Multi-Objective Evolutionary Computation for the Portfolio Optimization Problem with Respect to Environmental, Social, and Governance Criteria

Riley Herman and Malek Mouhoub[a]
Department of Computer Science, University of Regina, Regina, Canada

Keywords: Multi-Objective Evolutionary Computation (MOEA), Metaheuristics, Portfolio Optimization.

Abstract: A common problem facing many is the tension between doing what aligns with our values and doing what is fiscally best. We propose a system leveraging Multi-Objective Evolutionary Computation, specifically MOEA/D, to produce highly performant portfolios tailored to an individual's Environmental, Social, and Governance (ESG) preferences given a custom survey that we have designed. The survey is conducted to construct a weighting to normalize a given investor's own responses and allow a single portfolio from the collection of the best portfolios to be matched to that investor. We have adopted two potential architectures to build our proposed system: Architecture 1, where the optimization is run for each investor that takes the survey, and Architecture 2 where a multi-objective optimization is run less frequently and the investor is given a portfolio from the Pareto front. This subset consists of all the non-dominated portfolios. The user may have different experiences, including quality or response time, depending on the architecture chosen. The results of the experiments we conducted demonstrate that both architectures performed comparably and produced high-quality portfolios. However, the best portfolio from Architecture 2 was better in most respects than any portfolio from Architecture 1. All Architecture 1 portfolios were more significantly tailored to each of the individuals' preferences. For Architecture 2, a limited number of high performing portfolios was generated: as a result, more investors would potentially be recommended to the same few portfolios, especially in comparison to Architecture 1.

1 INTRODUCTION

A common problem facing many is the tension between doing what aligns with our values and doing what is fiscally best. In our age of mass information, the availability of a publicly traded company's ESG (Environment, Social, Governance) data has resulted in a demand for financial portfolios that reflect the values of consumers. With the gluttony of choices in stocks and other investment vehicles as well as the gluttony of data on publicly listed companies, it is an overwhelming task to compute portfolios that are performant as well as personally ethical. The majority of researchers exploring the portfolio optimization problem as an example of a multi-objective optimization problem are using two objective optimizations, as Markowitz first proposed: risk and return. Here, the problem is extended to better reflect the real world and to produce a practical and useful application. Independent objectives of environmental, social, and governance scores are added to further address the needs of data-savvy investors. In addition, risk is split to capture two types of risk measurement to better encapsulate different types of risk.

In addition, this paper presents a survey. This survey gathers and synthesizes data on the prospective investor's preferences towards ESG, as well as their risk tolerance. The survey was necessary to answer a question: What score defines a stock that performs "well" with respect to an ESG criteria? The same survey can be reused as input for the potential investor to be categorized and given a final singular portfolio from the collection of all the best performing portfolios.

In order for this research to be as practical as possible, two potential architectures of a recommendation system are presented. In much of the literature, the step from theory to practical application in the real world is not considered. Given the possible types of data one may attempt to use, the hardware available

[a] https://orcid.org/0000-0001-7381-1064

Herman, R. and Mouhoub, M.
Multi-Objective Evolutionary Computation for the Portfolio Optimization Problem with Respect to Environmental, Social, and Governance Criteria.
DOI: 10.5220/0013604500003970
In *Proceedings of the 15th International Conference on Simulation and Modeling Methodologies, Technologies and Applications (SIMULTECH 2025)*, pages 357-365
ISBN: 978-989-758-759-7; ISSN: 2184-2841

to a potential implementer, or even the market segment one may try to target, these two architectures could both be successful in different situations. This choice requires a modification to the model, and as such, both modifications are presented. In sum, the following contributions are presented.

- Two architectures for a portfolio recommender system
- A survey for the gathering of risk and ESG preferences
- A problem definition grounded in the literature

2 PORTFOLIO OPTIMIZATION

The portfolio optimization problem is as follows: Given a set of investment vehicles and some corresponding information about each investment vehicle (here, the potential risk, return, and ESG scores), compute a portfolio of investment vehicles such that the total return is maximized, risk is minimized, and the portfolio does not cost more than a given budget. This is also extended to maximize positive environmental, social and governance criteria.

This problem finds its roots in Markowitz. Many different routes have been traversed on this problem: This is not a comprehensive history of attempts at solving this problem, but should give ample justification for the techniques explored here. Mathematically speaking, of course, this is a solved problem. The issue is no longer whether an optimal portfolio can be found, but rather how quickly a sufficiently optimal portfolio can be found.

