2024 IEEE International Conference on Data Mining Workshops (ICDMW)

Neurosymbolic AI for Mining Key Aspects of Socially Responsible Investing

Wihan van der Heever
College of Computing and Data Science
Nanyang Technological University
Singapore
wihan001@e.ntu.edu.sg

Ranjan Satapathy
Institute of High Performance Computing (IHPC)
*Agency for Science, Technology and Research (A*STAR)*
Singapore
satapathy_ranjan@ihpc.a-star.edu.sg

Erik Cambria
College of Computing and Data Science
Nanyang Technological University
Singapore
cambria@ntu.edu.sg

Rick Siow Mong Goh
Institute of High Performance Computing (IHPC)
*Agency for Science, Technology and Research (A*STAR)*
Singapore
gohsm@ihpc.a-star.edu.sg

***Abstract*—Environmental, Social, and Governance (ESG) factors have become critical for assessing corporate sustainability and ethical responsibility. However, the vast volume of unstructured data available across corporate reports, social media, and news sources poses a challenge for systematic ESG analysis. This paper explores the application of neurosymbolic AI, which combines neural networks' pattern recognition capabilities with the structured reasoning of symbolic AI, to mine key aspects of ESG from large-scale, diverse data sources. By leveraging SenticNet for concept parsing and deep learning for sentiment analysis, we extract relevant ESG metrics, classify corporate practices, and identify trends. This hybrid approach enhances both the interpretability and scalability of ESG analysis, providing more accurate insights into corporate behaviors and their impact on sustainability goals. Results demonstrate that neurosymbolic AI not only improves the extraction of meaningful ESG aspects but also enables real-time monitoring, supporting data-driven decision-making for investors, regulators, and stakeholders.**

***Index Terms*—ESG, Sentiment Analysis, NLU, AI**

I. Introduction

In recent years, the integration of Environmental, Social, and Governance (ESG) criteria has gained substantial traction within the investment community. ESG refers to a framework used to evaluate the sustainability and ethical impact of an investment in a company. The Environmental component addresses how a corporation interacts with the natural environment, focusing on issues such as carbon emissions, climate change, resource depletion, waste management, pollution, and biodiversity. The Social dimension examines a company's relationships with employees, suppliers, customers, and communities, covering aspects such as labor practices, employee relations, diversity and inclusion, human rights, community engagement, and consumer protection. The Governance aspect evaluates the internal practices and policies that govern a company, including board composition, executive compensation, shareholder rights, ethical conduct, transparency, and anti-corruption measures.

The relevance of ESG considerations in investment decisions is underscored by a growing body of evidence suggesting that companies with strong ESG performance may achieve superior financial performance and reduced risk profiles over the long term. This paradigm shift reflects an increasing recognition of the interconnectedness between corporate sustainability, ethical conduct, and financial returns. Consequently, ESG criteria are not only shaping investment strategies but are also influencing corporate behavior and policy-making. This paper seeks to explore the mechanisms through which ESG factors contribute to enhanced corporate performance, the methodologies for assessing ESG criteria, and the implications for investors, corporations, and broader societal outcomes.

Socially Responsible Investing (SRI) represents a subset of the broader ESG framework, emphasizing the ethical implications of investment decisions alongside financial returns. SRI involves the deliberate selection of investments based on both financial performance and social responsibility criteria, aiming to promote positive social change while achieving competitive financial returns. This approach to investing incorporates various strategies, such as negative screening, positive screening, and impact investing, each designed to align investment portfolios with the investor's ethical values and social objectives.

Negative screening, one of the oldest and most widely practiced SRI strategies, involves the deliberate exclusion of companies or entire industries that engage in activities deemed harmful, unethical, or socially undesirable. These typically include sectors like tobacco, alcohol, weapons manufacturing, and fossil fuels, which are often associated with adverse societal and environmental impacts. By excluding such companies from investment portfolios, investors can align their financial decisions with their personal values or institutional mandates, avoiding the support of companies whose practices conflict with their ethical or moral standards. This approach also reduces potential reputational risks associated with investments in controversial industries.

2375-9259/24/$31.00 ©2024 IEEE
DOI 10.1109/ICDMW65004.2024.00070

Fig. 1. Overview of all things ESG: from environmental impact, to social responsibility, and governance practices.

Conversely, positive screening seeks to proactively identify and invest in companies that demonstrate strong ESG performance or contribute to positive social and environmental outcomes. This method rewards companies that adopt sustainable practices, such as reducing carbon footprints, promoting diversity and inclusion, and implementing ethical governance policies. By channeling capital into such companies, positive screening not only supports their growth but also signals to the broader market that sustainability and corporate responsibility are valued by investors. As a result, it encourages other companies to adopt similar practices to attract investment, ultimately promoting a shift toward more sustainable business operations across industries.

