

PAPER • OPEN ACCESS

Bayesian optimization of ESG (Environmental Social Governance) financial investments

To cite this article: Eduardo C Garrido-Merchán *et al* 2023 *Environ. Res. Commun.* **5** 055003

View the [article online](#) for updates and enhancements.

You may also like

- [The interplay of policy and energy retrofit decision-making for real estate decarbonization](#)
Ivalin Petkov, Christof Knoeri and Volker H Hoffmann
- [A Sustainable environmental change and ESG initiatives by the manufacturing and others service Industries during COVID19 Pandemic](#)
Monika Mittal, Manoj Pareek, Shubham Sharma et al.
- [A comparative analysis of green financial policy output in OECD countries](#)
Bjarne Steffen

Environmental Research Communications



PAPER

Bayesian optimization of ESG (Environmental Social Governance) financial investments

OPEN ACCESS

RECEIVED

13 February 2023

REVISED

21 April 2023

ACCEPTED FOR PUBLICATION

27 April 2023

PUBLISHED

5 May 2023

Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#).

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.



Eduardo C Garrido-Merchán, Gabriel González Piris and María Coronado Vaca*

Universidad Pontificia Comillas, Madrid, Spain

* Author to whom any correspondence should be addressed.

E-mail: ecgarrido@icade.comillas.edu, gabrielgonzalezpiris@alu.icaic.comillas.edu and mcoronado@icade.comillas.edu**Keywords:** environmental social governance, ESG, socially responsible investments, portfolio optimization, Bayesian optimization (BO), Genetic algorithm (GA), Simulated Annealing (SA)**Abstract**

Financial experts seek to predict the variability of financial markets to ensure investors' successful investments. However, there has been a big trend in finance in the last few years, which are the ESG (Economic, Social and Governance) criteria, due to the growing importance of investments being socially responsible, and because of the financial impact companies suffer when not complying with them. Consequently, creating a stock portfolio should consider not only its financial performance but compliance with ESG criteria. Portfolio optimization (PO) techniques previously applied to ESG portfolios, are all closed-form analytical ones. But the real world is rather a black box with unknown analytical expressions. Thus, in this paper we use Bayesian optimization (BO), a sequential state-of-the-art design strategy to optimize black-boxes with unknown analytical and costly-to-compute expressions, to maximize the performance of a stock portfolio under the presence of ESG criteria soft constraints incorporated into the objective function. And we compare it to two other black-box techniques widely applied for the optimization of 'conventional portfolios' (non-ESG ones): the metaheuristics Genetic algorithm (GA) and Simulated Annealing (SA). Although BO has many theoretical advantages over GA and SA, it has never been applied to PO. Thus, this paper investigates whether BO can be used in the ESG PO framework as an alternative and compares it with GA and SA. This is the research gap to which this paper responds. To show the empirical performance of BO, we carry out four illustrative experiments and find evidence of BO outperforming the baselines. Thus we add another different optimization approach to the world of ESG investing: a black-box non-heuristic optimization approach through BO. Our study is the first paper that leverages BO and ESG scores into a PO technique. This paper opens the door to many new research lines in (ESG) portfolio optimization.

1. Introduction

In the past 15 years, ESG criteria have increasingly become integrated into mainstream portfolio management. ESG investment -also called socially responsible investment (SRI)- has attracted much attention from both institutional and individual investors in capital markets [1–8]. Many long-term institutions such as pension funds, insurance companies, sovereign wealth funds, foundations and endowments have signed up to the UN Principles for Responsible Investment (PRI). The six PRI offer a menu of possible actions for incorporating ESG issues into investment practice. The PRI were developed in 2006 by an international group of institutional investors reflecting the increasing relevance of environmental, social and corporate governance (ESG) issues to investment practices. The process was convened by the United Nations Secretary-General. As of December 2022, the PRI has 5,179 signatories, representing US\$121 trillion of assets under management (AUM) -a huge increase from US\$6.5 trillion in 2006- [9]. In addition, individual investors show increasingly strong activism claiming for their money being invested in ESG assets and PRI signatories rely on shareholder activism to pursue responsible investing [10].

In addition, the latest United Nations Global Compact-Accenture CEO Study on Sustainability [11] provides insights and a resolute call to action from a record from over 2,600 CEOs across 128 countries and 18 industries. Sustainability has emerged as the core of resilience. Compared to the 2013 CEO Study, CEOs now unequivocally feel it is their role to make their business more sustainable (98% agree versus 83% in the 2013 edition). Sustainability is the only path to building a truly resilient company. We must accelerate our work in sustainability to build a more resilient future if we are to rescue the UN Sustainable Development Goals (SDGs) by 2030 [11]. The 2021 CEO Study found that business leaders were severely off track to delivering on their sustainability and climate goals. Today, the situation is even more tenuous. Facing continued fallout from the pandemic, coupled with the effects of Russia's war in Ukraine, broader geopolitical uncertainty, inequality, and climate change, CEOs report heightened frustration and uncertainty in preparing for what will happen next. As a result, CEOs are now making sustainability a top priority in their agendas, re-evaluating their investment criteria and developing innovative business models enabled by technology to drive change [11].

Thus, ESG criteria have become more significant due to the growing importance of investments being socially responsible, and because of the financial impact companies suffer when not complying with ESG criteria. The E of Environmental takes into account the direct and indirect environmental impact of the company's activities, the S of Social considers the impact of the company in society and its community, and the G of Governance takes into account the corporate governance of the company.

The interest of investors towards ESG funds and assets has grown lately. Morningstar highlights in its U.S. Sustainable Funds Landscape Report how there are 5 times as many sustainable funds in the U.S. today than a decade ago, and 3 times more than five years ago [12, 13]. Additionally, it is strongly believed by many investors that integrating ESG criteria when investing increases returns and profitability. The global asset management firm macquarie AM states in its report 'ESG survey' that three out of four institutional investors consider that taking into account ESG criteria in its investments increases financial returns. Empirical evidence in this respect is mixed as several studies show the existence of a negative relationship between the environmental and financial performance of portfolios while others argue in favor of a positive effect. For this, we refer to the most recent studies and meta-analysis conducted on this area [14, 15].

An investor's fundamental goal is optimally allocating his investments between different assets. This is called an 'asset allocation' or 'portfolio optimization' problem. An investor has to choose (a) in which assets to invest and (b) how much in each of these assets. In other words, how can (s)he get the optimal portfolio? What an optimal portfolio is must be defined. The pioneer Markowitz's model [16] assumes that investors only care about the portfolio's return and risk, not about ESG criteria or more variables other than the portfolio's risk and return). Thus, the optimal portfolio has minimum risk or maximum return in this model. In order to incorporate the ESG constraints, a few variants have been proposed from the original Markowitz's approach, as will be shown in section 2.

As a result of the situation described above, the relevance of ESG investing involves the need for more research on alternative portfolio optimization techniques within the ESG framework. Despite the voluminous literature that analyzes the financial performance of ESG portfolios compared to 'conventional' -or non-ESG- ones, we believe that the literature on comparing portfolio optimization frameworks used to determine the asset weights in ESG portfolios is underdeveloped.

Moreover, Oikonomou *et al* (2018) proved that different optimization techniques lead to different ESG portfolio performance [5]. Hence, apart from the ESG screening criteria, investors and portfolio managers also need to carefully consider the choice of asset allocation method. So, it is claimed to consider and investigate alternative optimization techniques within the ESG framework. A few different techniques for portfolio optimization have been applied to ESG portfolios, all closed-form analytical ones. But the real world is rather a black box. For this reason, we propose a different approach, Bayesian Optimization, which presents a highly flexible framework for optimizing ESG portfolio management criteria due to its adaptability in handling black-box functions. The main advantage of our approach is that it is 'expression-agnostic'. The flexibility of BO stems from its ability to adapt to different ESG criteria without requiring problem-specific adjustments, being capable of optimizing any function of any investor profile with any way of combining ESG and risk-return. Prior in the literature, other black-box techniques -including metaheuristics such as Genetic algorithm (GA) [17] or Simulated Annealing (SA) [18]- have also been proposed as alternative portfolio optimization techniques. Bayesian optimization (BO) is a state-of-the-art class of methods that optimize black-boxes. Although BO has many theoretical advantages over traditional black-box techniques, such as GA and SA, it has never been applied to portfolio optimization (and therefore not to ESG portfolio optimization). This is the research gap to which the present paper responds. This paper aims to analyze whether BO can be applied within the ESG portfolio optimization framework as an alternative technique and compare it with other traditional black-box techniques used in portfolio management: GA and SA. To show the empirical performance of BO, we carry out four illustrative experiments. The present study is the first paper (to the best of our knowledge) that leverages Bayesian optimization and ESG scores into a portfolio optimization technique. Our goal is to add another

different optimization approach to the world of ESG investing: a black-box non-heuristic optimization approach through BO. Our main contribution is to show that BO can be applied within the ESG portfolio optimization framework. In particular, this paper seeks to contribute to the literature on optimization techniques for ESG portfolios with an alternative approach by investigating whether BO can be used in the world of ESG investing, opening a further research line in portfolio allocation.

This paper is organized as follows. First, we begin with a state-of-the-art description of ESG portfolio optimization and Bayesian optimization methods. Then, we describe the components of the objective function that we will use in the illustrative experiment. In the following section, we describe Bayesian optimization in detail in order to understand the illustrative experiment section, that comes afterward. Finally, we close the manuscript with the conclusions and further work section.