Several metaheuristics (Lynn and Suganthan, 0 01; Darmstadt and Approximity, 2003; Korani and Mouhoub, 2021; Korani and Mouhoub, 2022; Sadreddin et al., 2022) have been applied to solve portfolio optimization problems. The first of which was PSO in the mid-1990s (Darmstadt and Approximity, 2003). Interestingly, in 2015 some improvements to the original PSO were published that apply much closer to the recommender system described here, namely Heterogeneous Multiple Population Particle Swarm Optimization. The idea is to break the initial population into smaller sub-variants of the larger PSO problem. (Lynn and Suganthan, 0 01)

Ant Colony Optimization underwent similar treatment on the same cycle; improvements to expand an early 1990s algorithm (Nayar et al., 2021) to better incorporate multiple objectives in the early 2010s. (Samantha Bastiani et al., 2015). Unique to the genetic algorithm approach is the production of multiple portfolios as a non-comparable Pareto front - other algorithms discussed above focus on giving a single portfolio as a solution. This makes these approaches universal and more efficient if the objective is to create portfolios for any kind of investor. The multi-objective focused algorithms SPEA-2, NSGA-II, and MOEA/D described above have all been applied to this problem. The first papers for all of these applications use two objectives - some measure of risk and some measure of return, later to be expanded. (Milhomem and Dantas, 1 01)

3 PROPOSED METHODOLOGY

3.1 Problem Definition

Given a set of investment vehicles and a budget, one must use risk and return to find an amount of each stock to purchase such that risk is minimal and return is maximal. Risk and return are not comparable. Thus, the solution is a set of options, each one best in its own way. Once the problem is formulated in terms of variables (the investment vehicles and how much to buy of each one, $v_0 ... v_n$), constraints (given in Equation 1

$$B \geq \sum v_i p_i \tag{1}$$

where B is the budget and p_i is the price of one instance of that investment vehicle), and objective functions (risk and return), the problem is ready to be modelled as a constraint optimization problem.

This is the standard definition of the problem: two objectives and one constraint, which is the minimum required for a problem to be a multi-objective constraint optimization problem. In order to capture the different types of risk, risk here is split into two non-comparable objectives: VaR and CVaR. Expanding on this, environmental, social and governance scores are each added as individual objectives, giving a total of six objectives in this formulation.

$$\sum v_i VaR_i \tag{2}$$

$$\sum v_i CVaR_i \tag{3}$$

Equation 2 and Equation 3 form the risk objectives, which should be minimized. Each investment vehicle also has a known environmental (e_i), social (s_i) and governance (g_i) score which should be maximized:

$$\sum v_i e_i \tag{4}$$

$$\sum v_i s_i \tag{5}$$

$$\sum v_i g_i \tag{6}$$

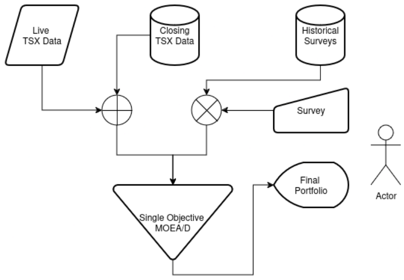

Figure 1: Flow chart for Architecture 1.

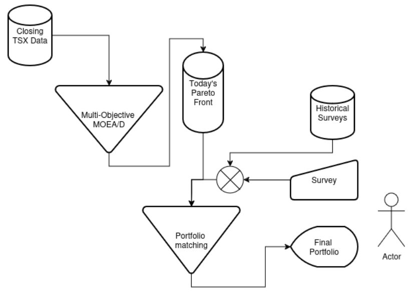

Figure 2: Flow chart for Architecture 2.

3.2 System Architecture

3.2.1 Architecture 1

Architecture 1 (depicted in Figure 1) runs the optimization each time an investor wishes to find a portfolio using their preferences as weights. This would be a weighted constraint optimization problem in which one optimal portfolio is produced for that investor - in other words, a single objective constraint optimization problem. The investor may sit at their desk, answer the survey, and then wait for the algorithm to finish computing their portfolio. Once the survey is answered, the following steps are executed.