Impact investing, a rapidly growing segment of SRI, goes beyond traditional financial metrics to prioritize measurable social and environmental impact. Impact investors aim to fund projects and organizations that address pressing global challenges, such as poverty, climate change, and inequality. This strategy is characterized by a dual focus on achieving financial returns and generating tangible outcomes for society and the environment. The evolution of SRI has been driven by a combination of investor demand, regulatory developments, and growing awareness of the long-term risks associated with unsustainable business practices. Investors are increasingly prioritizing investments that reflect their values and contribute to sustainable development goals. Additionally, regulatory bodies across the globe are implementing guidelines to encourage transparency and accountability in ESG reporting, further integrating SRI into mainstream investment practices.

Research has shown that SRI does not necessarily entail a trade-off between ethical considerations and financial performance. Numerous studies have demonstrated that companies with robust ESG practices often exhibit lower volatility, reduced risk of regulatory penalties, and enhanced reputation, all of which can contribute to superior financial performance. As a result, SRI is gaining recognition as a viable and attractive investment strategy for those seeking to achieve both financial success and positive societal impact.

This paper aims to leverage neurosymbolic AI to enhance SRI, its various strategies and approaches, the empirical evidence supporting its efficacy, and the challenges and opportunities it presents for investors and corporations alike. By examining the intersection of ethical investing and financial performance, this study aims to contribute to the ongoing discourse on the role of finance in fostering a more sustainable and equitable global economy.

In particular, we carry out a preliminary investigation on 𝕏 data to study the reasons for such volatility but also to understand how investors and people in general associate value to ESG. In particular, we collected about 300,000 tweets about ESG and employed state-of-the-art neurosymbolic AI tools to discover online conversation drivers and sentiments around ESG stocks and, hence, gain insights about what makes them valuable.

The remainder of this paper is organized as follows: Section II introduces our data collection methodology; Section III describes the data analysis approach undertaken; Section IV discusses results; finally, Section V offers concluding remarks.

II. Data Collection

We collected our ESG dataset from 𝕏 using the ten most popular hashtags listed below. In particular, we used the 𝕏 Pro API package between 1^{st} August to 1^{st} September 2024.

- `#ESG`: The most general and widely-used hashtag for content related to ESG, encompassing discussions on sustainability, corporate governance, and social responsibility across industries.
- `#SRI`: An hashtag about investing in companies that prioritize ESG criteria alongside financial returns, which is used by individuals and organizations discussing ethical investments, sustainability, and corporate responsibility.
- `#Sustainability`: A broad term used in the context of ESG, this hashtag focuses on environmentally sustainable practices, climate change solutions, and long-term resource management.
- `#CSR`: Corporate Social Responsibility (CSR) highlights corporate efforts to positively impact society, covering social responsibility initiatives like fair labor practices, philanthropy, and environmental projects.
- `#SustainableFinance`: Used for discussions about financing projects and investments that incorporate ESG criteria, with a focus on green bonds, ethical investing, and sustainable economic growth.
- `#GreenEnergy`: Related to clean energy sources such as wind, solar, and hydropower, often tied into conversations about the environmental side of ESG.
- `#ImpactInvesting`: Refers to investments made with the intention of generating positive social and environmental impact alongside a financial return, a growing focus area within ESG investing.
- `#NetZero`: A hashtag related to corporate and government commitments to reducing carbon emissions to net zero by a certain date, key to ESG goals in combating climate change.
- `#CircularEconomy`: This hashtag focuses on designing products and systems that minimize waste and reuse resources, supporting the environmental goals of ESG frameworks.
- `#EthicalInvesting`: Often used in conjunction with discussions on ESG, this term covers investments made based on ethical principles, avoiding companies that harm the environment or society.

Hashtag	Start Date	End Date	Tweet Count
#ESG	01-08-2024	01-09-2024	94,125
#SRI	01-08-2024	01-09-2024	81,403
#Sustainability	01-08-2024	01-09-2024	72,140
#CSR	01-08-2024	01-09-2024	33,088
#SustainableFinance	01-08-2024	01-09-2024	9,327
#GreenEnergy	01-08-2024	01-09-2024	6,103
#ImpactInvesting	01-08-2024	01-09-2024	4,510
#NetZero	01-08-2024	01-09-2024	2,762
#CircularEconomy	01-08-2024	01-09-2024	2,017
#EthicalInvesting	01-08-2024	01-09-2024	1,928
			Total: *307,403*

TABLE I
DISTRIBUTION OF COLLECTED TWEETS WITH RESPECT TO HASHTAGS.

Fig. 2. Word cloud representing the top keywords in the dataset.

Within one month, we collected a total of 1,000,000 tweets. After pre-processing (e.g., removal of irrelevant tweets, removal of duplicates, removal of re-tweets, etc.), we were left with about one third of it. The exact distribution of tweets with respect to hashtags is illustrated in Table I.

Fig. 2 proposes a visual representation of the most significant terms in the collected dataset (after stopword removal), where the size of each word is proportional to its frequency.

