Investment Portfolio Optimization and ESG Factors Impact on Financial Performance using Machine Learning

MSc Research Project
MSc in Data Analytics

Akansha Sharma
Student ID: x23194961

School of Computing
National College of Ireland

Supervisor: Prof. William Clifford

National College of Ireland
Project Submission Sheet
School of Computing

Student Name:	Akansha Sharma
Student ID:	x23194961
Programme:	MSc in Data Analytics
Year:	2024
Module:	MSc Research Project
Supervisor:	Prof.; William Clifford
Submission Due Date:	12/12/2024
Project Title:	Investment Portfolio Optimization and ESG Factors Impact on Financial Performance using Machine Learning
Word Count:	8453
Page Count:	23

I hereby certify that the information contained in this (my submission) is information pertaining to research I conducted for this project. All information other than my own contribution will be fully referenced and listed in the relevant bibliography section at the rear of the project.

ALL internet material must be referenced in the bibliography section. Students are required to use the Referencing Standard specified in the report template. To use other author's written or electronic work is illegal (plagiarism) and may result in disciplinary action.

Signature:	Akansha Sharma
Date:	29th January 2025

PLEASE READ THE FOLLOWING INSTRUCTIONS AND CHECKLIST:

Attach a completed copy of this sheet to each project (including multiple copies).	☐
Attach a Moodle submission receipt of the online project submission, to each project (including multiple copies).	☐
You must ensure that you retain a HARD COPY of the project, both for your own reference and in case a project is lost or mislaid. It is not sufficient to keep a copy on computer.	☐

Assignments that are submitted to the Programme Coordinator office must be placed into the assignment box located outside the office.

Office Use Only	
Signature:	
Date:	
Penalty Applied (if applicable):	

Investment Portfolio Optimization and ESG Factors Impact on Financial Performance using Machine Learning

Akansha Sharma
x23194961

Abstract

This study aims to examine the effects of incorporating ESG factors in the investment portfolio and improving the performance of the portfolio using machine learning. The rationale for the study arises from the increasing relevance of sustainable investment and the challenges of integrating ESG factors into investment decision-making to increase returns and manage risks profitably. By applying deep learning technologies, such as LSTM and SVM, the paper contrasts the efficiency of ESG portfolios with standard ones.

The study explores that comparable financial results can be achieved while aiding sustainability goals with ESG-integrated portfolios. The integration of ESG factors in the portfolio showed a Sharpe ratio of 0.538 and balanced performance metrics, close to the Sharpe ratio (0.541) of the non-ESG portfolio. Portfolios optimized for specific ESG components, such as Environmental Risk (Sharpe Ratio: 0.505; ESG Score: 0.841) provided a better fit to sustainability criteria while slightly lower financial parameters. Therefore, these results demonstrate the effectiveness of ESG factors in improving investment approaches and providing optimized results regarding the Sharpe ratio, volatility, and maximum drawdown.

Consequently, it becomes possible to tap into the potential of machine learning in expanding the ESG portfolio while delivering significant improvements in ethically and financially attainable results. Hence, for the creation of comprehensive, efficient, responsible, and sustainable investment practices, this research contributes to ESG integration frameworks. Future work may expand on these results by exploring broader data sets and more complex techniques to define portfolio optimization.

1 Introduction

1.1 Overview

Portfolio management is a crucial element of financial decision-making, which is mainly focused on increasing returns on investment whilst reducing the impacts of the risks. In recent years, the incorporation of ESG factors in the investment management process has become recognized as one of the groundbreaking concepts in finance. ESG factors are an extended list of factors for analyzing the firms, beyond mere financial performances; it also involves an ethical standpoint. This change is stimulated by the findings highlighting that the integration of ESG factors may result in higher long-term financial performance and lower risk associated with irresponsible business activities (Friede et al. (2015), Torricelli (2024)).

However, the inclusion of ESG characteristics in the optimization of a portfolio is still problematic. This is mainly because ESG data includes both qualitative and quantitative information that is difficult to model using mainstream financial modes Nundlall and Van Zyl (2023). Hence, the current research lacks exploration of the potential of modern approaches like machine learning in an integrated system. Machine learning provides tools for analysis of big datasets, and discovering relationships that are not apparent in the presented data Lakshminarayanan and McCrae (2019).

This research aims to fill this gap through examining the use of LSTM, SVM and, Heuristic models for the implementation of ESG factors in portfolio optimization. This work will benefit academic researchers who are interested in developing more effective methodologies for sustainable finance, and financial institutions that wish to adopt more effective and responsible investment models.

1.2 Research Questions and Objectives

This study aims to address the following research questions:

1. How investment portfolios be optimized using machine learning techniques that include ESG factors to improve financial performance?

2. Which is the most effective machine learning technique for optimization of investment portfolios, when considering ESG factors?

3. What is the effect of ESG factor integration on the financial performance of investment portfolios, as when compared with traditional optimization methods?

4. How do different ESG factors impact the risk-return profile of investment portfolios, individually and collectively over time?

The objectives of this research are to optimize portfolios with and without ESG integration by developing and implementing machine learning models, assessing these portfolios financial performance using metrics such as Sharpe Ratio, returns, volatility, and ESG scores. Also, the effect of individual Environmental, Social and, Governance factors on the performance of the portfolio will be explored and the trade-off between financial returns and sustainability will be determined.

The research will use performance metrics like volatility, Sharpe Ratios and risk-adjusted returns for evaluating whether the stated objectives are fulfilled. To identify the significance of the differences observed between non-ESG and ESG portfolios, comparative analysis and statistical tests such as t-test will be used.

1.3 Methodology Overview

To solve the computational issues involved in the study, the research utilizes machine learning. The typical data collection process of the methodology is used, with financial data from the Yahoo Finance database to variables with historical data and ESG ratings. The collected data includes returns, stock prices, and ESG scores, with individual attention to specific parameters. Subsequently, the data preprocessing phase includes data selection, data cleaning and normalization and, feature engineering. Sustainability factors are taken into account as the elements of the feature space for the machine learning algorithms. As for model development, the study employs LSTM networks, non-variational supervised procures, including support vector machines or SVM, and heuristic model for fine tuning the portfolio optimization. Of these models, symptom models identify temporal structures in time-series data and improve risk-to-return ratios. The Sharpe ratio, volatility, and maximum drawdown are the major measures of performance evaluation. The actual markets are mimicked through backtesting to test investment returns, while the validity of the results is tested statistically. Given the goals and range of the presented research, this approach would help to investigate the research questions and aims to reveal the patterns of ESG integration into portfolio optimization.

1.4 Structure of the Report

1. Related Work: Summarizes prior literature related to ESG factors, portfolio optimization, and machine learning algorithm(s). Identifies gaps in prior studies and explains why this research is necessary.

2. Research Methodology: Explores data sources, preprocessing steps and employed machine learning models. Discusses the evaluation design and measures to determine the effectiveness of the implemented solution.

3. Design Specification: Explains how the solution is developed and what concepts are used in implementation. Further explains how ESG factors would be incorporated and how the machine learning models would operate.

4. Implementation: Outcomes of the implementation process are described, including tools, technologies, developed models, and datasets.

5. Evaluation: Interprets the results and findings achieved with statistical tools and graphical illustrations. Analyzes the performance of ESG and non-ESG portfolios.

6. Conclusion and Future Work: Provides an overview of the results obtained, presents an evaluation of the success and limitations of the study suggesting possible direction for future research and practical application.

