# Aligning ESG Controversy Data with International Guidelines through Semi-Automatic Ontology Construction

Tsuyoshi Iwata\*<sup>1,2</sup>, Guillaume Comte<sup>2</sup>, Melissa Flores<sup>2</sup>, Ryoma Kondo<sup>†3,4</sup>, and Ryohei Hisano<sup>‡3,4</sup>

<sup>1</sup>University of Zürich (UZH), Plattenstr. 14, Zurich, 8032, Switzerland
<sup>2</sup>RepRisk AG, Stampfenbachstrasse 42, Zurich, 8006, Switzerland
<sup>3</sup>Graduate School of Information Science and Technology, The University of Tokyo, 7-3-1
Hongo, Bunkyo-ku, Tokyo, Japan
<sup>4</sup>The Canon Institute for Global Studies, ShinMarunouchi Building 5-1 Marunouchi 1-chome, Chiyoda-ku, Tokyo, Japan

September 16, 2025

#### Abstract

The growing importance of environmental, social, and governance data in regulatory and investment contexts has increased the need for accurate, interpretable, and internationally aligned representations of non-financial risks, particularly those reported in unstructured news sources. However, aligning such controversy-related data with principle-based normative frameworks, such as the United Nations Global Compact or Sustainable Development Goals, presents significant challenges. These frameworks are typically expressed in abstract language, lack standardized taxonomies, and differ from the proprietary classification systems used by commercial data providers. In this paper, we present a semi-automatic method for constructing structured knowledge representations of environmental, social, and governance events reported in the news. Our approach uses lightweight ontology design, formal pattern modeling, and large language models to convert normative principles into reusable templates expressed in the Resource Description Framework. These templates are used to extract relevant information from news content and populate a structured knowledge graph that links reported incidents to specific framework principles. The result is a scalable and transparent framework for identifying and interpreting non-compliance with international sustainability guidelines.

#### 1 Introduction

Environmental, social, and governance (ESG) investing has become a global priority as financial institutions, regulators, and stakeholders increasingly integrate non-financial factors into decision-making processes [23]. Recent policy developments such as the European Union's Corporate Sustainability Reporting Directive (CSRD) [11] and Sustainable Finance Disclosure Regulation [10] represent emerging regulatory frameworks aimed at standardizing ESG disclosures. In parallel, global

<sup>\*</sup>tsuvoshi.iwata@df.uzh.ch

<sup>\*</sup>kondor@g.ecc.u-tokyo.ac.jp

<sup>&</sup>lt;sup>‡</sup>hisanor@g.ecc.u-tokyo.ac.jp

normative frameworks such as the United Nations Global Compact (UNGC) [24] and Sustainable Development Goals (SDGs) [25] promote broader sustainability principles and ethical commitments. Together, these developments reflect a broad movement toward more comprehensive and transparent ESG reporting. As a result, there is a growing demand for reliable ESG data that capture emerging risks, sector-specific impacts, and developments across supply chains and financial markets [1].

Meeting this demand remains a complex challenge. Traditional ESG data sources, including corporate sustainability and annual reports [20], are typically self-reported and often emphasize positive performance, which can introduce bias and limit coverage. Independent sources such as news media have therefore become essential for uncovering controversies, verifying corporate claims, and identifying emerging risks [3]. ESG data derived from such news coverage are often referred to as ESG controversy data because they highlight adverse incidents that may signal underlying risks not captured in official disclosures [9]. However, news articles are written in free text, published in high volume, and distributed across a wide range of outlets and narrative threads. This complexity is further heightened by issues such as bias and misinformation. To manage this, commercial vendors annotate news using proprietary ESG classification systems. For example, MSCI applies 28 categories [16], Sustainalytics uses 22 [15], Refinitiv uses 24 [6], and RepRisk uses 108 [21]. These proprietary taxonomies differ significantly in structure and terminology. Because the design and application of these systems are often opaque, the resulting ESG ratings lack consistency across providers and are difficult to align with both formal regulatory frameworks and high-level normative frameworks [7, 8]. This is particularly problematic for principle-based frameworks such as the UNGC, which are normative in nature and do not prescribe specific metrics [2, 12]. In practice, this makes it necessary to analyze original news content directly to ensure the accurate identification of ESG events and alignment with relevant standards.