III. Data Analysis

In order to gain insights from the collected data, we leverage sentiment analysis, a field of natural language understanding (NLU) [1] in which computational methods are used to determine the polarity or emotional tone expressed in a piece of text. Different AI techniques have been leveraged to improve both accuracy and interpretability of sentiment analysis algorithms, including symbolic AI [2], subsymbolic AI [3], and neurosymbolic AI [4].

Besides traditional algorithms [5] focusing on English text, multilingual [6], [7] and multimodal [8], [9] sentiment analysis have also attracted increasing attention recently. Typical applications of sentiment analysis include social data analytics [10], recommender systems [11], financial forecasting [12], personalization [13], and mental health [14].

In this work, we use Sentic APIs[1], a suite of application programming interfaces available in 80 languages, which employ neurosymbolic AI to perform various sentiment analysis tasks in a fully interpretable manner [15] (Fig. 3). A short description of each API and its usage within this work is provided in the next 12 subsections.

[1]https://sentic.net/api

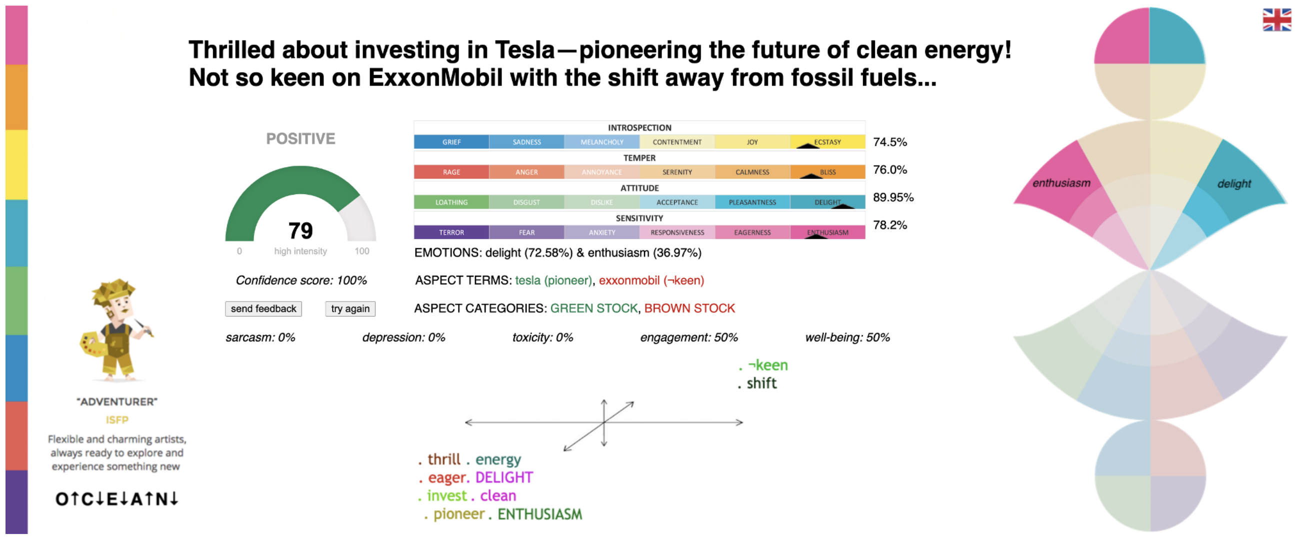

Fig. 3. Sentic API user interface sample.

A. Concept Parsing

This API provides access to Sentic Parser [16], a knowledge-specific concept parser based on SenticNet [17], which leverages both inflectional and derivational morphology for the efficient extraction and generalization of affective multiword expressions from text. In particular, Sentic Parser is a hybrid semantic parser that uses an ensemble of constituency and dependency parsing and a mix of stemming and lemmatization to extract 'semantic atoms' like pain_killer, go_bananas, or get_along_with, which would carry different meaning and polarity if broken down into single words. We use the API for extracting words and multiword expressions from text in order to better understand what are the key concepts related to ESG. As shown in Fig. 3, for example, concepts extracted are `thrill`, `clean`, `shift` and `not keen`.

B. Subjectivity Detection

Subjectivity detection is an important NLU task that aims to filter out 'factual' content from data, i.e., objective text that does not contain any opinion. This API leverages a knowledge-sharing-based multitask learning framework powered by a neural tensor network, which consists of a bilinear tensor layer that links different entity vectors [18]. We use the API to classify ESG-related tweets as either objective (unopinionated) or subjective (opinionated) but also to handle neutrality [19], that is, a tweet that is opinionated but neither positive nor negative (ambivalent stance towards the opinion target). All labels come with a confidence score based on how much SenticNet concepts contributed to the classification output. As depicted in Fig. 3, the confidence score of the proposed example is 100%. Finally, the *Subjectivity Detection module* is also responsible for identifying the language of the input, as indicated in the top-right corner of the UI.