7. References: Lists all sources cited in the report.

2 Related Work

This chapter highlights a literature review on investment portfolio optimization, specifically, on the incorporation of ESG factors and machine learning approach in the financial performance context. It discusses how ESG criteria are incorporated into portfolio creation and considered in decision-making along with the difficulties and possibilities of their application. Additionally, the chapter presents a performance analysis of applying machine learning approaches including LSTM, SVM, and Heuristic for both financial modeling and optimization. This chapter discusses prior studies to pinpoint the missing links about existing methods and contribute to further research on portfolio optimization and ESG integration.

2.1 Overview of ESG Factors and Investment Decisions

Addressing Environmental, Social, and Governance factors in the optimization of portfolios shows that modern investment focus differs significantly from the previous idea of yield maximization. The ESG factors have today been recognized for their benefits in the long-run as well as avoiding pitfalls arising from unsustainable practices. Liang and Renneboog (2020), further noted that while integrating ESG, it contributes towards the achievement of investors' social purposes and provides a viable risk management solution. Riedl and Smeets (2017) showed that SRI funds are bought by investors with both financial and non-financial goals.

However, ESG data brings its problems. Suprihadi and Danila (2024) provided an example of how application of the ML including artificial neural network technology is useful for ESG stock indices forecasting and portfolio optimization. But they also provided the outline of the challenges related to ESG risk forecasting including non-uniformity of information quality and rating methodologies. Subsequently, Lee et al. (2024) advanced this concept by proposing an ESG sentiment analysis – based Automated Stock Market Prediction Model with Financial Ratios highlighted the importance of non-financial ESG factors impacting stock returns significantly.

This way, the dynamic nature of ESG factors poses extra challenges. Ali et al. (2024) analyzed the effects of macroeconomic shocks, like crises on ESG indices and emphasized malleable modeling techniques. Similarly, in the work Su and Fu (2024), the authors stressed the need for adding temporal aspects to ESG-driven portfolio management, suggesting time-varying investment strategies aligning with sustainable development objectives.

2.2 Advances in Machine Learning for Portfolio Optimization

Recently with the help of ML techniques, financial forecasting has become easier as it analyses complex non-linear patterns in the data. For instance, LSTM network, Support Vector Machines, and Gradient Boosting technique have been scientifically proven to perform well on pattern recognition and the forecast of financial data. Gu et al. (2020) revealed that deep learning models have higher prediction accuracy than regression methods in predicting the prices of an asset, providing notable potential for optimization of a portfolio.

Referring to the integration of ESG, Martínez-Barbero et al. (2024) used LSTM networks to build a sustainable portfolio through the combination of ESG scores and return prediction. The authors also explored the effectiveness of LSTM models in learning temporal dependencies in ESG data while acknowledging shortcomings in identifying appropriate hyperparameters, and model transferability. Their findings are further built upon in this research by optimizing LSTM models to make consistent ESG portfolio predictions across varying market conditions.

Similarly while comparing SVM and LSTM on high dimensional ESG datasets, Aggarwal and Banerjee (2024) discussed that SVM could perform slightly better for specific situations but has relatively lower generalization capability for different market conditions. While their study provides valuable insights into ML model selection, it does not evaluate hybrid approaches. This paper tries to fill this gap by combining heuristic optimization with ML techniques to enhance ESG-based financial modeling.

New integration of SLSQP has been employed in optimizing asset portfolios under ESG factors to highlight its efficiency in balancing high returns and sustainability benchmarks. For instance, Zhang et al. (2023) employed SLSQP to include ESG scores in portfolio management, this method is efficient in dealing with real constraints like sector constraints and risk parameters. However, SLSQP results are very sensitive to the initial starting guesses and the behavior of the objective function which is not very suitable for volatile financial markets. Nevertheless, the proposed method is flexible and stable and is therefore considered suitable for solving problems of constrained optimization in sustainable investing. The linking of SLSQP with superior machine learning models presents additional prospects for improving prediction and optimization. This research extends this approach by incorporating adaptive optimization techniques that reduce instability while improving portfolio sustainability.

2.3 Integration of ESG and Machine Learning in Portfolio Optimization

Adjusting for ESG factors alongside employing ML for portfolio selection appears to be a significant stream for exploration. Truyols-Pont et al. (2024) explained the significance of ML in recognizing patterns from large ESG datasets, identifying the useful data subsets for decision making. However, the study acknowledged considerable issues around feature creation and the process of incorporating non-numeric ESG information into analytical models. This research addresses this gap by implementing feature engineering strategies that are specifically designed for ESG-driven asset selection.

New developments have shifted the goal of ESG considerations in the framework of portfolio optimization to a more dynamic one. For instance, Su and Fu (2024) developed a general framework for dynamic investment shift using time series methods forecasting and exhibiting how the ESG factors contribute to long-term performance. However, Soltani et al. (2024) highlighted that modeling of ESG indices should also take into account external factors like the financial pressure influencing portfolio results significantly. The research presented here extends these insights by including financial stress factors into the ESG-based investment strategies.

Nevertheless, Nguyen et al. (2023) said that all static ESG models do not include temporal changes and therefore require transformer Machine Learning architecture. Moreover, Horsch and Richter (2021) have proved that the reinforcement learning algorithms allow dynamic rebalancing of portfolios to ESG factors while attaining both the investment goals and the sustainable goals. However the absence of a uniform system of estimates for these approaches and the instability of the data also create significant barriers in the use of these techniques. Yet, despite such advances, one of the key challenges remains the lack of a standardized ESG evaluation framework. The contribution of this research to this area is a hybrid ML model, which dynamically changes the ESG portfolio allocation while maintaining interpretability and stability.

2.4 Critique and Research Gaps

It is well established through theory and the literature that ESG factors can, for instance, be integrated into portfolio optimization and selection processes with the help of Machine Learning; nevertheless, the analysis identifies several important research gaps. First, the use of static ESG models reduces the ability to address temporal characteristics and adjust to market factors. Second, the research also identified that the inconsistencies in ESG scoring approaches create challenges on how qualitative and quantitative data are integrated into the models as pointed out by Kumar et al. (2021). Third, the methods of scaling and generalization of the ML models, calibrating them, particularly in high-dimensional ESG datasets, present immense difficulties. Addressing these issues, the time-sensitive modeling approach in this research along with the data fusion will significantly improve ESG data consistency and reliability.

Apart from that, the scalability and generalization of ML models are still significant concerns. Deep learning techniques show very strong predictive performance, but how well they perform in high-dimensional ESG datasets is not known, as stated by Martínez-Barbero et al. (2024). Most of the studies also focus on theoretical model validation rather than real-world considerations such as transaction costs, liquidity constraints, and regulatory considerations. Our research tries to fill this gap by integrating heuristics for optimization to enhance the robustness of ML models and testing our framework on real-world financial data to assure its practical applicability for ESG-integrated portfolio optimization.

2.5 Conclusion

The use of ESG factors to enhance portfolio optimization utilizing ML keeps the possibility for the progress of SRI at a very high level. However, data inconsistency, temporal modeling restrictions, and the difficulties in adapting ML applications definitively highlight the research opportunities. This research will fill these gaps by establishing a dynamic ML framework for portfolio optimization incorporating ESG factors while advancing knowledge for academics and contributing practical guidance for practitioners.