Even with access to relevant ESG news content, aligning extracted information with ESG frameworks presents additional complexity. Normative frameworks such as the UNGC and SDGs, in addition to regulatory instruments such as the CSRD, are typically principle based and written in abstract, high-level language. These frameworks often lack standardized taxonomies and are subject to ongoing reinterpretation, amendment, or contextual adaptation. Mapping unstructured news content to such evolving frameworks requires more than static keyword tagging or rule-based systems. For example, assigning a general class such as "human rights violations" may overlook new distinctions introduced in revised guidelines, such as variations in due diligence expectations or cross-border responsibilities. These semantic mismatches contribute to the well-known discrepancies in ESG ratings among major data providers and limit the ability of stakeholders to draw consistent conclusions [7, 8].

To address the challenge of aligning unstructured ESG content with high-level normative frameworks, we propose a method for constructing ESG knowledge graphs using lightweight ontology design, pattern instantiation, and structured information extraction, all supported by large language models (LLMs). As illustrated in Figure 1, our approach begins with defining a meta-ontology that links ESG news, frameworks, and principles via violation patterns grounded in the UNGC framework (Figure 1(1)). For each principle, we use language models to generate concise Resource Description Framework (RDF) comments and three representative relational patterns, encoded as instances of an abstract violation pattern class (Figure 1(2)). Then we promote these patterns to 30 ESG violation templates (3 patterns  $\times$  10 UNGC principles), each modeled as a dedicated RDF class to enable type-level reasoning (Figure 1(3)). Finally, we apply this framework to ESG-related news articles, where language models extract named entities and identify text segments that match the predefined patterns. We link matched triples to the relevant UNGC principles and instantiate them in a compliance-aligned knowledge graph (Figure 1(4)).

Our approach differs from those in previous studies on ESG knowledge graphs within the semantic

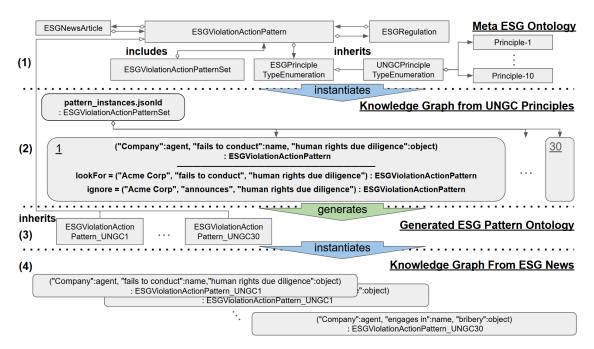


Figure 1: Overview of the proposed method.

web field because we align unstructured text directly with normative principles using automated, ontology-based techniques [3, 17, 20, 13]. By matching extracted content to violation patterns associated with specific principles, we support interpretable detection of ESG non-compliance. These capabilities were not addressed in previous studies. Our approach is also connected to recent work on LLM-based ontology learning [5]. We introduce ESG-aligned rdfs:Class definitions, which are integrated into existing ontologies through subclass inheritance, opening the door to a richer ESG knowledge graph. All code and data are publicly available: https://github.com/tsuyoshiiwataRR/NEWS\_LLM\_UT\_RR.

## 2 Methodology

Figure 1 presents our four-stage approach for constructing ESG knowledge graphs that align unstructured news content with formal ESG regulatory principles.