2. State of the art in portfolio optimization with ESG and Bayesian optimization

Different traditional and widely known models for portfolio optimization have been applied to ESG portfolios [5] such as the Markowitz mean-variance portfolio optimization approach [16], or the Black and Litterman asset allocation model [19]. Likewise, more recent portfolio optimization methods or with a less solid mathematical basis have been used for the construction of ESG portfolios (the naïve diversification approach - or 1/N portfolio strategy- [5], the risk-parity portfolio framework [10, 20] and the reward-to-risk timing portfolio strategy [21]).

Markowitz mean-variance portfolio optimization model only considers risk and return and does not allow for additional criteria. The need for portfolio selection to be able to include criteria beyond mean and variance is solved with multi-criteria portfolio selection. Many multiple criteria methods have already been applied in the field of portfolio selection since Lee and Lerro [22]; we refer to Aouni *et al* (2018) for a review [23]. Concerning studies that apply multi-criteria methods to optimize ESG portfolios we can cite, among others, the following: managing ESG portfolios from a linear multicriteria approach [24]. A multicriteria approach but in a classical utility theory under uncertainty framework, instead of a linear one [25]. A two-stage multi-objective framework for the selection of ESG portfolios by applying a ‘Hedonic Price Method’, selecting ESG portfolios using goal programming models and fuzzy technology [26]. A model that combines goal programming with ‘goal games’ against nature [27]. A tri-criterion framework for inverse optimization of ESG portfolios [28]. An integration of the ESG portfolio selection problem into a Decision Support System [29]. A multicriteria portfolio selection model for mutual funds based on the Reference Point Method [30]. Similarly, a Markowitz’ model modification through a new tri-criterion model enabling investors to custom-tailor their asset allocations and incorporate all personal preferences regarding return, risk and social responsibility [31]. Additionally, a formulation of the portfolio optimization problem as a multiple-objective problem, where the third objective corresponds to corporate social responsibility [32]. Analogously, three adaptations of multiobjective evolutionary algorithms to include a social screening preceding the optimization process [33]. Finally, a hybrid ESG portfolio selection model with multi-criteria decision making (MCDM) and multi-objective optimization problem (MOOP) techniques [34].

Other authors propose ESG-adjusted capital asset pricing models: the Sustainable CAPM model (S-CAPM) [35, 36]. This leads us to ESG factor models or the ESG factor investing strand of literature which considers ESG criteria as a traditional systematic risk factor, either as a standalone factor or as a subcomponent of factor strategies [1, 6, 37–43]. Additionally, a data envelopment analysis (DEA) model with quadratic and cubic terms to enhance the evidence of two or more aspects, as well as the interaction between the environmental, social, and governance attributes has been proposed [3]. They then combined the ESG scores with financial indicators to select assets based on a cross-efficiency analysis.

All these asset allocation methods listed above are closed-form analytical techniques. But the real world is rather a black box. For this reason, we propose a different approach, Bayesian Optimization, which presents a highly flexible framework for optimizing ESG portfolio management criteria due to its adaptability in handling black-box functions. The main advantage of our approach is that it is ‘expression-agnostic’. The flexibility of BO stems from its ability to adapt to different ESG criteria without requiring problem-specific adjustments, being capable of optimizing any function of any investor profile with any way of combining ESG and risk-return. As will be explained in section 3.4, BO is capable of learning the underlying structure and uncertainty of the objective function, enabling efficient exploration and exploitation in the search space. This feature, combined with the L-Lipschitz continuity of ESG risk-performance functions, allows Bayesian optimization to optimize various ESG criteria with minimal prior knowledge and minimal assumptions about the function’s structure, making it a highly adaptable and flexible approach for ESG portfolio management. Prior in the literature, other black-box models -including metaheuristics such as Genetic algorithm (GA) [17] or Simulated Annealing (SA) [18]- have also been proposed as alternative portfolio optimization techniques [44–53]. Bayesian optimization

(BO) is a state-of-the-art class of methods that optimize black-boxes. Although BO has many theoretical advantages over traditional black-box techniques, such as GA and SA, it has never been applied to portfolio optimization (and therefore not to ESG portfolio optimization). Thus, this paper investigates whether BO can be used in the ESG portfolio optimization framework as an alternative and compares it with other traditional black-box techniques applied in portfolio optimization: GA and SA. This is the research gap to which the present paper responds. Our goal is to add another different optimization approach to the world of ESG investing: a black-box non-heuristic optimization approach through BO. While metaheuristics (such as GA and SA) and BO share that both do not require a closed analytic expression (they can deal with black boxes), four theoretical advantages of BO over GA and SA can be highlighted: • Metaheuristics are used when finding the certified optimum solution in a reasonable amount of time is impossible, but they only find near-optimal solutions [54]. In contrast, BO is not a heuristic but a method based on probability theory -and furthermore Bayesian-, which uses a surrogate probabilistic model (in our case a Gaussian process, which is a precious model since it is non-parametric and analytically closed) that for each possible portfolio (combination of weights) is giving us a predictive distribution of the value of the black box (in our case the Sharpe ratio modified with ESG values). • When applying BO it is fulfilled that in the event of slight variations in the portfolio weights, our cost function (our ESG-constrained Sharpe ratio) is smooth. • Whereas BO can be applied even in the case that the evaluation is very costly (in terms of computing time), which occurs most often in the real world (for example, if our risk estimation was not a standard deviation but some very expensive monte carlo simulation method, or if we estimated the risk by training a deep neural network, or in the case we assess the ESG from social networks), GA and SA could not be applied in these more realistic situations as they need a very high number of observations for the objective function and thus require huge computing time. • While BO makes it possible to model the latent function contaminated by noise (with the same portfolio each time different results are obtained), GA and SA can not model it. These theoretical advantages lead us to expect BO to perform better than GA and SA empirically. Thus, we carry out four illustrative experiments to show the empirical performance of BO compared to GA and SA. The present study is the first paper (to the best of our knowledge) that leverages Bayesian optimization and ESG scores into a portfolio optimization technique. Our main contribution is to show that BO can be applied within the ESG portfolio optimization framework. In particular, this paper seeks to contribute to the underdeveloped strand of literature on optimization techniques for ESG portfolios with an alternative approach by investigating whether BO can be used in the world of ESG investing, opening a further research line in portfolio optimization.

Recently, another alternative approach to ESG portfolio optimization has been suggested that applies deep reinforcement learning (DRL) [55]. In particular, they proposed a deep reinforcement learning model-that contains a Multivariate Bidirectional Long Short-Term Memory (LSTM) neural network-to predict stock returns for constructing an ESG portfolio. They called their new model Deep Responsible Investment Portfolio (DRIP). We are willing to explore deep reinforcement learning approaches to ESG portfolio optimization as a further line of research, as stated in section 6.

3. The objective function: Sharpe Ratio subject to ESG Criteria

In this section we will describe and motivate our objective function. In particular, we will minimize the Sharpe ratio of a portfolio under the presence of soft constraints consisting of ESG criteria. As we will further see, we include the ESG soft constraints in the objective function, acting as penalization criteria for the objective function. Interestingly, it is useful to model the optimization problem like this as these constraints are not hard, in the sense that they penalize the possible solutions but do not make them unfeasible. The following subsections describe the two components of the objective function: the ESG criteria and the Sharpe ratio. After describing them, we explain how do we integrate them in the objective function.

3.1. Fundamentals of ESG Criteria

ESG criteria are a framework for analyzing and assessing an organization's performance in environmental, social, and governance matters in comparison with its competitors [56].

The six most prominent ESG rating agencies are Sustainalytics ESG Risk Rating, MSCI ESG Ratings, Moody's ESG (formerly Vigeo-Eiris), Refinitiv (formerly Asset4), Bloomberg ESG Disclosures Scores, and S&P Global ESG Scores (formerly RobecoSAM). These agencies belong to some of the largest financial groups, such as Morningstar, MSCI or Bloomberg. The scores given to the different companies are industry oriented, which means they must be compared with their industry competitors in order to evaluate their compliance with ESG criteria. Some rating agencies, such as Sustainalytics ESG Risk Rating, divide the risk of not meeting ESG criteria into manageable and unmanageable risks. Manageable risks are then divided into managed risks and management gap. Management gap represents the potential improvement in compliance with ESG criteria companies have.



Table 1. Different aspects or categories in each of the 3 key ESG areas, evaluated by ESG rating agencies. (*) In the European Union (EU) the Non-Financial Reporting Directive (NFRD) regulates and establishes the extent to which companies in the EU must comply with corporate transparency [57].

Environmental	Climate Change	Carbon Emissions	Carbon and greenhouse emissions in all aspects of a firm's operations: transport, production or manufacturing
		Product's Carbon Footprint	And their vulnerability to climate change
	Natural Capital	Water Stress	Dependency and management of water resources in a company
		Land Use	Involves the modification of natural habitats into built environments or semi-habitat environments, such as farmland.
		Dependency on Raw Materials	In a company's production, manufacturing and operations
	Pollution & Waste		Toxic emissions, use of packaging material or electronic waste.
	Environmental Opportunities		Renewable energy or green building
Social	Human Capital		Employee well-being in a firm: health, safety, labor development, remuneration, etc.
	Product Liability		Chemical safety, product safety and quality, privacy and consumer protection, etc.
	Stakeholder Opposition		Opposition of stakeholders, such as local communities or suppliers, to a company.
	Social Opportunities		Healthcare or communication, etc.
Governance	Corporate Governance		Independence and diversity of the Board of Directors, executive compensation
	Corporate Behavior	Business Ethics	The "World's Most Ethical Companies" award
		Transparency	Access to quality corporate information for all stakeholders (*)

Rating agencies look mainly at different aspects or categories in each of the three key ESG areas when evaluating a company's ESG performance, such as the ones shown in table 1.

