International Journal of Emerging Research in Engineering and Technology
Pearl Blue Research Group| Volume 3, Issue 4, 128-132, 2022
ISSN: 3050-922X | https://doi.org/10.63282/3050-922X.IJERET-V3I4P114

Original Article

AI-Driven Portfolio Optimization for ESG Investment

Venkata Sai Nageen Kanikanti
Director, Software Engineering, Fidelity Investments.

Abstract - *Environmental Social and Governance (ESG) considerations are fast becoming part and parcel of investment strategy. Meanwhile, artificial intelligence (AI) techniques, namely, natural language processing (NLP) and large language models (LLMs) are opening up the field to in-depth analysis of unstructured corporate disclosures and financial texts. This paper gives a model of an AI-based portfolio optimization taking into account the ESG factors becoming a direct part of the decision-making process. The methodology combines systematic historical market data together with unstructured textual evidence taken out of earnings calls, regulatory filings, and sustainability reports. The pipeline includes three major components: 1. ESG signal extraction by means of domain-adapted transformers; 2. Structured fact cogeneration and summarization with LLMs; 3. Multi-objective portfolio optimization that incorporates balancing returns, risk, and ESG alignment. The model is set with two conference contexts in view. In the track of ICDM, the task is to mine ESG-related information on earnings calls and 10-K reports using mining based on NLP. In DSAA track, the interest is in the use of LLMs to process financial documents to enable ESG assessment that can be audited. The framework is grounded conceptually in Adewale (2024), which argues for integrating ESG criteria in AI-based portfolio management, and in Rane, Choudhary, and Rane (2024), which surveys AI approaches that strengthen ESG in sustainable business practices. The proposed solution shows the promise of AI-based source of investments in giving sustainability-focused investing an edge on appropriate textual evidence to be converted into explainable and actionable portfolio allocation signals.*

Keywords - *Environmental, Social, And Governance (ESG), Artificial Intelligence (AI), Portfolio Optimization, Natural Language Processing (NLP), Large Language Models (Llms), Sustainable Investing, Financial Text Mining, Earnings Call Analysis, Regulatory Filings (10-K), ESG Signal Extraction, Multi-Objective Optimization.*

1. Introduction

According to the Roger, (2024), Sustainable investing has developed far beyond a negative screening into an array of approaches including best-in-class stock selection and portfolio applications of ESG factors to traditional portfolio theory. Institutional investors are increasingly of the view that ESG considerations are an ethical and economic necessity. In parallel with ESG integration, challenges are quite high. Disclosures are usually not uniform, incomplete and voluntary. One of the study's researchers stated that, different agencies that rate companies on the basis of ESG often disagree and this leaves portfolio managers uncertain on which ratings are valid and matter about financial performance. The challenges can be addressed through the AI techniques. Models of NLP trained on large quantities of writing may be very effective at identifying patterns, themes, and signals within corporate communications that could otherwise be difficult to discern manually. LMs can summarise and organise disclosures, and machine learning can pair ESG features to forward-looking estimates of returns and risk.

The insights gained using AI when paired with portfolio optimization structures can shove the pendulum of capital allocation to be used on a more sustainable investment without jeopardizing financial output one of the study's researchers. This study seeks to develop a holistic picture of AI-based portfolio optimization in the cases of integrating ESG issues. The manuscript elaborates on the methodological aspects with references to data mining and analytics applications to the Conference audiences, as well as the evaluation factors.

2. Background

Paul, (2023) stated that, Portfolio optimization traditionally approach to portfolio optimization employs the mean-variance criterion, a process in which investor's trade off expected returns with risks. More recently One of the study's researchers stated that, factor models have been used to generalize on this foundations, explaining the differences in returns in terms of exposure to factors like size, value, and momentum. The further addition of non-financial goals in SG integration complicates the situation by bringing in non-financial factors that might not be aligned with conventional factors. In the past ESG data was based on third-party ratings. These ratings are, however, marred by inter-provider disagreement and by backward-looking bias (Nair, 2024). As an example, ratings tend to record past performance or disclosures as opposed to future commitments. This has led to a new surge in the search of alternative data sources that can capture emerging ESG information directly out of the company's own communications.