(1) Meta ESG Ontology. We begin by defining a lightweight ontology that connects ESG-related news, regulations, and principles through structured violation patterns. To achieve this, we extend schema:NewsArticle and schema:Legislation from Schema.org [22], and hna:ActionPattern from our previous study [14], to introduce the classes ESGNewsArticle, ESGRegulation, and ESGViolationActionPattern. These classes are linked to ESG principles via ESGPrincipleTypeEnumeration, which we instantiate as UNGCPrincipleTypeEnumeration to represent the ten principles of the UNGC, covering human rights, labor, environment, and anti-corruption. Because of their abstract and normative phrasing, these principles are difficult to operationalize directly, which makes them the ideal benchmark for evaluating principled event extraction. To support this alignment, we use LLMs to generate concise rdfs:comment summaries for each principle. These serve as semantic anchors for downstream processing. We also define ESGViolationActionPatternSet to organize multiple patterns associated with each principle. All classes are aligned to upper ontologies (e.g., Schema.org) and typed with explicit domains/ranges, enabling RDFS/OWL reasoning over subclass

hierarchies and properties.

#### (2) Pattern Instantiation.

To operationalize these abstract principles, we first collected the official textual descriptions of each UNGC principle from the United Nations Global Compact website [24]. We then prompt LLMs to generate representative violation patterns from these descriptions. For each principle summary, the model outputs three relational patterns in the form (Entity A, Action, Entity B), which we design to capture both direct and indirect violations (e.g., (Company, violates, rights of indigenous communities)). We encode these patterns as instances of ESGViolationActionPattern in JSON-LD format. Each instance includes positive and negative examples, which we specify through the lookFor and ignore properties, respectively. This structured representation forms the foundation for fine-grained pattern matching and subsequent ontology construction.

- (3) Semi-Automatic Ontology Construction. Based on the instantiated patterns, we semi-automatically define 30 rdfs:Class entities (corresponding to the total number of patterns, i.e.,  $3\times10$  = 3 patterns for each of the 10 UNGC principles) by promoting each ESGViolationActionPattern instance to a subclass of ESGViolationActionPattern. These subclasses represent specific types of ESG violations and provide type-level constraints for annotating actions in ESG-related news. This intermediate ontology layer enables principled reasoning over ESG behavior and facilitates alignment with abstract regulatory norms.
- (4) Knowledge Graph Construction. To populate the graph, we use a large collection of negatively framed news articles sourced from Webz.io [26] because such articles are more likely to report instances of corporate misconduct. We apply a two-stage filtering process using multiple LLMs (GPT-40 mini [18], GPT-4.1 [19], and Claude 3.7 [4]). First, we retain only English-language articles and exclude those unrelated to corporate or ESG topics. Second, we use LLMs to further refine the corpus, selecting articles that mention a company and describe a plausibly ESG-relevant negative event. This results in a final candidate set of approximately 800 to 2,000 articles, which depends on the language model.

For each article, we apply a structured prompting strategy to extract named entities categorized into organizations, persons, and locations, while excluding non-informative content such as author bylines or boilerplate text. Then we evaluate the article against a predefined set of 30 violation patterns. For each matched pattern, we use a carefully designed prompt to instruct the language model to extract a corresponding triple as in [14]. We ground the resulting triple using the identified named entities and linked to the appropriate UNGC principle. This ensures that everything is based on textual evidence.

### 3 Evaluation

To assess the effectiveness of our ontology-informed extraction method, we conducted a manual evaluation using 200 human-labeled samples. We randomly selected these samples from the refined corpus produced by GPT-40 mini during the article selection stage. This choice was motivated by the observation that GPT-40 mini retained the largest set of candidate articles compared to other models, providing a broader and richer pool of candidate samples. In contrast, GPT-4.1 and Claude 3.7 exhibited significantly more conservative selection behavior, which limited their utility as sources to build a diverse and representative evaluation dataset. The samples represented matched triples aligned with the UNGC Principles. We evaluated the method using standard metrics: precision, accuracy, and recall.

Table 1 provides a comparison across language models. GPT-40 mini achieved the best per-

Table 1: Comparison table among different LLMs.

