyond company-provided data. Consultancies could also use the knowledge graph to benchmark ESG practices across peers, generating comparative analytics for advisory reports and assisting clients in preparing for regulatory disclosure mandates.
- Corporations could utilise the framework for proactive reputation management, understanding public sentiment regarding their ESG initiatives (positive or negative), benchmarking their media portrayal against peers, and identifying shifts in ESG focus within their sector. Additionally, corporate boards could incorporate these insights into strategic decision-making, such as tailoring sustainability campaigns, anticipating activist shareholder concerns, or aligning ESG communication strategies with investor and societal expectations.
- Insurance and Risk Management Firms: The framework can support insurers in evaluating ESG-related risks (e.g., climate litigation, governance scandals, or labour disputes) by monitoring media sentiment and thematic focus across industries. This could inform underwriting, risk-adjusted pricing, and exposure monitoring.
- Academia and Research: Researchers in sustainable finance, computational linguistics, and policy studies could leverage the framework as an open data resource for analysing ESG narratives, studying cross-cultural discourse on sustainability, and developing advanced explainable AI techniques for graph-based inference.

On the grounds of scalability, the framework was designed with modularity, demonstrating its potential for scalability. To handle larger volumes of data from more exchanges or to enable near real-time processing, architectural enhancements would be necessary. This could include optimising data ingestion pipelines, employing more robust API management and web scraping strategies, and potentially using distributed processing for NLP tasks or a clustered graph database environment. While the daily batch processing (constrained by resources in this study) proved effective for demonstrating the concept, a transition towards streaming data processing would be required for accurate real-time ESG monitoring. Further investigation would also be needed to assess the NLP model and graph update performance under such high-throughput conditions. In practice, real-time monitoring could enable scenario analysis, for example, regulators detecting ESG controversies within hours, or investors adapting portfolios intra-day based on ESG sentiment volatility. This shift from retrospective analysis to proactive intelligence represents a significant advancement in ESG analytics.

The presented framework offers a robust foundation for real-world deployment, with potential avenues for pilot testing in collaboration with ESG auditors or financial analysts to refine its practical utility. Envisioned future extensions include the development of longitudinal analysis capabilities to track evolving ESG narratives and corporate performance over extended periods, alongside the full implementation of real-time data ingestion and processing to provide truly dynamic ESG monitoring across diverse global markets. Ultimately, this work contributes a functional prototype and a methodological blueprint

for leveraging AI in the ESG domain, paving the way for continued innovation at the intersection of artificial intelligence, sustainable finance, and data science.

Author Contributions: Conceptualization, F.J.; Methodology, O.M.H.N.; Software, O.M.H.N. and C.J.; Validation, F.J.; Formal analysis, O.M.H.N.; Data curation, O.M.H.N. and C.J.; Writing—original draft, O.M.H.N.; Writing—review & editing, F.J. and C.J.; Visualization, O.M.H.N.; Supervision, F.J.; Project administration, C.J. All authors have read and agreed to the published version of the manuscript.

Funding: This research was supported by the UKFin+ Pilot Stream (2025) under the project "Greenwashing Analysis: A Data Science Approach to Evaluating ESG Integrity." This funding supported the foundational work on AI-driven ESG analytics and informed the methodology used in this study.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The full Knowledge Graph, associated datasets, results, and the demo for this study are stored in a private GitHub repository (https://github.com/WCKDNaz/KG-ESG-UEL). Access can be granted upon reasonable request.