C. Polarity Classification

Once an opinionated tweet is detected using the *Subjectivity Detection API*, the *Polarity Classification API* further categorizes this tweet as either positive or negative. This is one of the most important APIs we use to understand the stance of tweeters towards SRI. It leverages an explainable fine-grained multiclass sentiment analysis method [20], which involves a multi-level modular structure designed to mimic natural language understanding processes. As illustrated in Fig. 3, for example, the extracted polarity is POSITIVE.

D. Intensity Ranking

We also employ the *Intensity Ranking API* to infer the degree of negativity (floating-point number between -100 and 0) or positivity (floating-point number between 0 and 100) of ESG tweets. In particular, the API leverages a stacked ensemble method for predicting sentiment intensity by combining the outputs obtained from several deep learning and classical feature-based models [21]. As shown in Fig. 3, the extracted polarity of the proposed example is 79 (high intensity).

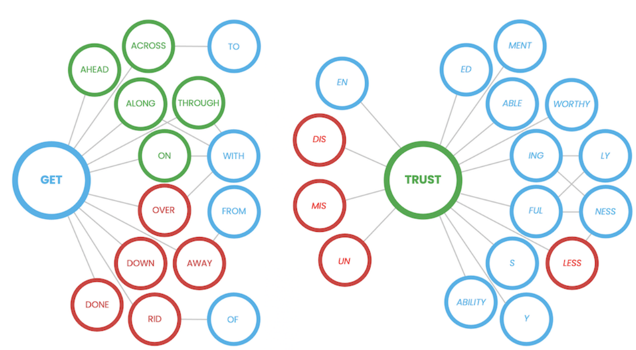

Fig. 4. Sentic Parser graph sample.

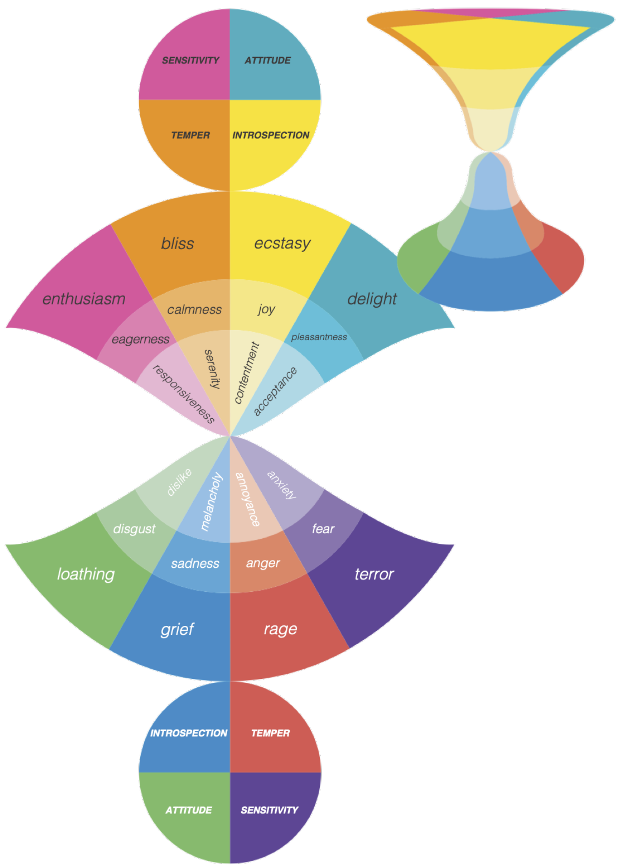

Fig. 5. The Hourglass of Emotions.

E. Emotion Recognition

This API employs the Hourglass of Emotions [22], an emotion categorization model that represents affective states both through labels and through four independent but concomitant affective dimensions describing the full range of emotional experiences that are rooted in any of us (Fig. 5). We leverage the API to move beyond just polarity and intensity, exploring the specific emotions evoked by SRI among both passionate supporters and outspoken critics. As depicted in Fig. 3, for example, the emotion spectrum of the input is visualized in terms of the Hourglass Model's affective dimensions, namely: 74.5% Introspection, 76% Temper, 89.95% Attitude, and 78.2% Sensitivity. From these, the API also extracts the two top resulting emotion labels, delight and enthusiasm, with an intensity of 72.58% and 36.97%, respectively.

F. Aspect Extraction

This API uses a meta-based self-training method that leverages both symbolic representations and subsymbolic learning for extracting aspects from text. In particular, a teacher model is trained to generate in-domain knowledge (e.g., unlabeled data selection and pseudo-label generation), where the generated pseudo-labels are used by a student model for supervised learning.

Then, a meta-weighter is jointly trained with the student model to provide each instance with sub-task-specific weights to coordinate their convergence rates, balancing class labels, and alleviating noise impacts introduced from self-training [23]. We use the API to better understand ESG and SRI in terms of subtopics or opinion targets. Instead of simply identifying a polarity associated with the whole tweet, the *Aspect Extraction API* deconstructs input text into a series of specific aspects or opinion targets to then associate a polarity to each of them. As illustrated in Fig. 3, the opinion targets extracted from the proposed example are `tesla` and `exxonmobil`, which belong to the aspect categories `GREEN STOCK` and `BROWN STOCK`, respectively. The UI also displays the affective concepts most relevant to each aspect term (in brackets) which are colored according to their respective polarities (green for positive and red for negative).