Model	TP	FN	FP	TN	Accuracy	Precision	Recall
GPT-40 mini	80	116	62	1742	0.91	0.56	0.41
GPT-4.1	31	165	24	1780	0.91	0.56	0.16
Claude 3.7 Sonnet	27	169	62	1742	0.88	0.30	0.14
GPT-40 mini ('one-shot')	166	30	$\bar{411}$	$\bar{1}\bar{3}\bar{9}\bar{3}$	0.78	0.29	0.85

formance across all evaluated metrics. Additionally, our ontology-guided method outperformed a simple one-shot approach, which attempted to assign principles directly based on their short descriptions and the full article text. Although the one-shot method achieved relatively high recall, this came at the cost of significantly lower precision and accuracy. The method frequently generated unsupported predictions, introducing noise and reducing reliability. These findings underscore the value of our structured method, particularly in settings where precision is more important than recall.

Table 2 presents the results for our structured prompting pipeline by GPT-40 mini. The model achieved consistently high accuracy across all ten UNGC Principles, with particularly strong performance in correctly identifying non-relevant content. This was reflected in the high number of true negatives, which contributed significantly to overall accuracy. Precision varied by principle, with the highest value of 0.86 observed for Principle 1. Principle 6 yielded the lowest precision, excluding Principle 9, which reflected differences in how clearly certain violations were expressed in the articles.

Although overall performance was strong, certain challenges remain. Precision was occasionally reduced due to broad interpretations of principle definitions. For example, general descriptions of poor working conditions may be classified as forced labor, thereby producing false positives for Principle 4. Recall performance also varied. Principles 1, 4, 6, and 8 achieved relatively low recall, often because of indirect or nuanced language that implies rather than states a violation. For instance, references to surveillance or safety concerns may suggest human rights issues, but are not always recognized as such. By contrast, Principles 2, 5, and 7 demonstrated the highest recall. Violations under Principle 5, such as child labor, were rare but clearly expressed, thereby making them easier to identify. Principle 7 appeared frequently because of its broader scope.

As an additional outcome of our experiments, we report the probability that a violation on one date is followed by another violation on a subsequent day by the same entity, as identified by our method, using the news dataset shown in Figure 2. The figure indicates a recurring pattern of violations within the same Principle, as well as cross-Principle infringement patterns. However, a larger dataset would be needed to confirm the robustness of these findings.

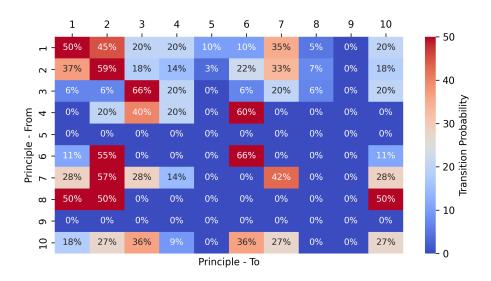
#### 4 Conclusion

In this paper, we presented an ontology-informed method to align unstructured ESG news with high-level principles through pattern-based extraction and knowledge graph construction. Our approach integrates lightweight ontology design and LLM-driven pattern generation to produce semantically consistent RDF triples of potential ESG violations. Applied to the UNGC principles, it outperformed a baseline one-shot approach in both accuracy and interpretability. The resulting knowledge graph supports applications such as tracking ESG trends [3] and predicting future risk events [13], which we leave for future work.