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A

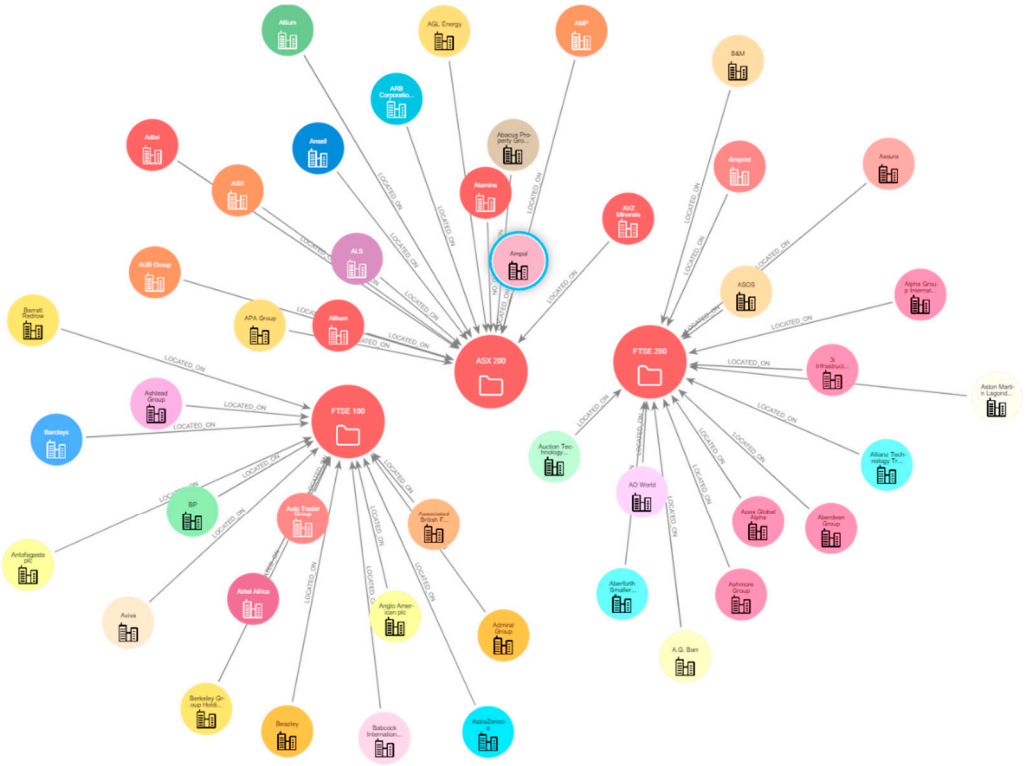


Figure A1. Network graph PoC.

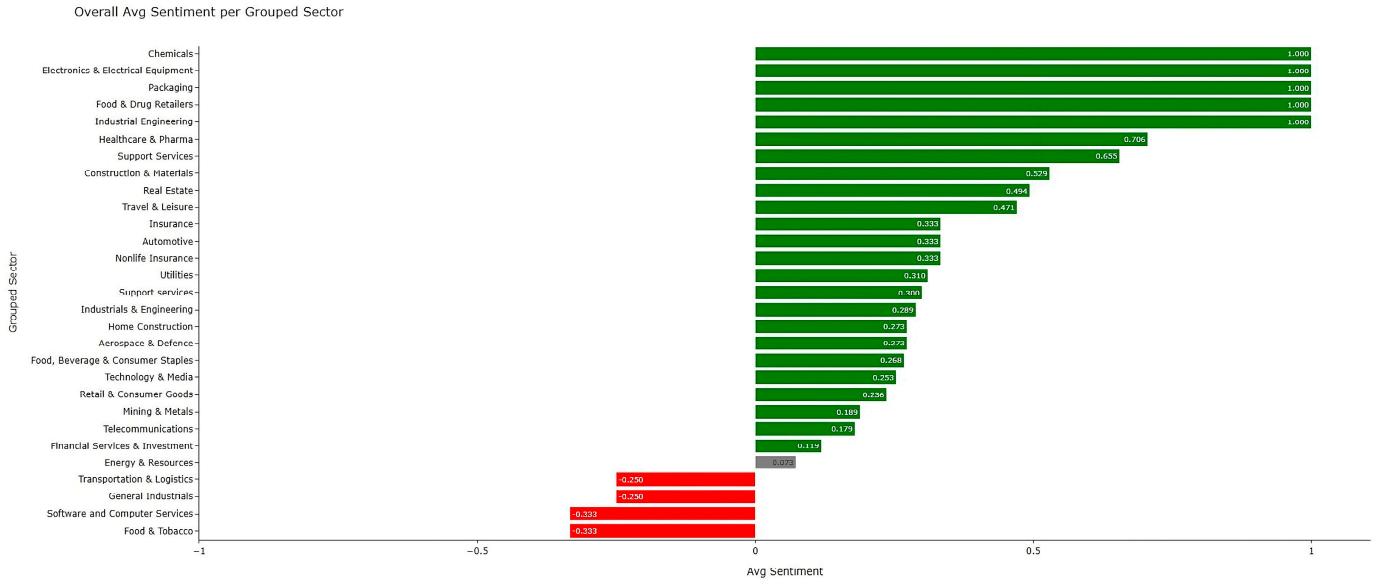


Figure A2. Overall average sentiment by grouped sector.