Earnings calls and regulatory filings can be a very valuable source. The calls record managerial informal remarks and analyst questions that have the potential to project the priorities, tonalities, and trustworthiness of the company. Reporting can provide information of environmental or governance issues through disclosure of risk and management discussions (10-K and 20-F reports). Sustainability reports and climate-related disclosures can elaborate on commitments and progress even further. It is here that the need arises to consider mining these disparate sources, providing highly advanced levels of AI usage, particularly within the use of unstructured, highly wordy, industry-specific text (Paul, 2023). The need to take into consideration AI-powered ESG inputs in the process of optimisation of portfolios is where sustainable finance begins to emerge. It particularly consists in the conversion of unstructured, qualitative cues into numerical indicators, and their insertion into optimization software.

3. Problem Formulation

The core issue is to build a portfolio that can maximize long-term risk-adjusted returns and at the same time satisfy investor preferences on ESG alignment. This will need the quantification of the ESG-related signals out of the unstructured data and their incorporation into the optimization scheme (Nair, 2024; Sun et al., 2024). Assuming a universe of financial assets with, among other data: returns, volatility information, and possibly a set of ESG indicators provided by training tasks of AI models. The optimization problem tries to find a weighting of a portfolio that maximizes the returns and simultaneously minimizes the risk and at the same time aligning with the preferences of the investor based on the Environmental, Social, and Governance (ESG) and finally controlling the transaction costs One of the Researcher. The ESG alignment measure is calculated on the basis of sustainable objectives and indicators extracted using AI and shows the extent to which the portfolio aligns with the objectives.

The formulation given is a generalization of the classical portfolio optimization in the sense that it does not merely impose a constraint on ESG but rather introduces the latter as a component (objective) of the problem. The method enables investors to put relative weightings on the ESG and the conventional financial objectives, and provides a portfolio that is economically appropriate while being socially responsible.

4. Data Sources

A robust AI-driven ESG system relies on multiple data sources:

- Financial Data: Prices, returns, corporate actions, and fundamentals serve as the quantitative backbone of the model.
- Corporate Communications: Earnings call transcripts provide forward-looking commentary and managerial tone.
- Regulatory Filings: 10-K and 20-F reports contain detailed risk disclosures and management discussions.
- Sustainability Reports: Voluntary disclosures highlight environmental, social, and governance initiatives.
- News and Controversies: Media reports provide external perspectives on company ESG practices and controversies.

By combining these sources, the system constructs a comprehensive view of ESG performance and prospects for each company.

5. ESG Signal Extraction through NLP

The first layer of the framework is concerned with ESG-relevant signal extraction of unstructured texts. Modern transformer-based NLP models may also be trained to classify text passages into ESG categories like climate risks, quality of governance or labor practices (Kannan & Seki, 2023; One of the Researcher. These models are able to pick up both the obvious statements and turn of phrases, as well as implicit clues of management responses, i.e. degree of certainty or specificity. Earnings calls are coded at the speaker turn level with a distinction between the two parties (the functions of an executive and an analyst). Executives can make forward-looking statements with the help of ESG goals commitments, and analysts may provide questions reflecting doubts or concerns. SEC filings are scanned text line-by-line with special attention to risk factor and management discussion sections one of the study's researchers. The end product is a table of structured ESG sub-scores per company and time with provenance behind each total score referring back to the passages. This layer constitutes a quantitative representation of the qualitative text that can be integrated into returns and risk prediction models.

6. LLM-Assisted Document Processing

As demonstrated by Perazzoli et al., (2022), the use of large language models in ESG-based portfolio optimization is an indication of a major transition away in the use of keyword-centric text mining to a more contextualized understanding of documents. In this layer, LLMs are applied to analyse unstructured financial documents e.g. annual reports, sustainability disclosures and regulatory filings. Rather than restrict the analysis to sentiment or frequency-based analysis, the models are able to produce naturally structured ESG insights due to the identification of key themes such as commitments to act on climate change, corporate governance frameworks, and diversity programmes (Raman et al., 2020; Castelberg & Flury, 2024). Heavily augmented retrieval strategy can keep the deployed information well-grounded in retrievable portions of texts, and the method is thus more transparent and reliable.