Rating agencies give scores in different metrics to represent a company's overall ESG risk. Just as an illustrative example, MSCI gives a score between CCC and AAA, as it can be seen in figure 1.

However, other agencies as Bloomberg give a [1–100] scale or Sustainalytics uses a 5 risk level range according to a [0–40] score from 10 to 10. This is because there is no benchmark to which comparing the accomplishment of the different criteria. Therefore, a company's ESG score doesn't provide much information about a firm's environmental, social and governance performance unless it is compared with the scores of other companies in the same sector or industry. Additionally, it is important to compare ESG scores of firms in the same sector, not across sectors, as there are industry related factors that affect a firm's ESG performance. For example, a company in an industry which requires high amounts of energy resources, such as a chemical company, is not comparable with a service company, such as a consultancy firm.

One of the criticisms made to ESG criteria derives from the fact that they fall short in taking into consideration the overall mission of corporations. Additionally, some of the criteria may be considered to be subjective, as analysts will have to score the companies according to their views and opinions on the performance on those particular areas. Therefore, in order to provide an ESG metric that encodes all, or the majority, of the

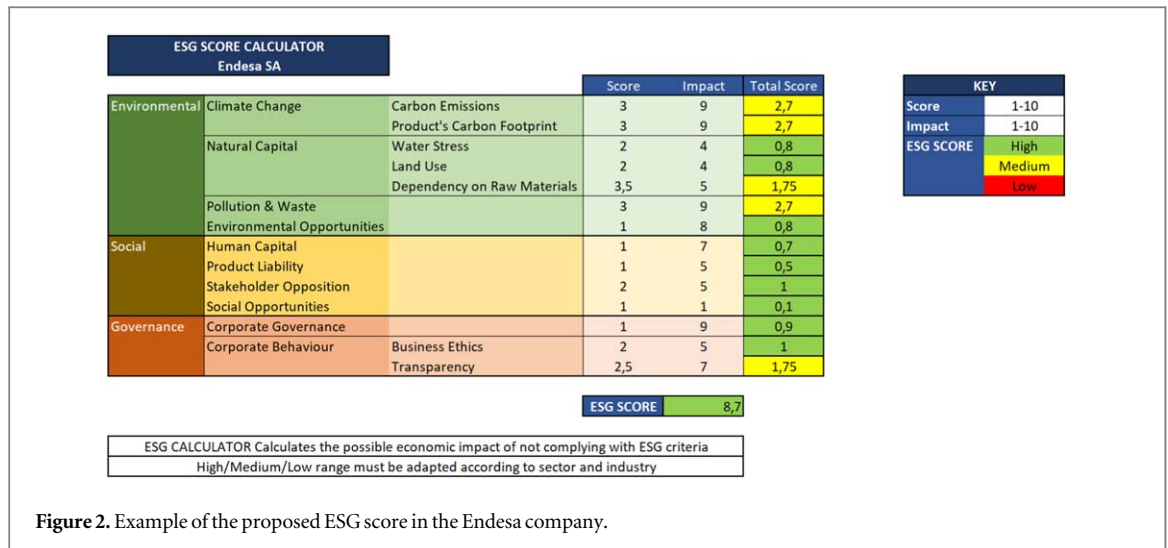


Figure 2. Example of the proposed ESG score in the Endesa company.

mentioned criteria, it is important to consider different ESG scores and approaches when assessing a firm’s performance in the three areas. For our work, as an illustrative example, we used the following ESG score, but we emphasize that our methodology is compatible with any subjective or objective ESG score, as we use Bayesian optimization that is able to deal with black-box functions whose gradients are unknown and whose values can be subjective.

In figure 2 the ESG Score of Endesa can be seen. Endesa received a score of 8,7, and had no category classified as ‘low’. The category where Endesa performs best according to ESG criteria is Product Liability, followed by Human Capital. Meanwhile, the categories where it performed worst are Carbon Emissions and Product’s Carbon Footprint. This is because, Endesa still has thermal plants operating, which are very pollutant, as they consume carbon. Additionally, it has combined cycles, which also emit greenhouse gases, although less than thermal plants, as they consume natural gas. Endesa’s ESG score is expected to improve in the following years, as it plans to close all its carbon businesses by 2027 and its combined cycles in 2040.

3.2. Sharpe ratio

The Sharpe ratio takes into consideration an asset’s return and its variance [58]. Therefore, it balances the trade-off between maximizing returns and minimizing the risk or volatility. Equation (1) shows the Sharpe Ratio used in the present research. The Sharpe ratio takes into consideration in the numerator the return and weight of the different assets in the portfolio and the risk-free rate. In the denominator it considers the covariance matrix of the portfolio and the weight of the different assets. The diagonal of the covariance matrix is the variance of the different assets. The covariance matrix is symmetric about the diagonal.

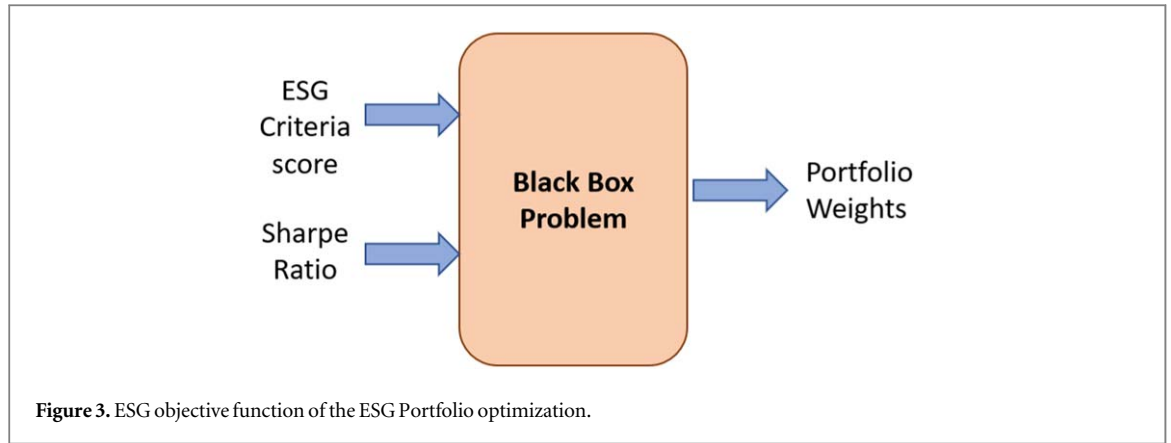
$$f_x = \frac{\sum_{i=1}^N w_i r_i - r_f}{\sigma_p} \quad s.t. \sum_{i=1}^N w_i = 1, 0 \leq w_i \leq 1, \quad (1)$$

where N is the number of different assets, w_i is the weight of each asset i in the portfolio, r_i is the return of asset i , r_f is the risk-free rate and σ_p is the standard deviation of the excess return of the portfolio.

3.3. ESG-constrained Sharpe Ratio

We combine the previous two concepts—ESG criteria and Sharpe ratio- in a single objective function that we will optimize with the method that will be illustrated in the following section, the bayesian approach. Most critically, we emphasize that this is only an illustrative example of how both criteria (risk and performance measured by Sharpe and ESG measured by the process mentioned in the previous section) can be combined. In practice, the critical added value of BO is that both ESG and Sharpe values can be obtained as expensive black-box processes. For example, social network analysis for ESG and Monte Carlo processes to estimate credit or market risks. For our illustrative case, in particular, the objective function tries to optimize the Sharpe ratio being penalized by the ESG criteria. The optimization of the objective function, being analytical or a potential black-box, will return the weights of the optimal portfolio in terms of performance and ESG compliance as figure 3 illustrates.

First, we assume that the money is fully invested and also without debts, so we ensure that the weights of the portfolio sum one $\sum_{i=1}^N w_i = 1$, where n is the number of titles in our portfolio, by performing a softmax



function on the weights vector \mathbf{w} . Geometrically, we are performing a bijection, transforming the hypercube $[0, 1]^n$ into a simplex space \mathcal{S}^n :

$$\mathbf{w}_i = \frac{\exp(\mathbf{w})}{\sum_{i=1}^N \exp(\mathbf{w}) + \epsilon} \quad (2)$$

where ϵ is a near-zero value to ensure computational robustness. Then, we compute the Sharpe ratio of the portfolio with respect to a risk-free asset. Combining both factors requires them to have the same magnitude in the objective function. For this purpose, they are normalized. The ESG score is normalized taking into consideration its maximum and minimum values which are 0 and 10. Analogously, the Sharpe ratio is also normalized using its maximum and minimum values. The ESG score e is simply obtained as a linear combination of the ESG scores of every single asset:

$$e = \sum_{i=1}^N w_i e_i. \quad (3)$$

In particular, the maximum and minimum values are estimated by their sample values on the dataset commented on in the illustrative experiment. Then, we simply add both factors. However, a logarithm factor can be added to the ESG factor and normalized again if we have low ESG scores to highly penalize the objective function and higher ESG scores for making the objective only slightly worse. More generally, due to the flexibility of BO, we could include any transformation of both metrics as our objective function, concretely, if $\phi(\mathbf{x})$ is any transformation of \mathbf{x} (following machine learning notation), like a logarithm or a cubic function, then, if o is the objective function, we could combine both factors as:

$$o(\mathbf{w}, \mathbf{e}) = \phi(\phi(r(\mathbf{w})), \phi(e(\mathbf{e}))), \quad (4)$$

where r is the Sharpe ratio of the portfolio, e is the ESG function and \mathbf{e} are the ESG scores. Recall that both the ESG and Sharpe factors could be added, multiplied, or combined as the investor decides. As the ESG score normalized and the Sharpe ratio normalized take values in the interval between 0 and 1, the fitness function takes values in the interval between 0 and 2.

