

Multi-Objective Bayesian Optimization of Deep Reinforcement Learning for Environmental, Social, and Governance (ESG) Financial Portfolio Management

E.C. Garrido Merchán; S. Mora-Figueroa Cruz-Guzmán; M. Coronado Vaca

Abstract-

Financial portfolio management focuses on the maximization of several objectives in a trading period related not only to the risk and performance of the portfolio but also to other objectives such as the environment, social, and governance (ESG) score of the portfolio. Regrettably, classic methods such as the Markowitz model do not take into account ESG scores but only the risk and performance of the portfolio. Moreover, the assumptions made by this model about the financial returns make it unfeasible to be applicable to markets with high volatility such as the technological sector. This paper investigates the application of deep reinforcement learning (DRL) for ESG financial portfolio management. DRL agents circumvent the issue of classic models in the sense that they do not make assumptions like the financial returns being normally distributed and are able to deal with any information like the ESG score if they are configured to gain a reward that makes an objective better. However, the performance of DRL agents has high variability, and it is very sensible to the value of their hyperparameters. Bayesian optimization is a class of methods that are suited to the optimization of black-box functions, that is, functions whose analytical expression is unknown and are noisy and expensive to evaluate. The hyperparameter tuning problem of DRL algorithms perfectly suits this scenario. As training an agent just for one objective is a very expensive period, requiring millions of timesteps, instead of optimizing an objective being a mixture of a risk-performance metric and an ESG metric, we choose to separate the objective and solve the multi-objective scenario to obtain an optimal Pareto set of portfolios representing the best trade-off between the Sharpe ratio and the ESG mean score of the portfolio and leaving to the investor the choice of the final portfolio. We conducted our experiments using environments encoded within the OpenAI Gym, adapted from the FinRL platform. The experiments are carried out in the Dow Jones Industrial Average (DJIA) and the NASDAQ markets in terms of the Sharpe ratio achieved by the agent and the mean ESG score of the portfolio. We compare the performance of the obtained Pareto sets in hypervolume terms illustrating how portfolios are the best trade-off between the Sharpe ratio and mean ESG score. Also, we show the usefulness of our proposed methodology by comparing the obtained hypervolume with one achieved by a random search methodology on the DRL hyperparameter space.

Index Terms-

Due to copyright restriction we cannot distribute this content on the web. However, clicking on the next link, authors will be able to distribute to you the full version of the paper:

[Request full paper to the authors](#)

If your institution has an electronic subscription to *Intelligent Systems in Accounting, Finance and Management*, you can download the paper from the journal website:

[Access to the Journal website](#)

Citation:

Garrido-Merchán, E.C.; Mora-Figueroa, S.; Coronado-Vaca, M. "Multi-Objective Bayesian Optimization of Deep Reinforcement Learning for Environmental, Social, and Governance (ESG) Financial Portfolio Management", Intelligent Systems in Accounting, Finance and Management, vol.32, no.2, pp.e70008-1-e70008-15, June, 2025.