The propensity to hallucination is one of the major problems of employing generative models on a financial task because they may not match with underlying information. As a means of countering this fact, there is programmatic verification of extracted facts applied by the system (Castelberg & Flury, 2024). Quantitative items, like emission reduction goals, or proportions of the work force, are just counted and verified that they appear to form a coherent whole, and qualitative results are triangulated against passages retrieved. The two checks will impose error-free precision and clarity, which are essential to investor confidence. By tethering every generated summary to real evidence, the system makes the auditability of the ESG lifecycle possible One of the investigators.

Finally, the resulting extraction layer as driven by the LLM we have considered expands upon the existing sub-scores which incorporate NLP-based elements by extending the signal that is captured (Joshi et al., 2024). It allows the system to go further than analyzing text on the surface level and generate facts that can be understood by investors and incorporated readily into the decision-making process with an evidence-base. This combination of machine reasoning and verifiable context promotes stewardship objectives, and portfolio construction in line with the sustainability imperatives.

7. From Signals to Financial Predictions

After the ESG information has been collected and organized it is time to connect the ESG signal with financial performance outcomes. The mechanism that connects the notions of sustainability features with traditional notions of risk-return consideration is the predictive modelling One of the Researcher. Machine learning models within this context combine ESG-derived metrics and the traditional signals of market capitalization ratio, leverage ratio, and industry-based benchmarks (Zhang & Zhang, 2024). The outcome is the system that has the ability to generate prospective measures of returns and risks that consider both financial fundamentals and sustainability exposures.

The effect of ESG on financial performance can occur in a variety of ways. Directly, ESG can deliver strong profiles that can attract upward price pressure or lower cost of capital against rising investor interest in sustainable investments (Talapaneni, 2024). Indirectly, companies with good governance or active environs policy may be less exposed to reputational and regulatory risks and as such, the cash flows are less at risk. In contrast, business firms with weak governance or those, which bring with them massive environmental risk, can be more volatile, amplifying the downside (Ruberg, 2021). These dynamics explain why it is not possible to dismantle ESG factors out of portfolio considerations.

To make these findings usable, predictive models eliminate confounding factors like industry sector, macroeconomic trends and regional policies. This will make ESG effects be distinguished without the influence of external factors (Talapaneni, 2024). The results of these models can give investors forecasts of risk-adjusted performance in which sustainability is clearly quantified. The incorporation of ESG in the development of predictive financial models can, therefore, promote two colinear objectives: maximizing returns and aligning with long- run market and environmental objectives. By doing that, portfolio optimization as an exclusively financial exercise no longer can exist: the responsibility is now in the very heart of the investment strategy framework.

8. Multi-Objective Portfolio Optimization

With efficient estimates of returns and risk, the portfolio optimization step is developed to trade-off between multiple goals in a formal, rule based manner. Konieczny, (2024) stated that, the process extends the conventional mean-variance models because it incorporates the aspect of sustainability to the optimization challenge. It has built a composite objective function that combines the three dimensions of risk-adjusted financial return, ESG alignment and transaction cost minimization (Taraj, 2024). ESG alignment is numerically displayed by comparing the cumulative ESG characteristic of the portfolio to specific and predetermined investor-defined metrics, which may include carbon emission reductions, governance goals or social assistance marks.

This optimization remains a relatively easy process through which investors may define a relative importance on a dimension. Consider, as an example, that some investors can be more concerned with carbon neutrality than others who value financial outperformance compared to strict adherence to ESG standards (Chen et al., 2020). It is possible to add constraints as per the investment policy such as sector exposure limits, minimum ESG rating constraints or liquidity constraints. With such constraints the optimizer would ensure that the final portfolio would be practical and consistent with real-world issues (Wang, 2024). The result is a set of portfolios, which combine financial prudence with sustainability requirements. Rather than trying to apply ESG factors as an overlay, it integrates them into decision making at the ground-level. Its holistic approach not only facilitates further convergence on investor values but also focuses portfolios toward being hedged against long-term risks and in a prime position to take the emerging opportunities in markets where sustainability is playing an increasingly prominent role.