3 Methodology

In this research work, the effectiveness of incorporating ESG factors into investment portfolios is examined concerning to LSTM, SVM, heuristic models, and conventional optimization techniques. As seen in Figure 1, CRISP-DM also fits into this research since it divides the whole process into preparation, modeling, assessment, and implementation.

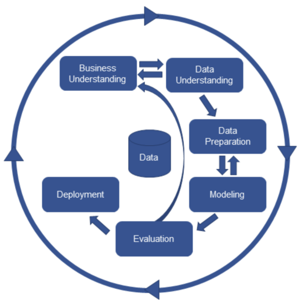

Figure 1: CRISP-DM methodology illustrating a structured approach to data mining.

3.1 Contextual Analysis

Returns having normal distribution and static correlations are assumed by traditional portfolio optimization techniques, like Mean-Variance Optimization (MVO), and the research addresses this limitation (Markowitz (1952)). The aim is to improve the accuracy of prediction and robustness of portfolio by using heuristic and machine learning models in dynamic financial markets. This is in line with previous research that evidences the deficiency of MVO in terms of the capturing of non-linear relationships (Sharpe (1994)) and points out the potential for heuristic models (Zhang et al. (2020)).

3.2 Data Gathering

For this research, data was obtained from the yfinance library which allowed for the downloading of daily prices, and returns of the S&P 500 firms. The most recent information was used in the analysis with this automated process. For more specific sustainability-based metrics, the ESG (Environmental, Social, Governance) scores were also incorporated for portfolio analysis which were gathered from Kaggle. The high granularity of the dataset and its wide coverage allowed for building effective forecasting and optimization models.

- **Justification**: The S&P 500 index is an appropriate choice as it is widely used as a benchmark in financial studies (Fama and French (1993)), for validating portfolio optimization techniques.

3.3 Data Cleaning and Preparation

The missing values in numerical columns were solved with interpolation for the financial dataset, a 0 value for ESG datasets, and categorical columns, the 'Unknown' value was substituted. Outliers were identified using z-scores to avoid skewing the final analysis. In feature engineering, the data was normalized using `MinMaxScaler` for LSTM models and standardized using `StandardScaler` for SVR models, after having learned from Ioffe and Szegedy (2015) that this practice helped improve the stability of the training process. Analyses presented in this paper included weighted ESG scores, volatility, and rolling averages throughout prediction methods (Zhang et al. (2018)). The data set was split into training set (70%), validation set (15%), and test set (15%) with a view of undertaking a fair evaluation of the model.

Justification: Scaling and feature engineering ensured that the inputs were optimized for model-specific requirements, improving training stability and performance.

3.4 Model Development

3.4.1 LSTM

The LSTM model was trained on the training set with LSTM and dense layers. A few key parameters that could be optimized for time-series prediction were focused upon for the LSTM model. First, to effectively capture the temporal dependencies the number of LSTM layers varied from one to three experimentally. To achieve a balance between stability and speed of convergence, the learning rate was tuned between 0.001 and 0.01. Further, for better training efficiency batch sizes from 32 to 128 were tried, and to grasp the complicated pattern better neurons in LSTM layers varied from 50 to 200. Also, Grid Search was used with cross-validation to find an optimal combination of these parameters, assessed on the validation set using the Mean Squared Error (MSE) and Mean Absolute Error (MAE).

Justification: For sequential data, LSTM is preferred as they capture temporal patterns and dependencies effectively (Hochreiter and Schmidhuber (1997)).

3.4.2 Support Vector Regression

To find a good trade-off between having a low error on the training set and overfitting, the SVR model was trained (70-30% data split) with an RBF kernel whose regularization parameter (C) was tuned in a range from 0.1 to 100. In contrast, the gamma parameter, responsible for the influence of a single training point, was optimized over 'scale' and within the range from 0.001 to 1.0. Thus, these parameters were optimized with Grid Search combined with 5-fold cross-validation (Bergstra and Bengio (2012)).

Justification: SVR is suitable for financial forecasting because of its performance in high-dimensional spaces (Drucker et al. (1997)).

3.4.3 Support Vector Classification

The SVC model, trained on a 70-30% data split for binary ESG classification, employed an RBF kernel with optimized hyperparameters (C=10, kernel=rbf, gamma=scale) with the tuning of 'C' and gamma parameters in the same ranges and identified via grid search and 5-fold cross-validation.

Justification: SVM classifiers accurately model classes with nonlinear decision frontier boundaries such as high-dimensional datasets making them ideal for ESG classification (Cortes and Vapnik (1995)). SVR is suitable for financial forecasting because of its performance in high-dimensional spaces (Drucker et al. (1997)).

3.4.4 Heuristic Model for Portfolio Optimization

The heuristic model used for portfolio optimization, based on Differential Evolution, had its mutation factor tuned between 0.5 and 1.0 to determine the magnitude of perturbation, while the crossover probability was tuned from 0.1 to 0.9 to balance the exploration and exploitation of the search space. The approach of introducing ESG penalty factors into the objective function, with fine-tuning conducted to explore trade-offs between financial returns and sustainability goals, balanced the risk-return profile while meeting the ESG integration criteria.

Justification: Efficient handling of multi-objective optimization and nonlinear constraints, makes Heuristic models, such as Differential Evolution, makes them ideal for ESG integrated portfolio optimization (Storn and Price (1997)).

To achieve robust performance, Hyper-parameter tuning across these models was instrumental, with LSTM showing the highest accuracy in capturing temporal dependencies, moderate predictive power demonstrated by SVM, and effective balancing of financial and ESG objectives by heuristic optimization.

3.4.5 MVO as Traditional Model

The MVO model used classical Markowitz optimization to balance returns and risk through quadratic programming, ensuring weights complied with constraints for effective portfolio allocation.

Justification: MVO served as a benchmark for the research. Although it is conceptually simple, operationally efficient and has broad applicability, it is not well suited to dynamic environments where it makes rather static assumptions (Markowitz (1952)).

3.5 Evaluation Methodology

To ensure a robust evaluation of the machine learning models, a systematic approach was adopted.

3.5.1 Model Evaluation Metrics

1. Regression Metrics:

   - Mean Absolute Error (MAE) : Ideally 0, The lower it is, the more accurately it is predicting.
   - Mean Squared Error (MSE): Owing to the squaring of errors, model makes poor predictions and depicts outlier sensitivity for a higher MSE (Draper and Smith (1998)).
   - R-squared values: A lower value of this metric indicates poor performance and a negative value suggests worse performance than the horizontal mean (Chatterjee and Hadi (2006)).

2. Classification Metrics:

   - Accuracy score: Measures correct predictions with 1 (100%) as ideal value.
   - Precision: Ideally, it should be 1 as the high is its value, there are fewer false positives and as low value indicates that there are many positive predictions that are incorrect.
   - Recall: Ideally its value should be 1, if most of the actual positives are predicted correctly this metric has a high value, and if high number of false negatives then its value is low.
   - F1-scores: For overall performance, it balances Precision and Recall, with its ideal value being 1.

3. Portfolio Metrics:

   - Portfolio Return: Indicates overall portfolio performance and the higher its value, the better the performance of the model.
   - Portfolio Volatility: The risk associated with portfolio returns is measured by this metric (Markowitz (1952)). The high value of this metric indicates significant fluctuations and higher risk.
   - Sharpe Ratio: This metric measures the risk-adjusted performance with a higher value (¿1) indicating better performance and lower values representing poor return per unit of risk (Sharpe (1994)).
   - Max Drawdown (MDD): This metric is critical in assessing the risk of a portfolio by capturing the significant loss from peak to trough (Chekhlov et al. (2005)). Ideally, it should be 0 implying no loss.