3.4. Adaptability of the Bayesian optimization framework for any ESG criteria

As we have seen, Bayesian optimization with Gaussian Processes presents a highly flexible framework for optimizing ESG portfolio management criteria due to its adaptability in handling black-box functions. In particular, more formally and given a sufficiently complex family of covariance functions for the Gaussian process model, Bayesian optimization can optimize any L-Lipschitz continuous function. Concretely, in the context of ESG portfolio management, let $f(\mathbf{x})$ denote a risk-performance function, like the Sharpe ratio, that takes into account ESG factors, and assume that $f(\mathbf{x})$ is L-Lipschitz continuous. This means that there exists a constant $L \geq 0$ such that for any \mathbf{x}_1 and \mathbf{x}_2 in the input domain, $|f(\mathbf{x}_1) - f(\mathbf{x}_2)| \leq L \|\mathbf{x}_1 - \mathbf{x}_2\|$. In other words, the Lipschitz continuity property ensures that $f(\mathbf{x})$ exhibits a controlled rate of change, which is advantageous when optimizing over diverse ESG criteria.

More informally, the flexibility of Bayesian optimization stems from its ability to adapt to different ESG criteria without requiring problem-specific adjustments. By constructing surrogate models of the unknown risk-performance function $f(\mathbf{x})$ using Gaussian Processes, Bayesian optimization is capable of learning the underlying structure and uncertainty of the objective function, enabling efficient exploration and exploitation in the search space [59]. This feature, combined with the L-Lipschitz continuity of ESG risk-performance functions, allows Bayesian optimization to optimize various ESG criteria with minimal prior knowledge and

minimal assumptions about the function's structure, making it a highly adaptable and flexible approach for ESG portfolio management.

4. Fundamentals of Bayesian optimization

Bayesian optimization is a state-of-the-art class of methods that optimize black-boxes, that is, unknown noisy analytical functions that are very expensive to evaluate whether in time or computational resources [60]. For example, the estimation of the generalization error of machine learning algorithms with respect to their hyper-parameters is considered to be a black-box function and the first successful application of Bayesian optimization (BO) [61]. In order to solve such a scenario, we need a method coping with the optimization of a black-box without using gradients, in a small number of steps, and considering noise in the evaluations. More formally, the purpose of BO is to retrieve the optimum \mathbf{x}^* of a black-box function $f(\mathbf{x})$ where $\mathbf{x} \in \mathcal{X}$ and \mathcal{X} is the input space where $f(\mathbf{x})$ can be observed. In other words, we want to retrieve \mathbf{x}^* such that,

$$\mathbf{x}^* = \arg \min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}), \quad (5)$$

assuming minimization. We can define a BO method by the following tuple

$$\mathcal{A} = (\mathcal{M}, \alpha(\cdot), p(f(\mathbf{x})|\mathcal{D})), \quad (6)$$

where $f(\mathbf{x})$ is the black-box that we want to optimize, \mathcal{M} is a probabilistic surrogate model, $\alpha(\cdot)$ is an acquisition, or decision, function, $p(f(\mathbf{x})|\mathcal{D})$ is a predictive distribution of the observation of \mathbf{x} and $\mathcal{D} = \{(\mathbf{x}_i, y_i) | i = 1, \dots, t\}$ is the dataset of previous observations at iteration t .

To successfully solve this task, BO uses a probabilistic surrogate model, being a common option the Gaussian process (GP), of the target function. Concretely, a GP is a set of random variables (of potentially infinite size), any finite number of which have (consistent) joint Gaussian distributions [62]. More formally, a GP is fully defined by a zero mean and a covariance function or kernel $k(\mathbf{x}, \mathbf{x}')$, that is, $f(\mathbf{x}) \sim \mathcal{GP}(\mathbf{0}, k(\mathbf{x}, \mathbf{x}'))$. More concretely, the covariance function of the GP receives two points as an input, \mathbf{x} and \mathbf{x}' . Given a set of observed data $\mathcal{D} = \{(\mathbf{x}_i, y_i) | i = 1, \dots, N\}$, where $y_i = f(\mathbf{x}_i) + \epsilon_i$ with ϵ_i some additive Gaussian noise, a GP builds a Gaussian predictive distribution $p(f(\mathbf{x}^*)|\mathcal{D})$ for the potential values of $f(\mathbf{x}^*)$ at a new input point \mathbf{x}^* . Concretely, $p(f(\mathbf{x}^*)|\mathcal{D}) = \mathcal{N}(f(\mathbf{x}^*)|\mu(\mathbf{x}^*), v(\mathbf{x}^*))$. Lastly, the mean $\mu(\mathbf{x}^*)$ and variance $v(\mathbf{x}^*)$ of the predictive distribution $p(f(\mathbf{x}^*)|\mathcal{D})$ are respectively given by:

$$\mu(\mathbf{x}^*) = \mathbf{k}_*^T (\mathbf{K} + \sigma^2 \mathbf{I})^{-1} \mathbf{y}, \quad (7)$$

$$v(\mathbf{x}^*) = k(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}_*^T (\mathbf{K} + \sigma^2 \mathbf{I})^{-1} \mathbf{k}_*, \quad (8)$$

where $\mathbf{y} = (y_1, \dots, y_N)^T$ are the observations collected so far; σ^2 is the variance of the additive Gaussian noise ϵ_i ; $\mathbf{k}_* = \mathbf{k}(\mathbf{x}_*)$ is a N -dimensional vector with the prior covariances between the test point $f(\mathbf{x}^*)$ and each of the training points $f(\mathbf{x}_i)$; and \mathbf{K} is a $N \times N$ matrix with the prior covariances among each $f(\mathbf{x}_i)$, for $i=1, \dots, N$. Each element $\mathbf{K}_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$ of the matrix \mathbf{K} is given by the covariance function between each of the training points \mathbf{x}_i and \mathbf{x}_j where $i, j = 1, \dots, N$ and N is the total number of training points.

Using this model, we can estimate a predictive distribution of the unknown function in areas of the space where it has not been evaluated yet. Using this distribution, BO computes, iteratively, an acquisition function $\alpha(\mathbf{x})$ that estimates, for every input space point \mathbf{x} , the expected utility of evaluating the objective. In particular, the point whose value maximizes the acquisition function is suggested for evaluation in an iterative fashion. Most critically, that point maximizes the compromise between exploration of unknown areas and exploitation of promising solutions evaluated before. The acquisition function $\alpha(\mathbf{x})$ is generally not difficult to maximize. In particular, we can compute the gradient $\nabla_{\mathbf{x}} \alpha(\mathbf{x})$ of the acquisition function and use it for its optimization. We can compute the gradient $\nabla_{\mathbf{x}} \alpha(\mathbf{x})$ because the acquisition function $\alpha(\mathbf{x})$ is cheap to evaluate, as it is only based on the GP predictive distribution $p(f(\mathbf{x})|\mathcal{D})$. An example of acquisition function is the expected improvement (EI). Let $\chi(\mathbf{x}) = \frac{\mu(\mathbf{x}) - \kappa - \epsilon}{\sigma(\mathbf{x})}$, EI is given by:

$$\text{EI}(\mathbf{x}) = (\mu(\mathbf{x}) - \kappa - \epsilon) \Phi(\chi(\mathbf{x})) + \sigma(\chi(\mathbf{x})) \phi(\chi(\mathbf{x})), \quad (9)$$

where $\Phi(\cdot)$ is the Gaussian CDF and $\phi(\cdot)$ is the Gaussian PDF. Concretely, it is a heavily based exploitative criterion.

Afterwards, the GP is updated with the suggestion and conditioned there, obtaining a new predictive distribution and making the BO method repeat the described instructions iteratively, until a number of evaluations is consumed, where BO gives the final recommendation. In particular, this point can be the one whose evaluation has the best observed value or the point that optimizes the GP predictive mean. We summarize the steps of the basic Bayesian optimization method in algorithm 1.

Algorithm 1. BO of a black-box function $f(\mathbf{x})$.

Input: Maximum number of evaluations (budget) T .
for $t = 1, 2, 3, \dots, T$ **do**
 1: **if** $N = 1$:
 Choose \mathbf{x}_t randomly from \mathcal{X} .
 else:
 Find \mathbf{x}_t by optimizing the acquisition function: $\mathbf{x}_t = \operatorname{argmax}_{\mathbf{x} \in \mathcal{X}} \alpha_t(\mathbf{x})$.
 2: Evaluate the black-box function $f(\cdot)$ at \mathbf{x}_t : $y_t = f(\mathbf{x}_t) + \epsilon_t$.
 3: Augment the dataset with the observed value: $\mathcal{D}_{1:t} = \mathcal{D}_{1:t-1} \cup \{\mathbf{x}_t, y_t\}$.
 4: Fit again the GP model using the augmented dataset $\mathcal{D}_{1:t}$.
end
5: Obtain the recommendation \mathbf{x}^* (best observed value).
Result: Recommended point \mathbf{x}^*

We illustrate the described process on figure 4. In particular, this figure is divided into 3 plots, that are 3 iterations of the BO algorithm. Plotted on dotted black is the unknown objective function that we want to optimize. We also plot in black the predictive mean of the conditioned Gaussian process on the previous observations of the objective (black dots). The blue clouds are the uncertainty coming from the GP predictive distribution. Please observe how the uncertainty grows in unexplored areas of the space. Additionally, we plot as a red dot the suggestion of the Bayesian optimization algorithm. This observation comes at the exact point of the previous maximization of the acquisition function, that is represented as a green area. Recall that the acquisition function is easy to maximize, for example by a quasi-Newton second order method as the L-BFGS algorithm, as we have gradients. The represented process continues in an iterative fashion until a budget of operations is consumed, where the recommended solution comes from the minimization of the predictive mean of the Gaussian process.