Table 2: Summary of s Accuracy Metrics, among human-labeled 200 samples (GPT-40 mini)

UNGC Principle	TP	FN	FP	TN	Accuracy	Precision	Recall
1 (Human Rights -Respect rights)	25	41	4	130	0.78	0.86	0.38
2 (Human Rights -Avoid abuse)	16	10	13	161	0.88	0.55	0.62
3 (Labour -Union rights)	8	10	12	170	0.89	0.40	0.44
4 (Labour -No forced labour)	4	10	3	183	0.94	0.57	0.29
5 (Labour -No child labour)	4	3	2	191	0.98	0.67	0.57
6 (Labour -Equal opportunity)	3	6	10	181	0.92	0.23	0.33
7 (Environment - Prevent harm)	12	9	8	171	0.92	0.60	0.57
8 (Environment - Act sustainably)	5	13	1	181	0.93	0.83	0.28
9 (Environment - Eco innovation)	0	1	1	198	0.99	0	0
10 (Anti Corruption)	3	13	8	176	0.90	0.27	0.19
All	80	116	62	$\overline{1742}$	0.91	0.56	0.41

Figure 2: Transition probability heatmap in percentage, From a violation in a Principle to another violation in another date by a same company. Total number of the pairs is 181.



### 5 Acknowledgments

R.H. is supported by JST FOREST Program (JPMJFR216Q), JST PRESTO Program (JPMJPR2469), Grant-in-Aid for Scientific Research (KAKENHI, JP24K03043), and the UTEC-UTokyo FSI Research Grant Program. R.K. is funded by JST, ACT-X Grant Number JPMJAX23CA, Japan. During the preparation of this work, the author(s) used ChatGPT, Grammarly in order to: Grammar and spelling check, Paraphrase and reword. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the publication's content. We also thank Edanz (https://jp.edanz.com/ac) for editing a draft of this manuscript.

#### References

- [1] Evolution of ESG Reporting Frameworks. https://scispace.com/papers/evolution-of-esg-reporting-frameworks-2mosv0oa3x, January 2020.
- [2] Nathan Andrews. The UN Global Compact: An Overview of the Promise and Pitfalls. In David Crowther and Shahla Seifi, editors, *The Palgrave Handbook of Corporate Social Responsibility*, pages 1–21. Springer International Publishing, Cham, 2020. ISBN 978-3-030-22438-7. doi: 10.1007/978-3-030-22438-7-75-1. URL https://doi.org/10.1007/978-3-030-22438-7-75-1.
- [3] Simone Angioni, Sergio Consoli, Danilo Dessì, Francesco Osborne, Diego Reforgiato Recupero, and Angelo Salatino. Exploring Environmental, Social, and Governance (ESG) Discourse in News: An AI-Powered Investigation Through Knowledge Graph Analysis. *IEEE access: practical innovations, open solutions*, 12:77269–77283, 2024. ISSN 2169-3536.
- [4] Anthropic. Claude 3.7 Sonnet: Hybrid reasoning AI model, 2025.
- [5] Hamed Babaei Giglou, Jennifer D'Souza, and Sören Auer. LLMs4OL: Large language models for ontology learning. In Terry R. Payne, Valentina Presutti, Guilin Qi, María Poveda-Villalón, Giorgos Stoilos, Laura Hollink, Zoi Kaoudi, Gong Cheng, and Juanzi Li, editors, *The Semantic Web – ISWC 2023*, pages 408–427, Cham, 2023. Springer Nature Switzerland. ISBN 978-3-031-47240-4.
- [6] Florian Berg, Kornelia Fabisik, and Zacharias Sautner. Is History Repeating Itself? The (Un)Predictable Past of ESG Ratings, August 2021. URL https://papers.ssrn.com/ abstract=3722087.
- [7] Florian Berg, Julian F. Kölbel, and Roberto Rigobon. Aggregate confusion: The divergence of ESG ratings. *Review of Finance*, 26(6):1315–1344, 2022.
- [8] Monica Billio, Michele Costola, Iva Hristova, Carmelo Latino, and Loriana Pelizzon. Inside the ESG ratings: (Dis)agreement and performance. *Corporate Social Responsibility and Environmental Management*, 28(5):1426–1445, 2021. ISSN 1535-3966. doi: 10.1002/csr.2177. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/csr.2177. \_eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/csr.2177.
- [9] Ahmed A. Elamer, Ahmed Alhaj-Ismail, Ali M. Gerged, and Collins G. Ntim. ESG controversies and corporate performance: The moderating influences of governance structures and ESG practices. *Corporate Social Responsibility and Environmental Management*, 31(4):3312–3327, 2024.
- [10] European Union. Regulation (EU) 2019/2088 of the european parliament and of the council of 27 november 2019 on sustainability-related disclosures in the financial services sector (sustainable finance disclosure regulation, SFDR). urlhttps://eur-lex.europa.eu/eli/reg/2019/2088/oj/eng, 2019.
- [11] European Union. Directive (EU) 2022/2464 of the european parliament and of the council of 14 december 2022 amending regulation (EU) no 537/2014, directive 2004/109/EC, directive 2006/43/EC and directive 2013/34/EU, as regards corporate sustainability reporting (corporate sustainability reporting directive, CSRD), 2022.