G. Personality Prediction

This API uses a novel hard negative sampling strategy for zero-shot personality trait prediction from text using both OCEAN and MBTI models (Fig. 6). In particular, the API leverages an interpretable variational autoencoder sampler, to pair clauses under different relations as positive and hard negative samples, and a contrastive structured constraint, to disperse the paired samples in a semantic vector space [24]. We use the API to study the different personalities and personas involved in ESG discussions and, hence, better understand the possible drivers of such discussions. As shown in Fig. 3, for example, the MBTI personality extracted is ISFP (Introversion, Sensing, Feeling, and Perceiving) and the OCEAN personality traits extracted are O↑C↓E↓A↑N↓, i.e., high Openness, low Conscientiousness, low Extraversion, high Agreeableness, and low Neuroticism.

H. Sarcasm Identification

This API combines commonsense knowledge [25] and semantic similarity detection methods to better detect and process sarcasm in text. It also employs a contrastive learning approach with triplet loss to optimize the spatial distribution of sarcastic and non-sarcastic sample features [26]. We use the API to understand how much SRI is subject to satire and critique but also to increase the accuracy and reliability of the *Polarity Classification API*. By expressing the opposite of the intended emotion, in fact, sarcasm can cause polarity misclassification. The sarcasm score goes from zero (no sarcasm detected) to 100 (extremely sarcastic content). As depicted in Fig. 3, no sarcasm was detected in the proposed example.

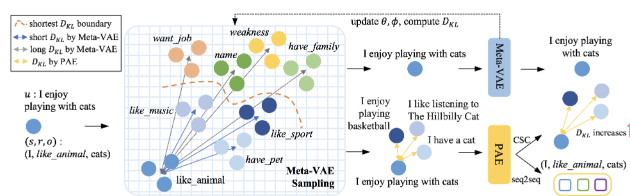

Fig. 6. Personality prediction visualization sample.

I. Depression Categorization

This API employs ensemble hybrid learning methods for automated depression categorization. In particular, the API combines symbolic AI (lexicon-based models) with subsymbolic AI (attention-based deep neural networks) to enhance the overall performance and robustness of depression detection [27]. We use it to study different reactions to green stock devaluations and depreciations by different users, e.g., those experiencing emotional distress or psychological challenges related to their ESG portfolio. The depression score ranges from zero (no depression detected) to 100 (severe depression). In the given example (Fig. 3), no depression was detected

J. Toxicity Spotting

Given the controversy associated with digital assets, it is important to measure the different types and intensities of toxicity associated with some ESG tweets. This API is based on a multichannel convolutional bidirectional gated recurrent unit for detecting toxic comments in a multilabel environment [28]. In particular, the API extracts local features with many filters and different kernel sizes to model input words with long term dependency and then integrates multiple channels with a fully connected layer, normalization layer, and an output layer with a sigmoid activation function for predicting multilabel categories such as 'obscene', 'threat', or 'hate' (Fig. 7). The toxicity score goes from zero (no toxicity detected) to 100 (highly toxic content). As depicted in Fig. 3, no toxicity was detected in the proposed example.

K. Engagement Measurement

Measuring engagement is important to understand which specific topics or events, e.g., ESG regulations, are more impactful for both SRI enthusiasts and skeptics. This API employs a graph-embedding model that fuses heterogeneous data and metadata for the classification of engagement levels.

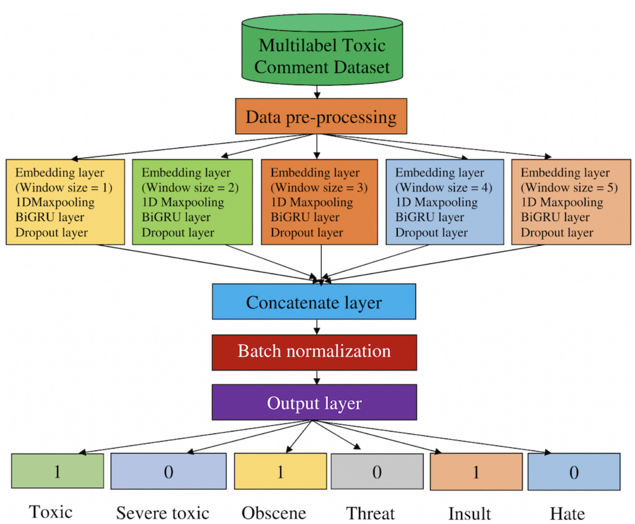

Fig. 7. Toxicity Spotting algorithm.

In particular, the API leverages hybrid fusion methods for combining different types of data in a heterogeneous network by using semantic meta paths to constrain the embeddings [29]. The engagement score ranges from -100 (high disengagement) to 100 (high engagement). As depicted in Fig. 3, for example, the engagement score is 50%.