References

1. Wang, M.P.M.; Casey, B. ESG-FTSE: A Corpus of News Articles with ESG Relevance Labels and Use Cases. Available online: https://www.londonstockexchange.com/ (accessed on 28 April 2025).
2. Kim, M.; Kang, J.; Jeon, I.; Lee, J.; Park, J.; Youm, S.; Jeong, J.; Woo, J.; Moon, J. Differential Impacts of Environmental, Social, and Governance News Sentiment on Corporate Financial Performance in the Global Market: An Analysis of Dynamic Industries Using Advanced Natural Language Processing Models. *Electronics* **2024**, *13*, 4507. [CrossRef]
3. Sanh, V.; Debut, L.; Chaumond, J.; Wolf, T. DistilBERT, a Distilled Version of BERT: Smaller, Faster, Cheaper and Lighter. Available online: https://www.researchgate.net/publication/336230729_DistilBERT_a_distilled_version_of_BERT_smaller_faster_cheaper_and_lighter (accessed on 27 April 2025).
4. Yin, W.; Hay, J.; Roth, D. Benchmarking Zero-shot Text Classification: Datasets, Evaluation and Entailment Approach. Available online: https://www.researchgate.net/publication/335599580_Benchmarking_Zero-shot_Text_Classification_Datasets_Evaluation_and_Entailment_Approach (accessed on 2 May 2025).
5. Schimanski, T.; Reding, A.; Reding, N.; Bingler, J.; Kraus, M.; Leippold, M. Bridging the gap in ESG measurement: Using NLP to quantify environmental, social, and governance communication. *Financ. Res. Lett.* **2024**, *61*, 104979. [CrossRef]
6. Girerd-Potin, I.; Taramasco, O.; Shahrour, M.H. Corporate social responsibility and firm default risk in the Eurozone: A market-based approach. *Manag. Financ.* **2021**, *47*, 975–997. [CrossRef]
7. Berg, F.; Kölbel, J.F.; Rigobon, R. Aggregate Confusion: The Divergence of ESG Ratings. *Rev. Financ.* **2022**, *26*, 1315–1344. [CrossRef]
8. Jones, H. EU Watchdogs See Greenwashing Across the Bloc's Financial Sector. *Reuters*, 1 June 2023. Available online: https://www.reuters.com/world/europe/eu-watchdogs-see-greenwashing-across-blocs-financial-sector-2023-06-01/ (accessed on 28 April 2025).
9. Poiriazi, E.; Zournatzidou, G.; Konteos, G.; Sariannidis, N. Analyzing the Interconnection Between Environmental, Social, and Governance (ESG) Criteria and Corporate Corruption: Revealing the Significant Impact of Greenwashing. *Adm. Sci.* **2025**, *15*, 100. [CrossRef]
10. Tahtinen, J.; Clements, G. Insights from the Reporting Exchange: ESG Reporting Trends. Available online: https://www.cdsb.net/sites/default/files/cdsb_report_1_esg.pdf (accessed on 16 September 2025).
11. Dia, H.; Pettersson, N. Evaluating the Accuracy of Sentiment Analysis Models When Applied to Social Media Texts. Ph.D. Thesis, School of Electrical Engineering and Computer Science (EECS), Stockholm, Sweden, 2024.
12. Koroteev, M.V. BERT: A Review of Applications in Natural Language Processing and Understanding. *arXiv* **2021**, arXiv:2103.11943. [CrossRef]
13. Mehra, S.; Louka, R.; Zhang, Y. ESGBERT: Language Model to Help with Classification Tasks Related to Companies Environmental, Social, and Governance Practices. *arXiv* **2022**, arXiv:2203.16788. [CrossRef]
14. Shahrour, M.H.; Dekmak, M. Intelligent stock prediction: A neural network approach. *Int. J. Financ. Eng.* **2023**, *10*, 2250016. [CrossRef]
15. Araci, D.T. FinBERT: Financial Sentiment Analysis with Pre-trained Language Models. *arXiv* **2025**, arXiv:1908.10063.

16. Qader, W.A.; Ameen, M.M.; Ahmed, B.I. An overview of bag of words; importance, implementation, applications, and challenges. In Proceedings of the 5th International Engineering Conference, Erbil, Iraq, 23–25 June 2019; pp. 200–204. [CrossRef]
17. Driller, J.; Trang, S.T.-N. Unlocking sustainable reporting: Leveraging knowledge graphs for ESG metrics extraction the role of knowledge graphs in sustainability reporting. In *Informatik 2024*; Gesellschaft für Informatik e.V.: Bonn, Germany, 2024. [CrossRef]
18. Angioni, S.; Consoli, S.; Dessì, D.; Osborne, F.; Recupero, D.R.; Salatino, A. Exploring Environmental, Social, and Governance (ESG) Discourse in News: An AI-Powered Investigation Through Knowledge Graph Analysis. *IEEE Access* **2024**, *12*, 77269–77283. [CrossRef]
19. FTSE 100 Dips as Tariff Worries Hit Sentiment. Available online: https://www.sharesmagazine.co.uk/news/shares/ftse-100-dips-as-tariff-worries-hit-sentiment (accessed on 7 May 2025).
20. LSEG FTSE UK Index Series. Available online: https://www.lseg.com/en/ftse-russell/indices/uk (accessed on 7 May 2025).
21. Lamont, D.; Mccarthy, H. Sustainability Stock Market League Table: Which Rank Best and Worst? Available online: https://www.schroders.com/en/global/individual/insights/sustainability-stock-market-league-table-which-rank-best-and-worst-/ (accessed on 7 May 2025).
22. Barnes, C.; Brien, J.; Zhu, S.; Abdul, M. Paying for Sustainable Growth. Available online: https://assets.kpmg.com/content/dam/kpmgsites/uk/pdf/2021/11/paying-for-sustainable-growth.pdf (accessed on 7 May 2025).

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.