Concretely, different acquisition functions, like the expected improvement or the upper confidence bound can be used. However, they are all trade-offs between exploration and exploitation. Analogously, different probabilistic surrogate models, like random forests or Bayesian neural networks can also be used. For more information about Bayesian optimization and its applications we conclude giving a reference to study this class of methods in more detail [59].

5. Illustrative experiments

We use the the Bayesian optimization method described in section 4 to optimize the ESG constrained Sharpe ratio objective function presented in section 3.3. However, for illustrative purposes, we use a function whose expression is known in this illustrative experiment. In a real-world scenario the expression for the performance of the stock portfolio may be only obtained by a Monte Carlo simulation of some events whose expression is unknown and the ESG score may be obtained by a webcrawler of the different companies at the current day of the simulation, which is also an unknown expression. Hence, only Bayesian optimization could be used in the described setting as a solution, as the analytical expression is unknown and costly. The illustrative experiment performed in this research consists of the optimization of a portfolio containing three companies from the Spanish electric utility and energy sector. These companies are: Endesa, Repsol and Iberdrola. The three companies compete on the electric utility sector, although there are some key differences between them, which will be analyzed. Repsol has a strong presence in the oil and gas sector, therefore it should have a lower ESG score, and the optimization should penalize Repsol with respect to the other two companies. The optimization algorithm considers two factors: the current ESG score of the three companies and their return in the last 12 months, from March 15th, 2021 to March 14th, 2022. We begin the section by analyzing its ESG score. We emphasize the notice already pointed out above: that the ESG scores cannot be computed using an analytical-form expression in other scenarios.

5.1. ESG analysis of the companies

Endesa, as seen in figure 2 and described in section 3.1, had an ESG score of 8,7. As it can be seen in figure 5, Repsol received an ESG score of 7,3, having the worst score out of the three companies. Repsol has three categories that were classified as having a ‘low’ ESG score, which were Carbon Emissions, Product’s Carbon Footprint and Pollution and Waste. This makes sense, as Repsol’s main business is the extraction, refinement and commercialization of hydrocarbons. The category in which Repsol performed best was Product Liability. Nevertheless, Repsol had some categories classified as ‘Medium’ in Social and in Governance. These

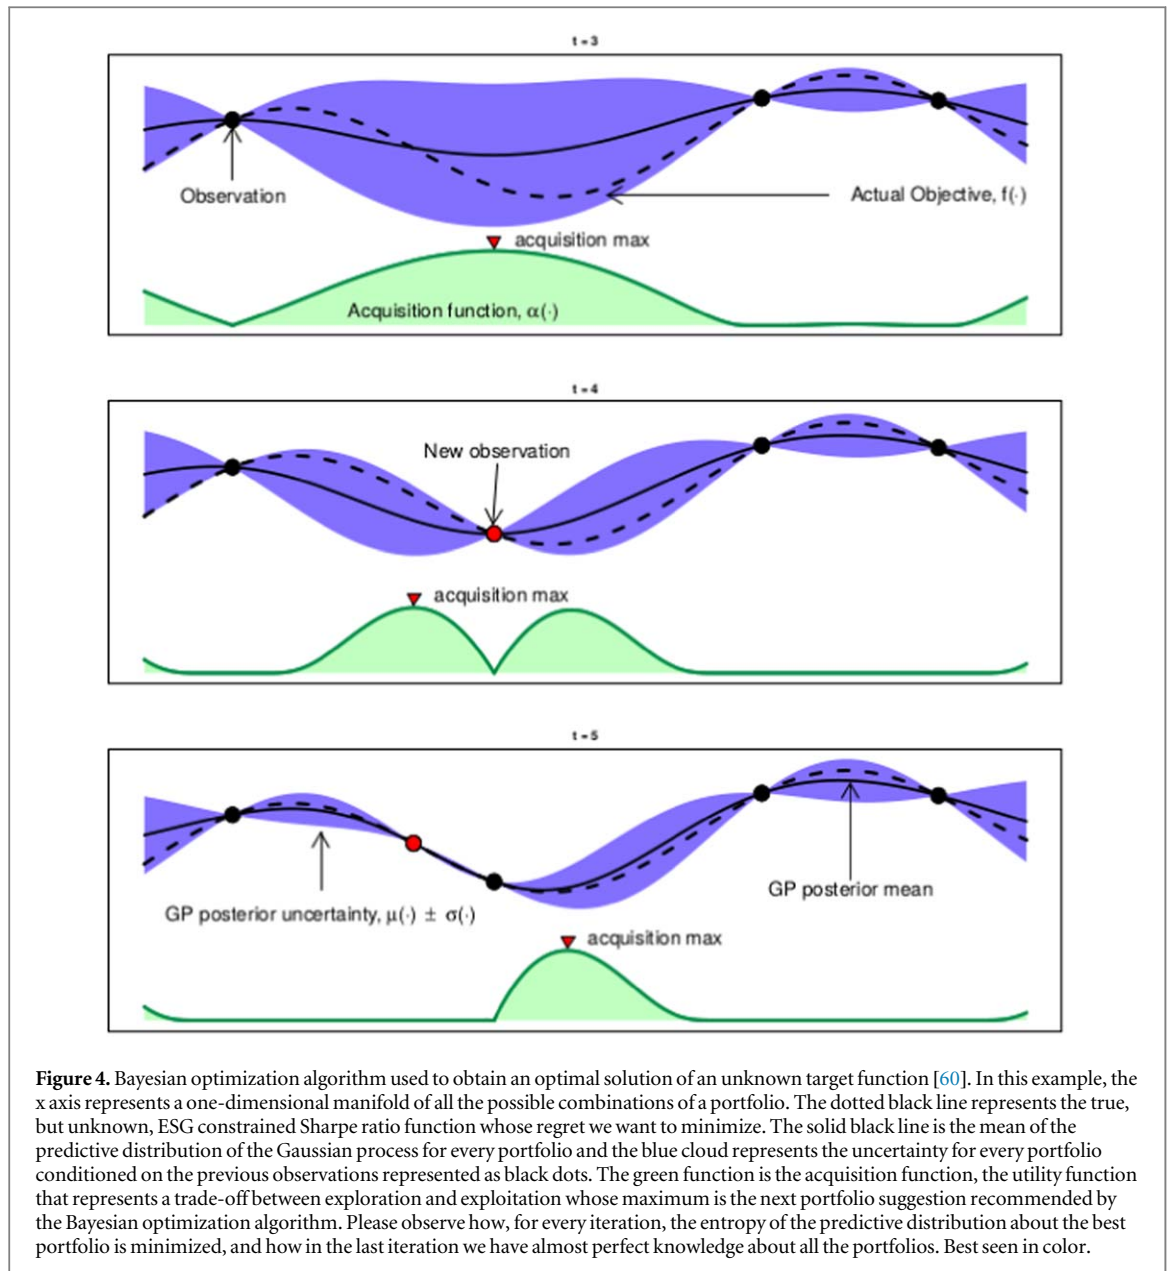


Figure 4. Bayesian optimization algorithm used to obtain an optimal solution of an unknown target function [60]. In this example, the x axis represents a one-dimensional manifold of all the possible combinations of a portfolio. The dotted black line represents the true, but unknown, ESG constrained Sharpe ratio function whose regret we want to minimize. The solid black line is the mean of the predictive distribution of the Gaussian process for every portfolio and the blue cloud represents the uncertainty for every portfolio conditioned on the previous observations represented as black dots. The green function is the acquisition function, the utility function that represents a trade-off between exploration and exploitation whose maximum is the next portfolio suggestion recommended by the Bayesian optimization algorithm. Please observe how, for every iteration, the entropy of the predictive distribution about the best portfolio is minimized, and how in the last iteration we have almost perfect knowledge about all the portfolios. Best seen in color.

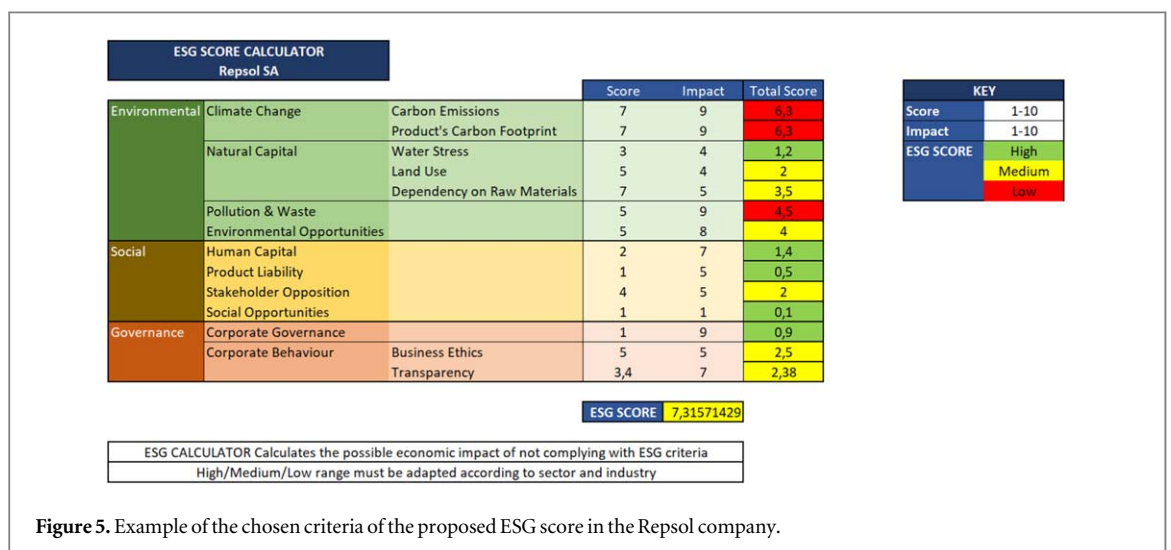


Figure 5. Example of the chosen criteria of the proposed ESG score in the Repsol company.