L. Well-being Assessment

Besides levels of toxicity and engagement, another important dimension for understanding SRI tweeters is their level of stress, e.g., anxiety caused by FOMO. This API leverages a mix of lexicons, embeddings, and pretrained language models for stress detection from social media texts [30]. In particular, the API employs a transformer-based model via transfer learning to capture the nuances of natural language expressions that convey stress in both explicit and implicit manners. The well-being score ranges from -100 (high stress) to 100 (high well-being). As illustrated in Fig. 3, the well-being score is 50% in the proposed example.

IV. Results

In this section, we discuss the most important insights gained through the use of Sentic APIs on the collected dataset. The *Concept Parsing API* enabled us to discover what are the current hot topics related to ESG. Below are the 20 most frequent concepts (words and multiword expressions) parsed.

- esg
- sri
- sustainability
- green_energy
- carbon_emission
- net_zero
- corporate_governance
- renewable_energy
- social_responsibility
- environmental_impact
- ethical_investing
- greenwashing
- climate_action
- esg metric
- carbon_footprint
- climate_crisis
- sustainable_finance
- clean_energy
- stakeholder_capitalism
- impact_investing

Through the *Subjectivity Detection API*, we realized that the vast majority of ESG tweets were opinionated. The unopinionated tweets were mostly promotional and advertising posts. This was further validated by the results of the *Intensity Ranking API*, which were high for both negative and positive spectrum, demonstrating that this is a quite contentious subject. By processing subjective text using the *Polarity Classification API*, we then realized that the number of SRI enthusiasts is far greater than the number of SRI detractors (at least for the time window of our analysis). For the former group, the most common MBTI personality type was ISFP and the predominant emotion was enthusiasm. The latter group (the naysayers), instead, was characterized by an ENTJ personality trait and a predominant emotion of anger.

The *Sarcasm Identification API* has flagged a subtle presence of sarcasm within the context of ESG. Unlike many other topics discussed on social media, ESG does not seem to lend itself well to sarcasm.

The *Depression Categorization API* observed minimal indications of depression concerning SRI. The only depression-related content was detected in tweets concering the upcoming US presidential election, in which users expressed anxiety and depression about the idea that Donald Trump may become president again and negatively affect their ESG portfolio.

The *Toxicity Spotting API* also did not pick up much toxic content. One significant factor is the nature of ESG, which inherently lends itself to more neutral or consensus-based discussions, minimizing the potential for conflict or toxicity. Most individuals approached the topic with openness, curiosity, and a willingness to listen to differing perspectives and, hence, fostered an environment conducive to rather constructive dialogues without personal attacks nor hostility.

The *Engagement Measurement API* exhibited high levels of interest and participation, mostly driven by SRI's relevance and controversy. ESG investing appeals to a growing segment of investors who seek not only financial returns but also a way to influence positive change in society and the environment.

The *Well-being Assessment API* detected medium-high levels of stress by users wondering whether socially responsible investments would perform as well as or better than traditional investments, especially if there is a trade-off between financial returns and ethical principles.

Finally, some very useful insights came from the *Aspect Extraction API*, which helped us individuate the key features of SRI that make them valuable in the eyes of investors and the wider ESG community. Below, we list the ten most frequent aspect categories, together with some aspect terms, along with a short elucubration on why such aspects emerged from the over 300,000 tweets as the most prominent.

- `Environment`: One of the three key aspect categories of ESG, regarding how companies or investments affect the environment, including pollution, carbon footprint, and resource use. Aspect terms related to this category are carbon emissions, waste management, renewable energy use, and climate change mitigation.
- `Social`: One of the three key aspect categories of ESG, regarding a company's initiatives to positively impact society, including philanthropy, community engagement, and fair labor practices. Concepts related to this aspect category are community development, employee welfare, and social contributions.
- `Governance`: One of the three key aspect categories of ESG, regarding the management structure and decision-making processes of a company, including board independence, shareholder rights, and executive compensation. Aspect terms related to this category are board accountability, voting rights, and ethical leadership.
- `Investing`: Another key aspect category concerning the financial performance and ethical alignment of investments, focusing on how capital is allocated to companies that prioritize long-term value creation alongside responsible business practices. Aspect terms related to this category include green stocks, brown stocks, financial returns, risk management, and portfolio diversification.
- `Sustainability`: This aspect category is about long-term environmental and social sustainability practices in a company's operations. Key aspects of this category are sustainable supply chains, product life cycles, and eco-friendly business models.
- `Ethics`: This category relates to discussions for ensuring that businesses follow ethical practices, including honesty, transparency, and integrity in operations. Aspect terms related to this category are transparency in reporting, avoiding corruption, and responsible marketing.
- `Diversity`: This aspect category is about efforts by organizations to promote diversity in the workforce and create an inclusive work environment. Aspect terms related to this category are gender equality, representation of minorities, and inclusive policies.
- `Compliance`: This category is about adherence to laws and regulations relevant to ESG and SRI, including environmental laws, labor laws, and anti-corruption measures. Aspect terms related to this category are legal compliance, anti-corruption policies, and regulatory challenges.
- `Efficiency`: This aspect category relates to the efficient use of resources, including water, energy, and raw materials, as part of an environmental focus. Aspect terms related to this category are energy management, operational efficiency, and resource optimization.
- `Engagement`: This category is about how well a company interacts with its stakeholders, including shareholders, customers, employees, and the community. Aspect terms related to this category are transparency, communication, and responsiveness to stakeholder concerns.