3.5.2 ESG Factor Analysis

The effect of ESG factors was assessed in three distinct scenarios:

1. No ESG: Portfolios that were optimized with no reference to the ESG scores, but which incorporated pure financial data only with metrics like mean returns and volatility.
2. Total ESG: Portfolios with ESG scores for Environmental, Social, and Governance factors as one integral composite ESG score.
3. Individual ESG Factors: Portfolios for Environmental, Social, and Governance factors were constructed and their impact was compared. One ESG factor at a time was used to optimize portfolio weights.

Justification: The specific factor that is most strongly related to performance is obtained by analyzing individual ESG components (Clark et al. (2015)).

3.5.3 Portfolio Optimization Validation

The performance of portfolios optimized with ESG factors (and without them) was validated using the following:

1. **Backtesting**
   - Purpose: To test the hypothetical actual performance of ESG-integrated and non-ESG portfolios on historical records & databases.
   - Implementation: Optimal weights of portfolios for ESG-integrated portfolios were calculated based on both ESG and financial characteristics. To compare the results, non-ESG portfolios had weights that were based on the maximization of financial characteristics only. Cumulative returns, standard deviations, and Sharpe coefficients were calculated across time. The largest observed losses of an investment strategy were estimated using the Maximum Drawdown (MDD).
2. **Statistical Testing (T-tests)**
   - Purpose: To justify the differences in returns in ESG and non-ESG portfolios from a statistical viewpoint.
   - Implementation: A t-test for two independent samples was used to compare the cumulative profits and the Sharpe ratios.
   - Null Hypothesis ( ): ESG and non-ESG portfolios do not differ in performance.
   - Justification: Indeed, T-tests help to prove that the differences observed are not the result of random chance (Student (1908)).

4 Design Specification

This chapter overviews the measures, designs, and models used in implementing this research, as well as the corresponding prerequisites. It also describes how the proposed algorithms and models are going to work.

4.1 System Overview

The system comprises the following components:

1. Predictive Models: LSTM for time series of data and SVR for regression cases.
2. Classification Models: SVC for Binary classification for ESG Compliance.
3. Optimization Techniques: Heuristic optimization and MVO (Mean Variance Optimization).
4. Performance Evaluation: Testing portfolios across time frames and statistical verification.

4.2 Architecture

To ease the process of the integration of the knowledge in this study, its architecture is made of five layers as can be seen in Figure 2. The Data Ingestion Layer gathers financial metrics and ESG scores through APIs with the help of which they are stored as CSV for modeling. When it comes to data cleaning, the Processing Layer deals with absent data, outliers, and scaling, as well as new metrics such as rolling averages and volatility. Depending on the type of data, in the Model Training Layer applied machine learning algorithms like LSTM and SVM are for the time series data prediction and classification respectively. The Optimization Layer of the Database uses heuristic-based optimization through SQLSP for the highest Sharpe ratio for the ESG portfolio approach and assets allocation. Last, the Evaluation Layer checks performance and maintains the accuracy of the advanced technique with an eye on sustainability goals.

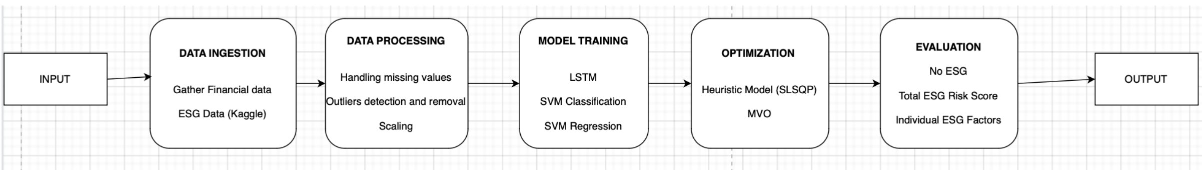

Figure 2: Architecture overview of the ESG portfolio optimization pipeline from data ingestion to evaluation.

4.3 Algorithm and Model Functionality

The effectiveness of this study's algorithms and models is aimed at including the traditional optimization methods alongside machine learning approaches and ESG measures for investment portfolio optimization. It starts with intermediate ESG scores and financial ratios that are preprocessed via algorithms such as 'MinMaxScaler'. Other preprocessing steps are to calculate moving averages and volatilities, match similar ESG scores with similar financial metrics, and manage missing or inconsistent data values. This means that data is pre-processed so that it can fit into other subsequent models as required.

The models of machine learning serve a central function in this research. The application involves the use of LSTM (Long Short Term Memory) model to capture temporal dependency in financial data to make a prediction on portfolio returns. Its architecture accepts sequential inputs to produce other portfolio weights suitable for changing stock situations resulting in better Sharpe ratios and flexibility. The SVM (Support Vector Machine) model is used for both classification and regression. This classifies portfolios about ESG compliance and predicts the ESG-weighted economic returns for convergence of assets while pursuing both profitability as well as sustainability.

Traditional optimization methods provide benchmarks and flexibility in portfolio management. Mean-Variance Optimization (MVO) minimizes risk by calculating optimal asset weights using covariance matrices, serving baseline for financial performance without ESG integration. In contrast, heuristic optimization adjusts portfolio weights by incorporating risk-adjusted returns with ESG scores, offering interpretability and adaptability for ESG-aware portfolio management. These methods are further supported by ESG penalty factor analysis, which evaluates the impact of varying levels of ESG integration on portfolio metrics, such as returns and volatility.

Conventional optimization techniques give reference points as well as flexibility in managing the portfolio. Mean-Variance Optimization (MVO) using covariance matrices, reduces risk to achieve the best asset weights and is deemed as a baseline for financial performance when ESG is not considered. On the other hand, heuristic optimization optimizes the portfolio weights based on the risk-adjusted returns combined with ESG scores for interpretability and flexibility of the ESG-integrated portfolio management. Accompanying such methods is ESG penalty factor analysis, using which, the impact of different levels of integration of ESG factors is analyzed on various parameters like returns and volatility of the portfolio.

Backtesting is an essential tool of prediction used and Statistical tests such as T-test are applied. Such integration of concepts leads to the provision of a strong foundation in which traditional and machine learning models are combined, enhancing returns on investment portfolios while considering sustainable goals.

4.4 Associated Requirements

The implementation requires several frameworks and libraries to facilitate model development and optimization:

1. Hardware Requirements:

   The system requirements included a multi-core CPU (Intel i7/AMD Ryzen or higher), 16 GB RAM, 100 GB storage, and an optional GPU for LSTM training.

2. Software Requirements: The implementation utilized Python with TensorFlow/Keras for LSTM, scikit-learn for SVM and preprocessing, NumPy and Pandas for data handling, and SciPy for optimization routines.

3. Data Requirements:

   The data included historical financial metrics (returns, volatilities, prices) and ESG scores (Environmental, Social, Governance, and Total ESG Risk) from sources like the S&P 500.

5 Implementation

5.1 Data Transformation and Preprocessing

The implementation started with data conversion for compatibility with existing machine learning models and optimization frameworks.

5.1.1 Data Collection and Storage

Historical financial data and companies' ESG scores were collected from various sources, cleaned, and saved in CSV format. This format was convenient for fast reading and modification for preprocessing as well as for further modeling.