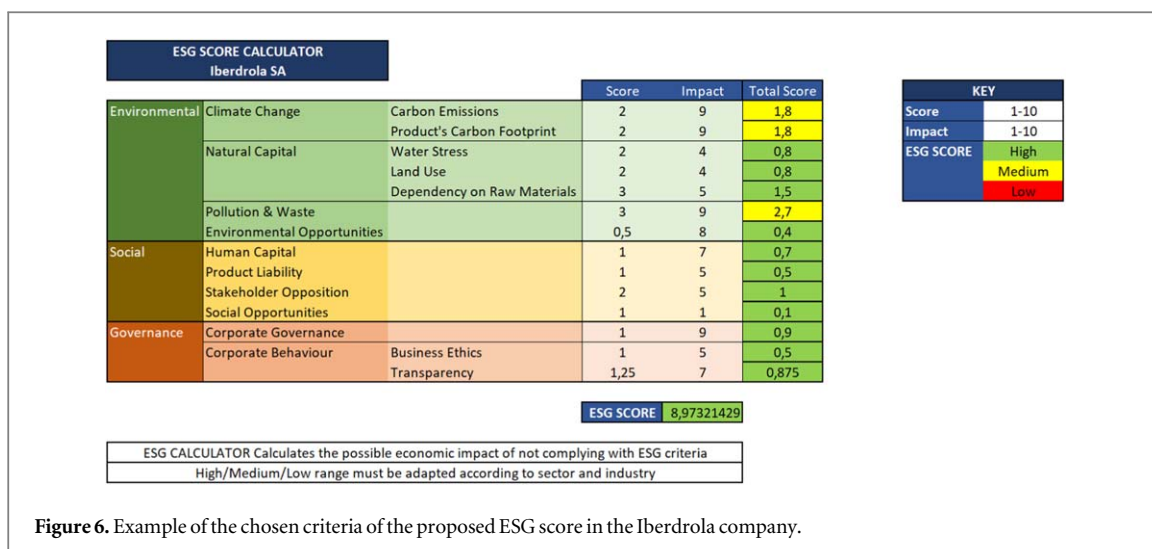


Figure 6. Example of the chosen criteria of the proposed ESG score in the Iberdrola company.

classifications have nothing to do with operating in the oil and gas sector, therefore it has a considerable margin of improvement in these categories. Repsol’s ESG score is also expected to improve in the coming years, as it plans to become carbon neutral in 2050.

In figure 6 the ESG score of Iberdrola can be seen. Iberdrola received a score of 8,97, being the company with the highest score. It had no category classified as ‘low’ and only three categories classified as ‘medium’. The categories in which Iberdrola performed best are Environmental Opportunities, Product Liability and Business Ethics. Iberdrola received an excellent score in Environmental Opportunities as it is the company which is planning to become carbon neutral earlier, in 2030, and is the one investing heavier in renewable energy. The categories in which Iberdrola performed worst are Carbon Emissions and Product’s Carbon Footprint, as Iberdrola still has combined cycles operating and commercializes electricity produced from all types of technologies, including thermal plants. Iberdrola performed better than Endesa in these categories, as it has a larger proportion of renewable energy resources and because it has no thermal plant operating. Iberdrola’s ESG score is expected to improve in the following years as it becomes carbon neutral.

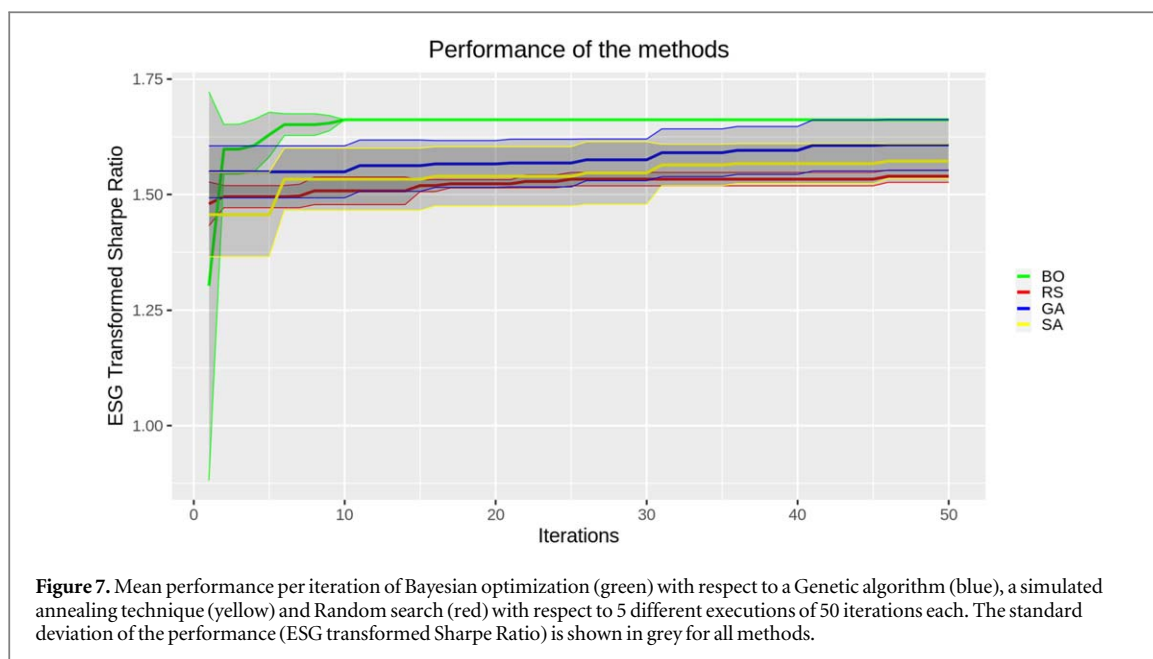
5.2. ESG constrained optimization of the Sharpe Ratio

Having computed the ESG scores, we now combine the companies and optimize the Sharpe ratio of the stock portfolio with respect to the participation of every company in the portfolio but constrained to the ESG scores. Concretely, the risk-free rate chosen for this experiment is 1.2% which was the average yield of the Spanish 10-year treasury bond during March 2022.

To normalize the Sharpe ratio, the maximum and minimum values need to be calculated. The formula used for the Sharpe ratio is represented and explained in section 3.3. To find the maximum and minimum value for the Sharpe ratio, the average returns and covariances for the three companies in the selected time period need to be computed. For Endesa we had an average return of -0.051%, for Repsol we had an average return of 0.036% and for Iberdrola we had an average return of -0.022%.

The maximum value of the Sharpe ratio is 3 and the minimum -60. We use this information to normalize the Sharpe ratio. The ESG score is normalized considering its maximum and minimum values, which are 10 and 0. In order to analyze the effectiveness of the Bayesian optimization algorithm, we compare the results of the optimization with the best score achieved when running 100 different random variables for each of the three weights. Both the optimization and the random search are run for a budget of 25 iterations. We use the upper confidence bound acquisition function. For the Bayesian optimization to be effective, the average value, which is the expected value, of the fitness function has to be higher than for the random search.

We propose two experiments with the same setting. In particular, we run Bayesian optimization and, as baselines that are able to adapt to any subjective analytical expression or a function without analytical expression, we run a Genetic algorithm and the Simulated Annealing technique that have both been previously used for portfolio optimization. Finally, we also run a Random search as another baseline. Each experiment is executed 5 times with different random seeds. Every repetition of the experiments runs 50 iterations of all the methods. We hope that the expected value, approximated by the empirical mean, of Bayesian optimization is higher than the one of the Genetic Algorithm, the Simulated Annealing, and the Random Search methods. Also, we hope that the standard deviation is lower in the case of Bayesian optimization, being a signal of a robust method.



In the case of Bayesian optimization we use the upper confidence bound acquisition function, in the genetic algorithm case we use a population of 5 individuals and 10 iterations, for a fair comparison with respect to Bayesian optimization. Critically, all the methods are only able to observe the objective function the same number of times. The simulated annealing technique has an inner loop of 5 iterations and an outer loop of 10 iterations. Finally, 50 points are retrieved randomly in a random search.

The first experiment involves the following ESG scores, for Endesa, 8.7, for Iberdrola, 8.97 and for Repsol, 7.32. In particular, we can see that the variability of these scores is not high according to its range. Figure 7 shows the mean performance of the best observed result in every iteration by Bayesian Optimization, in green, compared to the performance displayed by Random Search, in red, Genetic algorithms in blue and simulated annealing in yellow.