9. Evaluation Framework

According to the Wang, (2024), the evaluation framework would provide a whole-of-system lens to evaluate the quality and impacts of the system regarding several aspects. On the lower rate, the NLP layer is evaluated to determine its capacity to distinguish ESG text and calibrate probabilities in a way that both reflects precision and recall. Extending on this, fidelity verified with the structured facts drawn can substantiate the reliability of the LLM layer and the consistency of the outputs with the named evidence can be ensured (Taraj, 2024). Beyond text processing, the process of evaluation then proceeds to the predictive modeling phase whereby the incremental predictive power of ESG features is tested. This comprises an evaluation of the incremental usefulness of the ESG variables in explaining/predicating the financial returns and risks, compared to those factors (Ruberg, 2021). Outcomes provided at the portfolio level are both financial and ESG-related measures, such as the level of returns, volatility, the Sharpe ratio, drawdown, carbon intensity, and controversy exposure.

A stress test is conducted under such routes to investigate the effect of environmental crises or governance scandals on the system. This assures that portfolios will be resilient to not just the financial impacts of volatility but also on conviction-driven shocks as well. When integrating this multilayered assessment framework, the framework allows the development of ESG data extraction and modeling to translate directly into credible, reliable, and actionable investment results.

10. Engineering and Implementation

The system is made in the form of a modular pipeline. Point-in-time integrity is facilitated by the ingestion of data so that its use of future information is prohibited. Kannan & Seki, (2023) stated that, the NLP models are used to convert the documents into ESG sub-scores but LLMs will produce structured facts. In a form of feature repository files with lineage tracking and version control, the outputs are saved. The optimization procedures lend themselves to portfolio weights that would be tested during backtests or used in practice One of the investigators. The continuous monitoring would identify the drift in the performance of the model as it is exposed to new disclosures. Flagged cases are reviewed by human analysts as a check against interpretability and resulting accountability. The system, therefore, balances computerization and control and comes into harmony with the technical and regulatory soundness.

11. Positioning Relative to Adewale (2024) and Rane et al. (2024)

Adewale (2024) argues that integrating ESG criteria directly into AI-based portfolio management can enhance performance and risk control when ESG is treated as a first-class objective rather than an after-the-fact screen. Our framework operationalizes this by embedding ESG alignment into the optimization objective, not merely as a constraint, and by using AI to generate timely, company-specific signals from disclosures, thus reducing reliance on lagging ratings. Rane, Choudhary, and Rane (2024) review AI approaches that strengthen ESG practice, emphasizing that advanced analytics can standardize and scale ESG assessment while improving transparency. Our work leverages that insight to construct a practical LLM-assisted pipeline with programmatic verification, linking unstructured text to structured ESG facts that can be audited. In combining these perspectives, the manuscript delivers a blueprint that is methodologically rigorous and purpose-built for institutional use.

12. ICDM Focus: NLP for Earnings Calls and Filings

In the domain of ICDM, ESG-relevant data is researched by mining it out of canonical financial texts. Examples are calibrated ESG classifiers, task-specific updated language models and contrastive learning related to passages and the ESG taxonomies (Castelberg & Flury, 2024). The experiments show that the increase in the precision of text mining allows increasing the predictive strength of the signals and, ultimately, in portfolios.

13. DSAA Focus: LLMs for Financial Document Processing

To the audience of the DSAA, the new aspect will be the use of LLMs to financial disclosures. The pipeline exemplifies retrieval-enhanced generation with citations enforcing, schema organized retrieval, and formal validation of the programming kind. Findings indicate the enhancement of reliability and interpretability of ESG signals by the test of LLM-generated facts (Castelberg & Flury, 2024). The study is not only methodological but also practical because it presents the instruments of auditability and transparency of sustainable investments.

14. Limitations

There are limitations to the framework The framework is not unlimited The framework has shortcomings The way that the corporates report may present selectivity and PR propaganda as opposed to actual performance (Zhang & Zhang, 2024). The algorithms could misunderstand or exaggerate on small data. Backtests are never able to capture trading expenses completely as well as the behavior of investors. In addition, the definition of ESG differs across markets and makes it more challenging to implement it in different jurisdictions. These shortcomings make it all the more necessary to interpret it carefully and improve it constantly.

15. Future Directions

This conceptual framework could be expanded going forward to fixed income and/or other types of investments as well. The causal inference techniques can assist in distinguishing between correlation and causation on ESG-return relationships One of the Researcher. Stress testing would benefit by integration with climate scenario analysis. Roam and growth of open datasets and benchmarks will aid reproducibility and progress in the community.