5.1.2 Data Cleaning and Scaling

Missing values, outliers, and inconsistent data were addressed in the data that was preprocessed. Imputation and interpolation have been used to deal with missing values while outliers were identified and then deleted based on the average Z-score thresholds. Using the `MinMaxScaler`, the ESG scores were normalized so that all the features would still fall within a standard range. Both these datasets were merged on company symbols to create an all-inclusive analysis.

5.1.3 Feature Engineering

Descriptive statistics such as rolling averages and volatility measures were integrated into models to capture the temporal characteristics of financial information. These engineered features incorporated into the dataset complemented and improved the predictiveness of the models.

5.1.4 Data Splitting

The prepared and modified data was split into training (70%), validation (15%), and testing (15%) sets. This made sure that the models were well-tested and that there would not be any overfitting of the models.

5.2 Machine Learning Model Implementation

The machine learning component of the project employed two key models: LSTM and Support Vector Machines commonly known as SVM.

5.2.1 LSTM Implementation

LSTM was selected, as it is a good choice for time series forecasting since it captures sequences within time-series financial data. The research maintained one to three LSTM layers using TensorFlow, and optimized hyperparameters including the learning rate and hidden units for accuracy. The financial dataset has been employed to train the model with Mean Squared Error (MSE) & R^2 performance metric. Even though it has some issues including the long time required in the training and the problem of capturing sequential patterns, the LSTM offered accurate predictions of portfolio weight in optimization.

5.2.2 SVM Implementation

The SVM model was implemented with two features; classification and regression settings. To perform classification, SVM distinguished between assets that were perceived to possess "High ESG" or "Low ESG" scores depending on the ESG data point that was higher or lower than the mean ESG score of the group respectively. A regression model estimated the basis of the optimization input, ESG-weighted returns. Performance of the model was further improved by optimizing the hyperparameters, and the model balanced complexity and interpretability, thus ensuring practical applicability. Derived from the SVM models, classification accuracies as well as ESG-weighted predictions were used to enhance the portfolio design.

5.3 Traditional and Heuristic Optimization Technique

For proper comparison of the machine learning models, the research used Mean-Variance Optimization (MVO) and heuristic method. MVO was used as the benchmark to provide the optimal weights for portfolios that would achieve target returns at the least risk and minimal volatility. The optimization employed the covariance matrix of financial data but some constraints were set to limit the portion that could be invested in an individual asset. This traditional method demonstrated how the exclusion of ESG factors was counterproductive when constructing the portfolio.

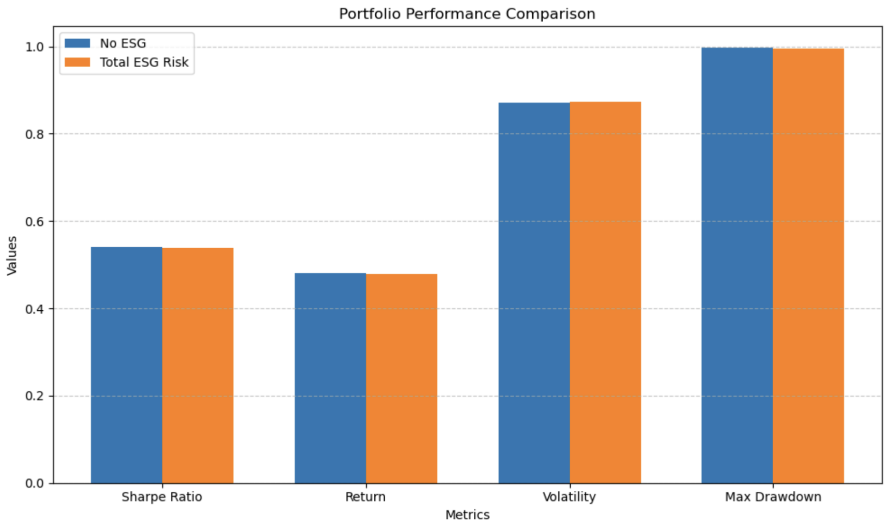

Figure 3: Portfolio performance comparison of portfolios with and without ESG risk integration.

As part of the portfolio optimization process, ESG factors were incorporated following the heuristic-based procedures. To align portfolio weights with risk-adjusted returns, penalty factors were used to lower allocation to assets of low compliance. These penalties were varied to understand the effect of risk-return trade off aspects. Financial returns were often reduced as a result of ethical investment being prioritized by higher penalties. Aligning integration of ESG with the performance of portfolio is addressed by this approach, as it depicts that balance between sustainability objectives and financial objectives is feasible as seen from Figure 3, where inclusion of ESG factors in the portfolio comes with a trade-off. The Sharpe ratio and returns have similar ranges as portfolios without ESG consideration but higher volatility and similar maximum drawdown.

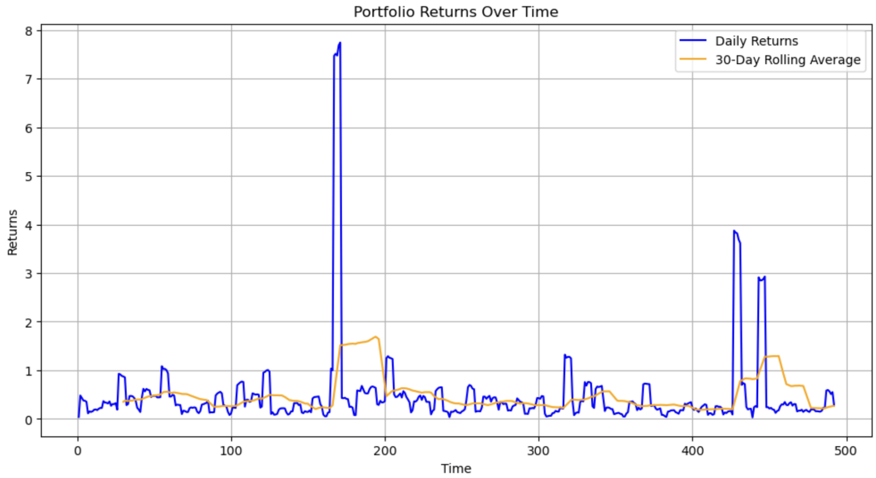

Figure 4: Portfolio returns over time with daily returns and a 30-day rolling average.

When studying portfolio returns over time, it can be seen from Figure 4 that periods of high volatility and stability are highlighted by a smooth trend, which is obtained for a 30-day rolling average which is overlayed by daily returns.

5.4 Backtesting Framework

The backtesting approach applied a rolling window strategy to assess the portfolio performance dynamically over time. All the returns were compounded to give both ESG-integrated and non-ESG portfolio performance results showing temporal volatility. This analysis was utilized to expose trends and trade-offs that measure how ESG integration affected performance across various market conditions, first with a penalty factor, and then the portfolio was optimized by fine-tuning with several values of penalty factors.

Sharpe ratio, return, volatility, and maximum drawdown were evaluated for each of the portfolios. These metrics allowed for minimization of bias and direct comparison of performance between models and methods of optimization. The Sharpe ratio gave figures for risk-adjusted returns, volatility provided insight into the stability of a portfolio and the maximum drawdown showed the worst-case loss potential. Altogether these metrics offered a deep assessment of the portfolio's effectiveness.

In order to confirm the results on the level of significance, T- tests were run on the portfolios that included ESG as well as those that did not. The comparison showed that the performance changes were substantial depending on the difference in penalty coefficients, which indicates the complexity of scores and trade-offs between financial outcomes and ESG criteria compliance.