V. CONCLUSION

ESG and SRI are gaining importance for their role in guiding ethical, sustainable, and responsible innovation. As AI increasingly influences decision-making across industries, it is essential to ensure that companies use AI technologies ethically, considering data privacy, fairness, and societal impacts. ESG principles provide a framework for ensuring transparency, accountability, and reducing risks associated with AI, such as bias, discrimination, or job displacement. The social aspect of ESG is crucial for managing the societal impacts of AI, including the potential for labor disruption and inequality. Strong governance ensures companies use AI ethically and align with emerging regulations, while environmental considerations push for sustainable AI practices, such as reducing the energy consumption of large-scale AI models.

ESG and SRI also play a significant role in building trust and accountability. Companies that prioritize these frameworks are more likely to earn public trust and avoid scandals related to unethical AI use. Investors are increasingly applying ESG criteria to their portfolios, driving funding toward companies that demonstrate responsible AI development, innovation, and alignment with societal goals. Furthermore, AI can enhance ESG reporting by providing real-time insights into a company's environmental, social, and governance performance, improving transparency for investors.

Fig. 8. Integration of AI and ESG for the future of humanity.

In this study, we gathered about 300,000 ESG tweets and utilized neurosymbolic AI to enhance our comprehension of the factors influencing online conversations and sentiments related to SRI. Our objective was to glean insights into the factors contributing to their perceived value. Through Sentic APIs we discovered the following ten key aspect categories: environment, social, governance, investing, sustainability, ethics, diversity, compliance, efficiency, and engagement. Future work will focus on conducting controlled experiments to investigate how such categories, and their relative aspect terms, influence the perceived value of ESG stocks.

In this brave new world of AI, companies that integrate ESG and SRI principles are better positioned to meet consumer and investor demands for ethical AI, avoid regulatory risks, and gain a competitive market advantage. These frameworks help ensure AI is used not only for profit but for the benefit of society and the environment.

Acknowledgments

This research/project is supported by the Ministry of Education, Singapore under its MOE Academic Research Fund Tier 2 (STEM RIE2025 Award MOE-T2EP20123-0005) and by the RIE2025 Industry Alignment Fund – Industry Collaboration Projects (IAF-ICP) (Award I2301E0026), administered by A*STAR, as well as supported by Alibaba Group and NTU Singapore.

References

[1] E. Cambria, *Understanding Natural Language Understanding*. Springer, ISBN 978-3-031-73973-6, 2024.

[2] Q. Liu, S. Han, Y. Li, E. Cambria, and K. Kwok, "PrimeNet: A framework for commonsense knowledge representation and reasoning based on conceptual primitives," *Cognitive Computation*, 2024.

[3] I. Chaturvedi, Y.-S. Ong, I. Tsang, R. Welsch, and E. Cambria, "Learning word dependencies in text by means of a deep recurrent belief network," *Knowledge-Based Systems*, vol. 108, pp. 144–154, 2016.

[4] E. Cambria, R. Mao, M. Chen, Z. Wang, and S.-B. Ho, "Seven pillars for the future of artificial intelligence," *IEEE Intelligent Systems*, vol. 38, no. 6, pp. 62–69, 2023.

[5] E. Cambria, B. Schuller, B. Liu, H. Wang, and C. Havasi, "Statistical approaches to concept-level sentiment analysis," *IEEE Intelligent Systems*, vol. 28, no. 3, pp. 6–9, 2013.

[6] D. Vilares, H. Peng, R. Satapathy, and E. Cambria, "BabelSenticNet: A commonsense reasoning framework for multilingual sentiment analysis," in *IEEE SSCI*, 2018, pp. 1292–1298.

[7] H. T. Nguyen, P. H. Duong, and E. Cambria, "Learning short-text semantic similarity with word embeddings and external knowledge sources," *Knowledge-Based Systems*, vol. 182, no. 104842, 2019.

[8] E. Cambria, N. Howard, J. Hsu, and A. Hussain, "Sentic blending: Scalable multimodal fusion for continuous interpretation of semantics and sentics," in *IEEE SSCI*, Singapore, 2013, pp. 108–117.

[9] L. Stappen, A. Baird, E. Cambria, and B. Schuller, "Sentiment analysis and topic recognition in video transcriptions," *IEEE Intelligent Systems*, vol. 36, no. 2, pp. 88–95, 2021.