16. Conclusion

In conclusion, this study works out a comprehensive framework of an AI-based portfolio optimization with ESG integration. The methodology works by using NLP-driven signal extraction, LLM-driven document processing and solving the proposed multi-objective optimization problem to convert unstructured corporate disclosures into structured investment signals that can be explained to the end user. It furthers the idea of sustainable investing by having ESG as a core part of the portfolio goals and not an un-prioritized consideration. Tailored to data mining and analytics conferences, the framework shows how at scale AI can align both financial and sustainability performance and, in this way, a means of capital allocation is possible that is more knowledgeable and more responsible.

References

[1] Adewale, T. (2024). Integration of ESG (Environmental, Social, Governance) Criteria into AI-Based Portfolio Management.

[2] Castelberg, D., & Flury, F. (2024). *Llm assisted development* (Doctoral dissertation, OST Ostschweizer Fachhochschule).

[3] Chen, C. C., Huang, H. H., & Chen, H. H. (2020). NLP in FinTech applications: past, present and future. *arXiv preprint arXiv:2005.01320*.

[4] Joshi, R., Bubna, Y., Sahana, M., & Shruthiba, A. (2024, November). An Approach to Intelligent Information Extraction and Utilization from Diverse Documents. In *2024 8th International Conference on Computational System and Information Technology for Sustainable Solutions (CSITSS)* (pp. 1-5). IEEE.

[5] Kannan, N., & Seki, Y. (2023). Textual evidence extraction for ESG scores. In *Proceedings of the Fifth Workshop on Financial Technology and Natural Language Processing and the Second Multimodal AI For Financial Forecasting* (pp. 45-54).

[6] Konieczny, K. (2024). *Wisdom from Words: Evidence from Earnings Conference Calls and Analysts' Research Reports using State-of-the-Art Natural Language Processing Methodologies* (Doctoral dissertation, Technische Universität München).

[7] Nair, V. (2024). AI-powered investment strategies: Enhancing portfolio management through machine learning. *Journal Of Recent Trends in Computer Science and Engineering (JRTCSE)*, *12*(1), 1-5.

[8] Paul, C. (2023). The Role of AI in Creating Sustainable and ESG-Focused Hedge Fund Portfolios.

[9] Perazzoli, S., Joshi, A., Ajayan, S., & de Santana Neto, J. P. (2022). Evaluating environmental, social, and governance (esg) from a systemic perspective: An analysis supported by natural language processing. *Social, and Governance (ESG) from a Systemic Perspective: An Analysis Supported by Natural Language Processing (October 11, 2022)*.

[10] Raman, N., Bang, G., & Nourbakhsh, A. (2020). Mapping ESG trends by distant supervision of neural language models. *Machine Learning and Knowledge Extraction*, *2*(4), 453-468.

[11] Rane, N., Choudhary, S., & Rane, J. (2024). Artificial intelligence driven approaches to strengthening Environmental, Social, and Governance (ESG) criteria in sustainable business practices: a review. *Social, and Governance (ESG) criteria in sustainable business practices: a review (May 27, 2024)*.

[12] Roger, J. D. D. (2024). AI-driven financial modeling techniques: transforming investment strategies. *The Journal of Applied Business and Economics*, *26*(4), 63-74.

[13] Ruberg, N. (2021). Bert goes sustainable: an NLP approach to ESG financing.

[14] Sun, Z., Satapathy, R., Guo, D., Li, B., Liu, X., Zhang, Y., ... & Goh, R. S. M. (2024, December). Information extraction: Unstructured to structured for esg reports. In *2024 IEEE International Conference on Data Mining Workshops (ICDMW)* (pp. 487-495). IEEE.

[15] Talapaneni, A. (2024). Enhancing ESG Reporting and SDG Mapping Using Natural Language Processing.

[16] Taraj, I. (2024). Financial and Environmental Narratives in Earnings Calls: Investor Processing Costs and Market Reactions.

[17] Wang, Y. (2024). NLP applications—clinical documents. In *Natural Language Processing in Biomedicine: A Practical Guide* (pp. 325-349). Cham: Springer International Publishing.

[18] Zhang, A. Y., & Zhang, J. H. (2024). Renovation in environmental, social and governance (ESG) research: the application of machine learning. *Asian Review of Accounting*, *32*(4), 554-572.