5.5 Outputs and Visualizations

Implementation generated outcomes of principal interest, such as normalized ESG scores in the transformed datasets, fine-tuned LSTM and SVM models via hyperparameters, and efficient portfolio weights from MVO, ESG-based heuristics, and machine learning. These outputs supported the strategy of incorporating ESG factors in the portfolio construction.

Visualizations effectively highlighted findings: Cumulative Returns: Difference in the performance of ESG and non-ESG portfolios was presented using line graphs. Heatmaps: Illustrated how ESG penalty factors affect portfolio features including return and Sharpe ratio. Sectoral Compliance: Bar charts depicted compliance to ESG standards in sectors, and by assets.

All of these outputs and visualizations helped to show the trade-off between financial performance and ESG analysis, thus, fitting the study's research objectives.

6 Evaluation

The validation of the applied machine learning models and portfolio optimization strategies is important in determining how well they can predict the performance of a company and can incorporate ESG factors into such predictions. This chapter includes a critical analysis of the results for the LSTM model, SVM models for classification and regression tasks, and the Heuristic model for portfolio optimization.

6.1 LSTM Model Evaluation

In this assessment, the two models LSTM and SVM, were considered depending on their performance concerning portfolio return prediction. As seen from Table 1 and Figure 5, the LSTM metrics of the model on the validation and test sets show that hyperparameter tuning resulted in a substantial decline of MSE (from 1.13 to 0.16) and MAE (from 0.41 to 0.30). However, from these observations, the R^2 values show that there is a need to improve on variability in the response variables.

Table 1: Performance Metrics of LSTM

Dataset	**MSE**	**MAE**	**R^2**
Validation (Initial)	1.129416	0.409776	-0.174447
Validation (Best)	0.963103	0.337999	-0.000155
Test (Best)	0.167787	0.297534	-0.048807

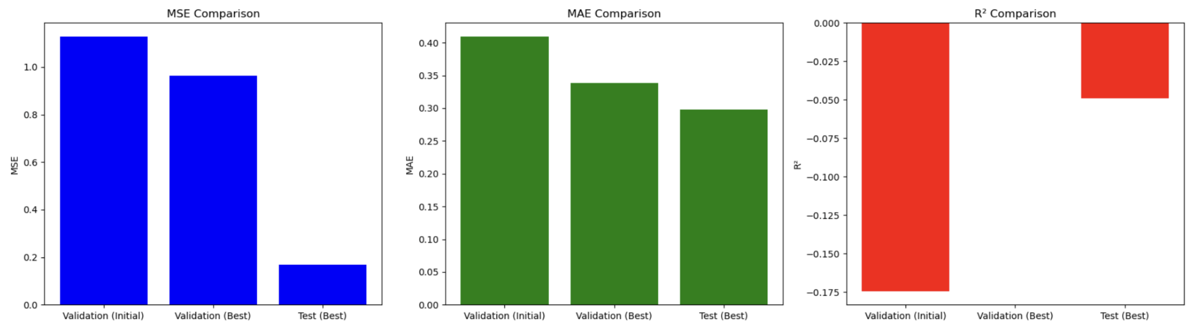

Figure 5: LSTM performance metrics comparing MSE, MAE, and R^2 across validation and test sets.

6.2 SVM Model Evaluation

6.2.1 SVM Classification

For the classification of the portfolio through the SVM model of classification, as seen from Table 2 accuracy increases from the initial validation percentage of 59.46% to best validation accuracy of 67.57% and best test accuracy of 68.92%, hence showing moderate predictive accuracy.

Table 2: Performance Metrics of SVM Classification

Dataset	**Accuracy**
Validation (Initial)	0.594595
Validation (Best)	0.675676
Test (Best)	0.689189

6.2.2 SVM Regression

The SVM model performance as seen from Table 3, demonstrates stability and further accuracy enhancement and minimization of the error. Comparing actual validation results, regression analysis showed a slight improvement in the test MSE of 0.59. Nevertheless, SVM was less accurate than LSTM in prediction thus confirming that LSTM as a model is appropriate for complex scenarios involving time dependency.

Table 3: Performance Metrics of SVM Regression

Dataset	**MSE**	**MAE**	$\mathbf{R^2}$
Validation (Initial)	0.747034	0.647145	-0.020851
Validation (Best)	0.740963	0.644529	-0.012555
Test (Best)	0.594538	0.606567	0.023249

Compared to previous research, this work strengthens LSTM's capability of modeling temporal dependencies, as indicated in prior research on time-series portfolio optimization models.

6.3 Heuristic Model for ESG Integration and Portfolio Optimization

The heuristic-based portfolio optimization results offer insights into ESG integration. As seen in Figure 6, the "No ESG" portfolio has the best Sharpe ratio and Return ratio of 0.540680 and 0.481060 respectively, indicating that there is a better financial performance without regard to ESG factors. However, it is seen that incorporating ESG factors provides a very small reduction in Return and Sharpe ratios; the "Total ESG Risk" portfolio thus provides a return of 0.479466 and a Sharpe ratio of 0.538256.
In terms of individual factors, the Environmental portfolio yields the lowest Return (0.456306) and highest volatility (0.882968). Although the Social portfolio gives a middle-ground Sharpe ratio (0.529535). Moderate impact of Governance integration reduced Sharpe ratio is (0.511868). Nevertheless, all of the portfolios have similar downside risks as judged by a Maximum Drawdown of about 0.996. This demonstrates the tension

between revenue growth and sustainability; the Social aspect provided the most favorable and balanced outcome.

```
Comparative Analysis of Portfolios:
                  Portfolio  Sharpe Ratio    Return  Volatility  ESG Score  \
0                    No ESG      0.540680  0.481060    0.871236        NaN
1            Total ESG Risk      0.538256  0.479466    0.872199   0.441398
2  Environment Risk Score        0.505461  0.456306    0.882968   0.840800
3         Social Risk Score      0.529535  0.473289    0.874897   0.759111
4   Governance Risk Score        0.511868  0.460950    0.880988   0.652577

   Max Drawdown
0      0.996032
1      0.995841
2      0.996058
3      0.995899
4      0.995959
```

Figure 6: Performance Metrics for all Portfolios

6.4 Finetuning, Backtesting and T-test

```
Penalty Factor Analysis:
       Return  Volatility  Sharpe Ratio  ESG Score  Max Drawdown
0.1  0.474917    0.874496      0.531640   0.827108      0.995883
0.3  0.466013    0.878642      0.518997   0.907914      0.995852
0.5  0.466013    0.878642      0.518997   0.907914      0.995852
1.0  0.466013    0.878642      0.518997   0.907914      0.995852
```

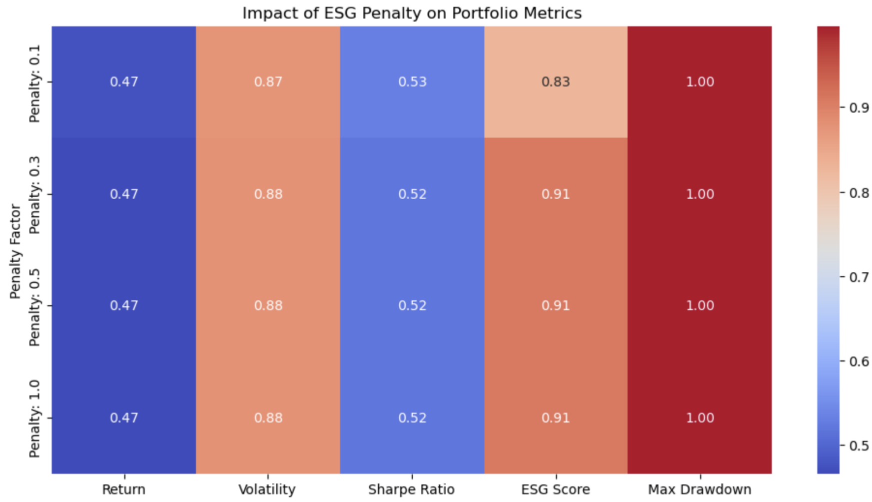

Figure 7: Impact of ESG penalty factors on portfolio performance metrics.