As can be seen, Bayesian optimization outperforms the baselines and Random Search, not only in performance but also in robustness, measured by the standard deviation on the mean performance. In particular, the best observed mean performance of Bayesian Optimization has been 1.662, the mean performance of the Genetic algorithm has been 1.597, the mean performance of Simulated Annealing has been 1.565 and the performance of Random Search has been 1.538. Moreover, we can observe how Bayesian optimization repetitions all converge into an optimal portfolio.



In particular, the optimum portfolio where all the repetitions of Bayesian optimization converge will be composed of 57,6% shares of Endesa, 21,2% shares of Iberdrola and 21,2% shares of Repsol.

We now give the details of a second experiment, where we set ESG scores with high variability. In particular, the new ESG scores are the following ones: [9,5,2]. We also execute this second experiment 5 executions where every different execution includes 50 iterations of Bayesian Optimization, a Genetic Algorithm, Simulated Annealing and Random Search. Figure 8 summarizes the obtained results:

We can see how, independently of the variability of the ESG scores, Bayesian Optimization outperforms Genetic Algorithm, Simulated Annealing and Random Search, in that order, both in performance (mean of the ESG Transformed Sharpe Ratio) and robustness. Again, Bayesian Optimization repetitions converge into an optimal portfolio, whose ESG Transformed Sharpe Ratio is lower than in the case of the previous portfolio as a result of considering a lower sum of ESG scores, which penalizes the ESG Transformed Sharpe Ratio fitness function.


Consequently, given the empirical evidence shown by the previous illustrative experiments, we can conclude that, for the ESG constrained portfolio optimization that makes any possible objective function be able to use, Bayesian optimization could be used to design an optimal portfolio.

In a third experiment, we run 5 more executions with the first configuration of ESG score but this time we let the metaheuristic techniques have a budget of 100 evaluations to see whether they achieve better results than BO if more observations of the objective function are considered. Figure 9 shows the mean performance of the best observed result in every iteration by Bayesian Optimization, in green, compared to the performance displayed by Random Search, in red and Genetic algorithms in blue. We skip simulated annealing as it delivers worst results than genetic algorithms in previous experiments.



Finally, we carried out another synthetic experiment where we simulated the returns of 5 companies and randomly sampled their ESG scores to see whether, on the 5-dimensional case, BO still outperforms the genetic algorithm and random search. Interestingly, BO is still able to outperform genetic algorithms and random search as shown in figure 10.

Critically, we have only used the upper confidence bound vanilla Bayesian optimization method, but several enhancements like predictive entropy search may be used to enhance the performance of Bayesian optimization. In particular, the vanilla Bayesian optimization method has great empirical evidence for a maximum of 7, or 8 dimensions [60], for more dimensions, meaning more entities in the portfolio, we will explore the possibility of high-dimensional BO and compare it to the performance of genetic algorithms in further work. In particular, while our study has primarily focused on portfolio optimization involving a small number of assets (3 to 5 companies), we acknowledge that in practical scenarios, investors often need to allocate resources across a much larger number of assets. To address this, several high-dimensional Bayesian optimization (BO) methods have been developed that can effectively handle problems with dimensions beyond the empirical limit of 7 dimensions, or in this case, assets, which is the common threshold for vanilla Gaussian process BO and that we describe in the following paragraph.



Concretely, one such approach is known as Bayesian optimization in high dimensions via random embeddings [63]. More formally, in this method, the high-dimensional optimization problem is transformed into a series of lower-dimensional problems through random projections, which can then be efficiently solved using standard BO techniques. Mathematically, given a high-dimensional function $f(\mathbf{x})$, where $\mathbf{x} \in \mathbb{R}^D$ and D is the number of dimensions, random embeddings define a lower-dimensional function $g(\mathbf{z}) = f(\Phi\mathbf{z})$, where $\mathbf{z} \in \mathbb{R}^d$, $d \ll D$, and Φ is a D^*d random matrix. Then, BO is performed on $g(\mathbf{z})$ in the lower-dimensional space, yielding more efficient optimization while maintaining the overall structure of the original problem. Another approach to tackle high-dimensional optimization is the use of additive Gaussian process (GP) models [64]. The main idea is to decompose the objective function into additive components, allowing the GP model to learn the structure of each component separately. Formally, the objective function $f(\mathbf{x})$ is approximated as $f(\mathbf{x}) = \sum_i f_i(\mathbf{x}_i)$, where $f_i(\mathbf{x}_i)$ are lower-dimensional functions. This decomposition allows the BO algorithm to focus on optimizing one component at a time, significantly reducing the computational complexity and making it suitable for high-dimensional problems.

Further approaches include the use of sparse GPs [65] and deep GPs [66] approximated by variational inference, which enable scalable BO by introducing structured approximations to the standard GP models. By leveraging these advanced techniques, high-dimensional BO methods can be effectively utilized for optimizing portfolios with a large number of assets, providing a more comprehensive solution to real-world portfolio management problems. We believe that the comparison and adaptation of these methods for portfolio optimization of a high number of assets is a great further line of research.

6. Conclusions and further work

This paper has applied Bayesian optimization to an ESG constrained stock portfolio scenario. The optimization's effectiveness has been proved in four illustrative experiments. The proposed method can be used to penalize the behavior of a portfolio according to any criteria, such as Corporate Social Responsibility (CSR) criteria or a firm's exposure to a certain country. For example, these penalizations could be used to take into account the consequences of being present in the Russian and Ukrainian markets during the current war for multinational firms

Another interesting line of further research consists of decoupling the ESG constraints from the main objective and considering them as black-box constraints. Then, we will be able to avoid choosing at all costs different stock portfolio configurations using constrained multi-objective Bayesian optimization [67]. We would also like to compare the performance of Bayesian Optimization with respect to other black-box optimization methods such as Tree Parzen Estimator and more complex Bayesian Optimization methodologies. Additionally, a great further line of research consisting of applying and adapting several high-dimensional BO methods for optimizing portfolios with a high number of assets, comparing their performance and also comparing them to the performance of genetic algorithms in further work. Finally, we are willing to explore

more complex Bayesian Optimization strategies and deep reinforcement learning approaches to ESG portfolio optimization as a further line of research.

Acknowledgments

We thank the Area Editor and two anonymous reviewers who provided very insightful comments to help improve the manuscript.

Data availability statement

The data that support the findings of this study are available upon reasonable request from the authors.

Disclosure statement

The authors declare no conflict of interest.

ORCID iDs

María Coronado Vaca  <https://orcid.org/0000-0002-0315-3753>

References

- [1] Fan J H and Michalski L 2020 Sustainable factor investing: where doing well meets doing good *International Review of Economics & Finance* **70** 230–56
- [2] Gibson Brandon R, Glossner S, Krueger P, Matos P and Steffen T 2022 Do responsible investors invest responsibly? *Review of Finance* **26** 1389–432
- [3] Chen L, Zhang L, Huang J, Xiao H and Zhou Z 2021 Social responsibility portfolio optimization incorporating ESG criteria *Journal of Management Science and Engineering* **6** 75–85
- [4] Gutsche G and Ziegler A 2019 Which private investors are willing to pay for sustainable investments? empirical evidence from stated choice experiments *Journal of Banking & Finance* **102** 193–214
- [5] Oikonomou I, Platanakis E and Sutcliffe C 2018 Socially responsible investment portfolios: does the optimization process matter? *The British Accounting Review* **50** 379–401
- [6] Melas D, Nagy Z and Kulkarni P 2017 Factor investing and ESG integration *Factor Investing. From Traditional to Alternative Risk Premia* ed E Jurczenko (Oxford: Elsevier) pp 389–413
- [7] Hitchens R, McCullagh S and Parks C 2015 Finding alpha in ESG. Credit Suisse, ESG- α series
- [8] Bollen N P 2007 Mutual fund attributes and investor behavior *Journal of Financial and Quantitative Analysis* **42** 683–708
- [9] United Nations principles for responsible investment 2022 <https://unpri.org/about-us/what-are-the-principles-for-responsible-investment>
- [10] Anderson R M, Bianchi S W and Goldberg L R 2012 Will my risk parity strategy outperform? *Financial Analysts Journal* **68** 75–93
- [11] United Nations 2023 *The XII United Nations Global Compact-Accenture CEO Study on Sustainability. Unlocking the global pathways to resilience, growth, and sustainability for 2030* <https://accenture.com/content/dam/accenture/final/accenture-com/document/Accenture-CEO-Study-United-Nations-Global-Compact.pdf>
- [12] Morningstar 2022 *Sustainable funds U.S. landscape report 2021* <https://morningstar.com/lp/sustainable-funds-landscape-report,2022>
- [13] Macquarie Asset Management 2022 *Macquarie Asset Management. MAM 2021 ESG survey report* <https://mirafunds.com/au/en/our-approach/sustainability/macquarie-asset-management-2021-ESG-survey.html>
- [14] Atz U, Van Holt T, Liu Z Z and Bruno C C 2022 Does sustainability generate better financial performance? review, meta-analysis, and propositions *Journal of Sustainable Finance & Investment* **13** 802–825
- [15] Whelan T, Atz U, Van Holt T and Clark C 2021 *ESG and financial performance: uncovering the relationship by aggregating evidence from 1,000 plus studies published between 2015-2020* NYU Stern Center for Sustainable Business and Rockefeller Asset Management
- [16] Markowitz H 1952 Portfolio selection *Journal of Finance* **7** 77–91
- [17] Holland J H 1975 *Adaptation in Natural and Artificial Systems* (Ann Arbor: University of Michigan Press)
- [18] Kirkpatrick S, Gelatt C D Jr. and Vecchi M P 1983 Optimization by simulated annealing *Science* **220** 671–80
- [19] Black F and Litterman R 1992 Global portfolio optimization *Financial Analysts Journal* **48** 28–43
- [20] Baker M, Bradley B and Wurgler J 2011 Benchmarks as limits to arbitrage: Understanding the low-volatility anomaly *Financial Analysts Journal* **67** 40–54
- [21] Kirby C and Ostdiek B 2012 It's all in the timing: simple active portfolio strategies that outperform naive diversification *Journal of Financial and Quantitative Analysis* **47** 437–67
- [22] Lee S M and Lerro A 1973 Optimizing the portfolio selection for mutual funds *The Journal of Finance* **28** 1087–101
- [23] Aouni B, Doumpos M, Pérez-Gladish B and Steuer R E 2018 On the increasing importance of multiple criteria decision aid methods for portfolio selection *Journal of the Operational Research Society* **69** 1525–42
- [24] Hallerbach W, Ning H, Soppe A and Spronk J 2004 A framework for managing a portfolio of socially responsible investments *European Journal of Operational Research* **153** 517–29