- [12] Ion Frecautan and Andreea Nita Danila. WHO IS GOING TO WIN: THE EU ESG REGULATION OR THE REST OF THE WORLD? A CRITICAL REVIEW. The Annals of the University of Oradea. Economic Sciences, 31(volume 31):109–120, December 2022. ISSN 1582-5450, 1222-569X. doi: 10.47535/1991AUOES31(2)011. URL http://anale.steconomiceuoradea.ro/en/wp-content/uploads/2023/01/Volume-2\_AU0ES\_december-2022-112-123.pdf.
- [13] Ryohei Hisano, Didier Sornette, and Takayuki Mizuno. Prediction of ESG compliance using a heterogeneous information network. *Journal of Big Data*, 7(1):22, March 2020. ISSN 2196-1115.
- [14] Riona Matsuoka, Hiroki Matsumoto, Takahiro Yoshida, Tomohiro Watanabe, Ryoma Kondo, and Ryohei Hisano. Hierarchical narrative analysis: Unraveling perceptions of generative AI. In Proceedings of Jinmoncon 2024, IPSJ SIG Computers and the Humanities, volume 2024, pages 91–98, 2024.
- [15] MORNINGSTAR. Sustainalytics Introduces Signif-Morningstar URL icant Enhancements to its ESG Risk Ratings, 2024.https://www.sustainalytics.com/esg-news/news-details/2024/05/31/ morningstar-sustainalytics-introduces-significant-enhancements-to-its-esg-risk-ratings.
- [16] MSCI ESG Research. MSCI ESG Controversies and Global Norms Methodology. 2024.
- [17] Keane Ong, Rui Mao, Frank Xing, Ranjan Satapathy, Johan Sulaeman, Erik Cambria, and Gianmarco Mengaldo. ESGSenticNet: A Neurosymbolic Knowledge Base for Corporate Sustainability Analysis, January 2025.
- [18] OpenAI. GPT-40 mini: Advancing cost-efficient intelligence, 2024.
- [19] OpenAI. GPT-4.1: OpenAI's flagship large language model, 2025.
- [20] Cristina-Claudia Osman, Ana-Maria Ghiran, and Robert Andrei Buchmann. Towards a knowledge management capability for ESG accounting with the help of enterprise modeling and knowledge graphs. In Forum and Doctoral Consortium Papers at the 17th IFIP WG 8.1 Working Conference on the Practice of Enterprise Modeling (PoEM 2024) Forum Proceedings, volume 3855 of CEUR Workshop Proceedings, pages 112–123. CEUR-WS.org, 2024.
- [21] RepRisk. RepRisk methodology overview, 2025.
- [22] Schema.org. Schema.org, 2025.
- [23] Matthew W. Sherwood and Julia Pollard. Responsible Investing: An Introduction to Environmental, Social, and Governance Investments. Routledge, London, September 2018. ISBN 978-0-203-71207-8.
- [24] United Nations. United nations global compact: The ten principles, 2000.
- [25] United Nations General Assembly. Transforming our world: The 2030 agenda for sustainable development, 2015.
- [26] Webz.io Ltd. Webz.io: Open, deep, and dark web solutions, 2025.