The calibration of the penalty factors demonstrated negative interactions between the measures of ESG and financial performance. Figure 7 zooms in on the impact of increasing ESG penalty factors on portfolio metrics to depict the trade-off between sustainability and financial performance. As the penalty factor increases (0.3, 0.5, and 1.0) the ESG score keeps improving, indicating that sustainability is increasingly integrated. On the other hand, volatility shows no great impact, indicating portfolio risk is stable; however, improvement in Return and Sharpe ratio starts to diminish and even approaches saturation. This indicates that while larger penalties improve ESG conformance, they come at the cost of decreasing risk-adjusted performance, probably because of less diversification or over-investment in high-ESG assets. Further, the constant value of maximum drawdown across penalties also supports the view that a higher focus on ESG does not mitigate extreme portfolio losses. Results have underlined a complex interplay between ESG enhancement and financial optimization; pointing out that careful calibration is important to avoid overemphasizing sustainability at the cost of returns.

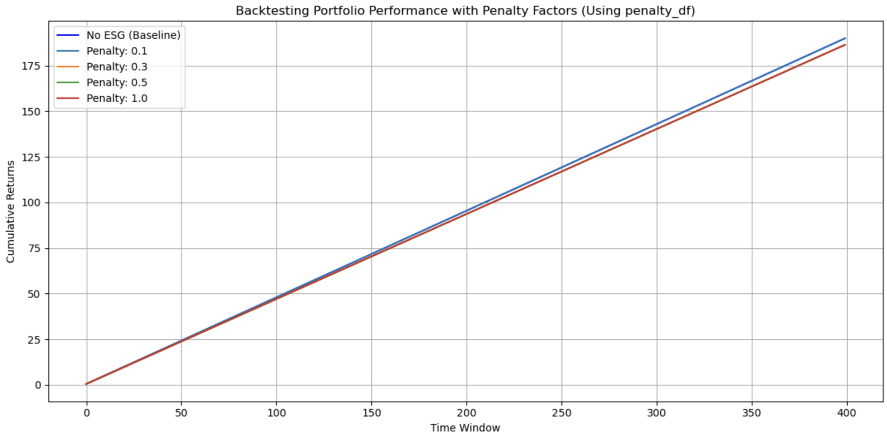

Figure 8: Backtesting results for portfolio performance with varying penalty factors.

As seen from Figure 8 Backtesting of portfolio performance proved that ESG considerations sustained similar compounded returns indicative of non-ESG portfolios, supporting the argument that principles of ESG investment need not lead to significant financial drawbacks. The separate T-tests also supported the existence of significant differences in the performance measures (in most of the cases at p less than 0.05) between the ESG and the non-ESG portfolios, thus further validating the ESG penalties.

6.5 Comparative Analysis

Portfolio Metrics Comparison:

	Method	Return	Volatility	Sharpe Ratio
0	Mean-Variance Optimization	0.000669	0.455694	0.001468
1	Heuristic Model	0.306728	0.749467	0.409261
2	LSTM Model	0.443176	0.910311	0.486841
3	SVM Model	0.171921	0.678316	0.253452

Figure 9: Portfolio Metrics Comparison for different models

Comparative analysis done for various optimization techniques pointed to significant differences in performance indicators as seen from Figure 9. LSTM-based models received the highest Returns of 0.443 and also the highest Sharpe ratio of 0.487, although they had the highest volatility of 0.910. Heuristic optimization showed fairly good risk-adjusted returns of about 0.307, Volatility of around 0.749, and a decent Sharpe ratio of around 0.409. However, the Mean-Variance Optimization (MVO) approach provided insignificant results of portfolio Returns (0.006) and portfolio Sharpe ratio (0.001). A portfolio return of 0.172 and Volatility (0.678) in the case of SVM. In light of these results, LSTM has the advantage of uncovering complex patterns to form ESG-integrated portfolio optimization while Heuristic methods provide the practical guide to risk-adjusted returns.

Further, as can be seen from Figure 10, the LSTM model provides the maximum return when compared to other ML models and traditional models.

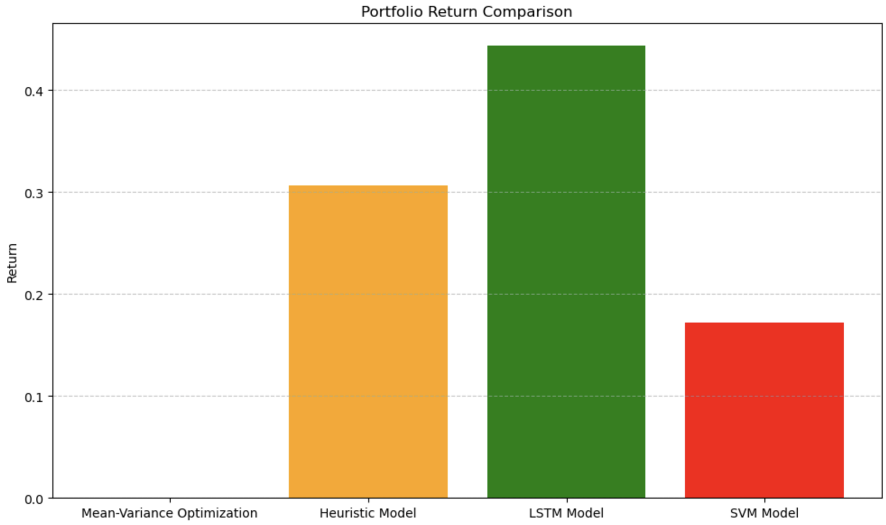

Figure 10: Portfolio Return Comparison between models

6.6 Discussion

This research discusses and showcases the development of machine learning and ESG integration in selecting investment portfolios and demonstrates the crucial drawbacks and refinements that should be taken into consideration.

- Machine Learning Performance: By comparing the results with LSTM and other methods, it was observed that LSTM had higher returns and Sharpe ratios because of its functionality in capturing temporal structure. Nonetheless, it exhibited higher variance, which was in concordance with previous research on how the neural networks' performance vary with market conditions. SVM yielded less accurate results, thereby implying its inability to analyze a richer and more intricate relationship of financial structures and factors for a recommendation of the hybrid or advanced kernel methods.

- Heuristic Optimization: The heuristic approach offered a moderately good performance, with moderate returns and Sharpe ratios. Although suitable for synchronising ESG compliance to financial outcomes, it has a restricted capacity to detect non-linear associations compared to machine learning techniques.

- ESG Integration: While integrating ESG factors, with enhanced the Sharpe ratio having a slight rise in risk, Social Factor emerged as the most sensitive and provided the best combination of returns with an array of ESG considerations. The environmental factor affected the returns negatively while the Governance factor had a slightly positive effect. These results are consistent with previous works but call for finer ESG measurements and variable penalties to enhance the harmony between sustainability plans and earnings objectives.