[10] E. Cambria, D. Rajagopal, D. Olsher, and D. Das, "Big social data analysis," in *Big Data Computing*, R. Akerkar, Ed. Chapman and Hall/CRC, 2013, ch. 13, pp. 401–414.

[11] P. Kazienko and E. Cambria, "Towards responsible recommender systems," *IEEE Intelligent Systems*, vol. 39, no. 3, pp. 5–12, 2024.

[12] L. Malandri, F. Xing, C. Orsenigo, C. Vercellis, and E. Cambria, "Public mood–driven asset allocation: the importance of financial sentiment in portfolio management," *Cognitive Computation*, vol. 10, no. 6, pp. 1167–1176, 2018.

[13] L. Zhu, R. Mao, E. Cambria, and B. J. Jansen, "Neurosymbolic AI for personalized sentiment analysis," in *Proceedings of HCII*, 2024.

[14] S. Ji, X. Li, Z. Huang, and E. Cambria, "Suicidal ideation and mental disorder detection with attentive relation networks," *Neural Computing and Applications*, vol. 34, pp. 10 309–10 319, 2022.

[15] E. Cambria, A. Hussain, C. Havasi, and C. Eckl, "Sentic computing: Exploitation of common sense for the development of emotion-sensitive systems," ser. LNCS. Springer, 2010, vol. 5967, pp. 148–156.

[16] E. Cambria, R. Mao, S. Han, and Q. Liu, "Sentic parser: A graph-based approach to concept extraction for sentiment analysis," in *Proceedings of ICDM Workshops*, 2022, pp. 413–420.

[17] E. Cambria, X. Zhang, R. Mao, M. Chen, and K. Kwok, "SenticNet 8: Fusing emotion AI and commonsense AI for interpretable, trustworthy, and explainable affective computing," in *International Conference on Human-Computer Interaction (HCII)*, 2024.

[18] R. Satapathy, S. Rajesh Pardeshi, and E. Cambria, "Polarity and subjectivity detection with multitask learning and bert embedding," *Future Internet*, vol. 14, no. 7, p. 191, 2022.

[19] A. Valdivia, V. Luzón, E. Cambria, and F. Herrera, "Consensus vote models for detecting and filtering neutrality in sentiment analysis," *Information Fusion*, vol. 44, pp. 126–135, 2018.

[20] Z. Wang, Z. Hu, S.-B. Ho, E. Cambria, and A.-H. Tan, "MiMuSA—Mimicking human language understanding for fine-grained multi-class sentiment analysis," *Neural Computing and Applications*, vol. 35, no. 21, pp. 15 907–15 921, 2023.

[21] M. S. Akhtar, A. Ekbal, and E. Cambria, "How intense are you? Predicting intensities of emotions and sentiments using stacked ensemble," *IEEE Computational Intelligence Magazine*, vol. 15, no. 1, pp. 64–75, 2020.

[22] Y. Susanto, A. Livingstone, B. C. Ng, and E. Cambria, "The Hourglass Model revisited," *IEEE Intelligent Systems*, vol. 35, no. 5, pp. 96–102, 2020.

[23] K. He, R. Mao, T. Gong, C. Li, and E. Cambria, "Meta-based self-training and re-weighting for aspect-based sentiment analysis," *IEEE Transactions on Affective Computing*, vol. 15, 2024.

[24] L. Zhu, W. Li, R. Mao, V. Pandelea, and E. Cambria, "PAED: Zero-shot persona attribute extraction in dialogues," in *ACL*, 2023, pp. 9771–9787.

[25] E. Cambria, A. Hussain, C. Havasi, and C. Eckl, "Common sense computing: From the society of mind to digital intuition and beyond," in *Biometric ID Management and Multimodal Communication*, ser. Lecture Notes in Computer Science. Berlin Heidelberg: Springer, 2009, vol. 5707, pp. 252–259.

[26] T. Yue, R. Mao, H. Wang, Z. Hu, and E. Cambria, "KnowleNet: Knowledge fusion network for multimodal sarcasm detection," *Information Fusion*, vol. 100, p. 101921, 2023.

[27] L. Ansari, S. Ji, Q. Chen, and E. Cambria, "Ensemble hybrid learning methods for automated depression detection," *IEEE Transactions on Computational Social Systems*, vol. 10, no. 1, pp. 211–219, 2023.

[28] J. A. Kumar, S. Abirami, T. E. Trueman, and E. Cambria, "Comment toxicity detection via a multichannel convolutional bidirectional gated recurrent unit," *Neurocomputing*, vol. 441, pp. 272–278, 2021.

[29] I. Chaturvedi, K. Thapa, S. Cavallari, E. Cambria, and R. E. Welsch, "Predicting video engagement using heterogeneous deepwalk," *Neurocomputing*, vol. 465, pp. 228–237, 2021.

[30] A. Rastogi, Q. Liu, and E. Cambria, "Stress detection from social media articles: New dataset benchmark and analytical study," in *IJCNN*, 2022.