- [25] Ballester E, Bravo M, Pérez-Gladish B, Arenas-Parra M and Pla-Santamaria D 2012 Socially responsible investment: A multicriteria approach to portfolio selection combining ethical and financial objectives *European Journal of Operational Research* **216** 487–94
- [26] Bilbao-Terol A, Arenas-Parra M and Cañal-Fernández V 2012 Selection of socially responsible portfolios using goal programming and fuzzy technology *Information Sciences* **189** 110–25
- [27] Trenado M, Romero M, Cuadrado M L and Romero C 2014 Corporate social responsibility in portfolio selection: a goal games against nature approach *Computers & Industrial Engineering* **75** 260–265
- [28] Utz S, Wimmer M, Hirschberger M and Steuer R E 2014 Tri-criterion inverse portfolio optimization with application to socially responsible mutual funds *European Journal of Operational Research* **234** 491–8
- [29] Calvo C, Ivorra C and Liern V 2015 Finding socially responsible portfolios close to conventional ones *International Review of Financial Analysis* **40** 52–63
- [30] Méndez-Rodríguez P, Pérez-Gladish B, Cabello J M and Ruiz F 2015 Mutual funds' socially responsible portfolio selection with fuzzy data *Multiple Criteria Decision Making in Finance, Insurance and Investment*. ed M Al-Shammari and H Masri (Cham: Springer) 229–47
- [31] Gasser S M, Rammerstorfer M and Weinmayer K 2017 Markowitz revisited: social portfolio engineering *European Journal of Operational Research* **258** 1181–90
- [32] Qi Y 2018 On outperforming social-screening-indexing by multiple-objective portfolio selection *Annals of Operational Research* **267** 493–513
- [33] Liagkouras K, Metaxiotis K and Tsihrintzis G 2020 Incorporating environmental and social considerations into the portfolio optimization process *Annals of Operations Research* **316** 1493–1518
- [34] Wu Q, Liu X, Qin J, Zhou L, Mardani A and Deveci M 2022 An integrated multi-criteria decision-making and multi-objective optimization model for socially responsible portfolio selection *Technological Forecasting and Social Change* **184** 121977
- [35] Pedersen L H, Fitzgibbons S and Pomorski L 2021 Responsible investing: the ESG-efficient frontier *Journal of Financial Economics* **142** 572–97
- [36] Zerbib O D 2022 A sustainable capital asset pricing model (S-CAPM): evidence from environmental integration and sin stock exclusion *Review of Finance* **26** 1345–88
- [37] Kaiser L 2020 ESG integration: value, growth and momentum *Journal of Asset Management* **21** 32–51
- [38] Halbritter G and Dorfleitner G 2015 The wages of social responsibility—where are they? a critical review of ESG investing *Review of Financial Economics* **26** 25–35
- [39] Herwall J, Guenster N, Bauer R and Koedijk K 2005 The eco-efficiency premium puzzle *Financial Analysts Journal* **61** 51–63
- [40] Görgen M, Wilkens M and Ohlsen H 2020 CARIMA—a capital market-based approach to quantifying and managing transition risks (UA and VfU) *Case Studies of Environmental Risk Analysis Methodologies*. ed M Jun et al (Paris: Central Bank and Supervisors Network for Greening the Financial Sector (NGFS)) pp 540–59
- [41] Roncalli T, Guenedal T L, Lepetit F, Roncalli T and Sekine T 2020 Measuring and managing carbon risk in investment portfolios. arXiv:2008.13198
- [42] Alessi L, Ossola E and Panzica R 2021 What greenium matters in the stock market? the role of greenhouse gas emissions and environmental disclosures *Journal of Financial Stability* **54** 100869
- [43] Gimeno R and González C I 2022 The Role of a Green Factor in Stock Prices. When Fama & French Go Green. Banco de España Working Paper No. 2207, Available at SSRN: <https://ssrn.com/abstract=4064848> or [10.2139/ssrn.4064848](https://ssrn.com/abstract=10.2139/ssrn.4064848)
- [44] Chang T-J., Meade N, Beasley J E and Sharaiha Y M 2000 Heuristics for cardinality constrained portfolio optimisation *Computers & Operations Research* **27** 1271–302
- [45] Chang T-J., Yang S C and Chang K J 2009 Portfolio optimization problems in different risk measures using genetic algorithm *Expert Systems with Applications* **36** 10529–37
- [46] Chan M C, Wong C C, Cheung B K and Tang G Y 2002 Genetic algorithms in multi-stage portfolio optimization system. *Proceedings of the Eighth International Conference of the Society for Computational Economics, Computing in Economics and Finance, Aix-en-Provence, France*
- [47] Vasiani V, Handari B and Hertono G 2020 Stock portfolio optimization using priority index and genetic algorithm *Journal of Physics: Conference Series* **1442** 012031
- [48] Frausto Solis J, Purata Aldaz J L, González del Angel M, González Barbosa J and Castilla Valdez G 2022 Saipo-taipo and genetic algorithms for investment portfolios *Axioms* **11** 42
- [49] Crama Y and Schyns M 2003 Simulated annealing for complex portfolio selection problems *European Journal of Operational Research* **150** 546–71
- [50] Serafini P 1994 Simulated annealing for multi objective optimization problems. *Multiple Criteria Decision Making: Proceedings of the Tenth International Conference: Expand and Enrich the Domains of Thinking and Application* (Springer) 283–92
- [51] Friesz T L, Anandalingam G, Mehta N J, Nam K, Shah S J and Tobin R L 1993 The multiobjective equilibrium network design problem revisited: a simulated annealing approach *European Journal of Operational Research* **65** 44–57
- [52] Rubio-García A, García-Ripoll J J and Porras D 2022 Portfolio optimization with discrete simulated annealing. arXiv:2210.00807
- [53] Lang J, Zielinski S and Feld S 2022 Strategic portfolio optimization using simulated, digital, and quantum annealing *Applied Sciences* **12** 12288
- [54] di Tollo G and Roli A 2016 Local search algorithms for portfolio selection: search space and correlation analysis *Recent Advances in Computational Optimization: Results of the Workshop on Computational Optimization WCO 2015*. (Cham: Springer) (Part of the Studies in Computational Intelligence book series) 655 21–38
- [55] Vo N N, He X, Liu S and Xu G 2019 Deep learning for decision making and the optimization of socially responsible investments and portfolio *Decision Support Systems* **124** 113097
- [56] Hvidkjær S 2017 ESG investing: a literature review *Report Prepared for Dansif*
- [57] Directive 2014/95/EU of the European Parliament and of the Council of 22 October 2014, amending Directive 2013/34/EU as regards disclosure of non-financial information and diversity information by certain groups of companies and certain groups 2014
- [58] Sharpe W F 1994 The Sharpe ratio *The Journal of Portfolio Management* **1** 49–58
- [59] Shahriari B, Swersky K, Wang Z, Adams R P and De Freitas N 2015 Taking the human out of the loop: a review of bayesian optimization *Proceedings of the IEEE* **104** 148–757352306
- [60] Garrido Merchán E C 2021 *Advanced methods for bayesian optimization in complex scenarios*. Doctoral Thesis. Universidad Autónoma de Madrid (UAM), Higher Polytechnique School, Computer Science Department
- [61] Snoek J, Larochelle H and Adams R P 2012 Practical bayesian optimization of machine learning algorithms arXiv:1206.2944v2

- [62] Williams C K and Rasmussen C E 2006 *Gaussian Processes for Machine Learning*. (Cambridge, MA: MIT Press) (<https://doi.org/10.7551/mitpress/3206.001.0001>)
- [63] Wang Z *et al* 2013 Bayesian optimization in high dimensions via random embeddings. *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence (IJCAI)*, vol 13, 1778–84
- [64] Kandasamy K, Schneider J and Póczos B 2015 High dimensional bayesian optimisation and bandits via additive models. *Proceedings of the 32nd International Conference on Machine Learning (PMLR)* 37 295–304
- [65] Titsias M 2009 Variational learning of inducing variables in sparse gaussian processes. *Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics, (AISTATS)* 5 567–74
- [66] Damianou A and Lawrence N D 2013 Deep gaussian processes. *Proceedings of the Sixteenth International Conference on Artificial Intelligence and Statistics (AISTATS)* 31 207–15
- [67] Fernández-Sánchez D, Garrido-Merchán E C and Hernández-Lobato D 2020 Improved max-value entropy search for multi-objective bayesian optimization with constraints arXiv:[2011.01150](https://arxiv.org/abs/2011.01150)