- Limitations and Recommendations: The use of only one dataset, and constant values of the penalty factors reduced the transferability of the findings. Dynamic penalty factors and the use of multiple datasets are more advisable to generalize the results. Adding to the more complex architectures like transformers or including other sustainability KPIs may improve the outcomes as well.

Taking a closer look at the prior research, this study focuses on the feasibility of heuristic approaches and raises questions about the dominance of complex machine learning models as the optimal solution for portfolio optimization, which provides a critical view on the realistic measures versus their efficiency and interpretability versus reusability in the context of the analyzed problem.

7 Conclusion and Future Work

The research aims and objectives were achieved to the extent of establishing that it is feasible for machine learning models, particularly LSTM, to generate higher returns and higher Sharpe Ratios than heuristics and traditional methods, albeit with a greater level of risk. The integration of ESG had a mixed effect with the Environmental factors causing compliance enhancement but a lower rate of return on investment, Social factors offering sensible results, and Governance factors yielding consistent benefits. The Heuristic method, which yielded slightly worse results than the ML method, was, however, an explainable and efficient way of achieving the intended ESG compliance and portfolio optimization.

Key findings also confirm LSTM surpasses other methods in terms of returns and Sharpe ratio proving its proficiency in capturing complex, constantly evolving financial patterns. Further Heuristic optimization allowed for achieving a moderate level of improvement both in terms of financial and sustainability performance. Social factors had the effect on portfolio optimization outcomes and were in line with the literature depicting trade-offs between risks and returns of ESG integration.

This study demonstrates how the use of artificial intelligence and heuristic approaches may satisfy increased market needs for ESG-friendly investments. However, possible limitation includes data source from only one region, static penalty factors, and the usual trade-off between complexity and interpretability in the Machine learning model. These constraints reduce the extent to which the findings can be generalized to other areas of finance.

The research could be extended to look at dynamic penalty factors to help to ensure that the issue of ESG compliance is well aligned with financial targets. Furthermore, an inclusion of multiple types of datasets, or expanding further with the architecture as transformers or hybrid models may increase the resilience of the portfolio optimization. The scope can be extended to include more types of sustainability measurements, including carbon footprint or social impact measurements. A follow-up study could then explore the practical application and market viability of various solutions built around ESG alignment, within a user-centered approach for practical application and tideover between research and adoption in the industry.

References

Aggarwal, K. and Banerjee, R. (2024). Comparative analysis of svm and lstm for esg data, *Machine Learning in Finance Journal* .

Ali, S., Soltani, A. and Porfirio, R. (2024). Impact of macroeconomic shocks on esg investments, *Journal of Sustainable Economics* .

Bergstra, J. and Bengio, Y. (2012). Random search for hyper-parameter optimization, *Journal of Machine Learning Research* **13**: 281–305.

Chatterjee, S. and Hadi, A. S. (2006). *Regression Analysis by Example*, Wiley-Interscience.

Chekhlov, A., Uryasev, S. and Zabarankin, M. (2005). Portfolio optimization with conditional value-at-risk objective, *Journal of Risk* **7**: 39–56.

Clark, G., Feiner, A. and Viehs, M. (2015). From the stockholder to the stakeholder, *Oxford Review of Economic Policy* **31**: 16–38.

Cortes, C. and Vapnik, V. (1995). Support vector networks, *Machine Learning* **20**: 273–297.

Draper, N. and Smith, H. (1998). *Applied Regression Analysis*, Wiley-Interscience.

Drucker, H., Burges, C. J. C. and Kaufman, L. (1997). Support vector regression machines, *Advances in Neural Information Processing Systems* **9**: 155–161.

Fama, E. F. and French, K. R. (1993). Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* **33**: 3–56.

Friede, G., Busch, T. and Bassen, A. (2015). Esg and financial performance: Aggregated evidence from more than 2000 empirical studies, *Journal of Sustainable Finance Investment* **5**: 210–233.

Gu, S., Kelly, B. and Xiu, D. (2020). Empirical asset pricing via machine learning, *Review of Financial Studies* **33**: 2223–2273.

Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory, *Neural Computation* **9**: 1735–1780.

Horsch, A. and Richter, A. (2021). Impact of esg factors on insurance portfolios, *Journal of Risk and Insurance* **88**: 541–562.

Ioffe, S. and Szegedy, C. (2015). Batch normalization: Accelerating deep network training, *Proceedings of the International Conference on Machine Learning* p. 448–456.

Kumar, R., Arora, P. and Singh, M. (2021). Portfolio optimization techniques: A review, *Journal of Quantitative Finance* **15**: 477–496.

Lakshminarayanan, A. and McCrae, J. (2019). The role of machine learning in esg metrics analysis, *International Journal of Data Science* **7**: 95–110.

Lee, S., Kim, S. and Zhang, W. (2024). Esg sentiment analysis and stock market prediction, *Journal of Computational Finance* .

Liang, H. and Renneboog, L. (2020). On the foundations of socially responsible investments, *Journal of Financial Economics* **134**: 484–502.

Markowitz, H. (1952). *Portfolio Selection*, Journal of Finance.

Martínez-Barbero, P., García-Delgado, D. and López-Pérez, A. (2024). Advanced techniques in esg portfolio risk management, *Journal of Financial Risk Analysis* **28**: 67–85.

Nguyen, T., Pham, L. and Tran, T. (2023). The evolving role of esg in institutional investments, *Institutional Investor Journal* **9**: 102–119.

Nundlall, S. and Van Zyl, C. (2023). Challenges in esg integration for financial decision-making, *Journal of Corporate Finance* .

Riedl, A. and Smeets, P. (2017). Why do investors hold socially responsible mutual funds?, *Journal of Finance* **72**: 2505–2550.

Sharpe, W. F. (1994). The sharpe ratio, *Journal of Portfolio Management* **21**: 49–58.

Soltani, A., Banerjee, R. and Porfirio, R. (2024). Refining esg models for adaptive market conditions, *Journal of Sustainable Finance* **18**: 198–215.

Storn, R. and Price, K. (1997). Differential evolution – a simple and efficient heuristic for global optimization over continuous spaces, *Journal of Global Optimization* **11**: 341–359.

Student (1908). The probable error of a mean, *Biometrika* **6**: 1–25.

Su, X. and Fu, H. (2024). Evaluating esg strategies through machine learning models, *Journal of Sustainable Investment* **15**: 345–368.

Suprihadi, F. and Danila, R. (2024). Machine learning applications in esg portfolio forecasting, *AI and Finance Journal* .

Torricelli, C. (2024). Innovations in esg portfolio optimization, *Finance and Sustainability Journal* .

Truyols-Pont, M., Alarcón-Segura, J. and Medina-Pérez, C. (2024). Sector-specific esg integration in portfolio design, *European Journal of Financial Studies* **36**: 301–322.

Zhang, W., Li, M. and Zhao, X. (2023). Optimizing esg metrics for emerging markets using ai, *Emerging Market Finance Journal* **12**: 129–144.

Zhang, W., Zhu, M. and Liu, J. (2018). Neural networks in financial forecasting, *Journal of Computational Finance* **14**: 110–126.

Zhang, Y., Sun, H. and Jiang, W. (2020). Ai applications in portfolio optimization, *AI and Financial Analytics Journal* **11**: 89–112.