1. Takes a snapshot of the TSX market data.
2. Computes the single optimized portfolio using a simplification of MOEA/D
   - Firstly, random solutions are generated using the snapshotted data.
   - The optimization is run per the configurations from section 4.2.1, however, this is a single objective optimization because the six objectives are able to be combined into one via the weights already provided by the user.
3. Presents this portfolio to the end user.

The biggest advantage of this is also connected to its biggest problem: if live pricing is used, then the portfolio may be out-of-date by the time it has finished computing. This is slightly mitigated by the speed of a single optimization when compared to a multi-optimization. If it uses the closing price data for the previous day, there are two scenarios depending on the time the investor chooses to create their portfolio:

1. Use the service when the markets are closed and act once the markets are open, or
2. Accept the risk that the prices are inaccurate by the day's market activity.

3.2.2 Architecture 2

The second architecture (Figure 2) runs a multi-objective optimization with all investors in mind. Once the individual investor sits down and requests a portfolio, that portfolio is computed as the highest scoring member of the Pareto front with respect to that investor's weighting. This would have linear growth based on the number of portfolios in the Pareto front, making it extremely quick. A day in the life of this architecture would look like this:

1. Any time after the market closes for the day, a snapshot of TSX data is taken
2. Using MOEA/D, a Pareto front of Pareto optimal portfolios is computed and saved
   - Firstly, random solutions are generated using the snapshotted data.
   - The optimization is run per the configurations from section 4.2.1. This is a six objective optimization, but it is only run one time per day.
   - The final Pareto front produced by MOEA/D is saved, the former day's discarded.
3. Every time an investor completes the survey, their results are converted to weights (per section 3.3.4)
4. Take each portfolio from today's Pareto front, multiply each objective value by its corresponding weight from the investor, and pick the maximal.

The biggest speed constraint for this method is, of course, that the optimization must run for a shorter amount of time than the time between the market closing and the market's next-day opening. Thus, the accuracy of the Pareto front solutions could be greater than the accuracy of the significantly more speed-conscious Architecture 1 because of the realignment of priorities when considering the parameters.

3.3 Survey

The survey in its entirety can be viewed in Appendix 5. It is divided into four parts: risk profiling, short, medium, and long.

3.3.1 Risk Profiling

Several publicly available risk assessment questionnaires were consulted before writing the risk section of the survey. These are commonplace and often follow similar patterns: many multiple choice questions where the answers ascend from low to high risk tolerance (an example one from a well known investment firm can be found in (Morningstar,)). Intentionally, there are a mix of even number (four) answer questions and odd number (five) answer questions. The even number of answers forces an indecisive investor to choose which side they lean on even if they want to remain relatively balanced.

3.3.2 ESG Survey

The short survey consists of one question: in essence, does the investor care about ESG metrics. If the investor does not, the survey endeavours to respect your time. More information is provided if the participant requires it to make this choice. If the participant is willing to give their preferences, the medium survey is presented.

The medium survey is intended to give a brief overview of the participants' preferences. This consists of an open text box that is currently unused in analysis (see Future Work for further information on why this is included), a ranking question, and three questions each of the following form: environmentally (or socially or governance) conscious companies should represent what percent of the portfolio. Then another checkpoint is presented, and if the participant wishes to continue, they can do the long survey.

The long survey is based on the Sustainable Development Goals of the United Nations. There are three ranking questions, each referencing one of environment, social, and governance, and finally another open text box that is unused for the same reason as expressed in the Future Work section. In the European Union (EU), by regulation, ESG must be incorporated into the investment advisory process. Thus, the development of surveys that conform to MiFID II sets the benchmark for surveys on the topic. A representative sample questionnaire that conforms to MiFID II that was consulted for the development of the survey presented here can be found here (Initiative, 2022).

3.3.3 Converting a Survey Result into Weights

Each survey question carries the same weight in the final weighting. The risk scoring is as follows for each question in the Risk Survey, starting from a score of 0:

- Each answer of *A* subtracts 2 from the overall score
- Each answer of *B* subtracts 1 from the overall score
- Each answer of *C* on a question with 5 possible answers adds 0. If there are 4 possible answers, a *C* adds 1
- Each answer of *D* on a question with 5 possible answers adds 1. If there are 4 possible answers, a *D* adds 2
- Each answer of *E* adds 2 to the overall score.

As is evident from the multiple choice answers and the scoring matrix, a higher score indicates a higher risk tolerance and vice versa.

For any final answer *C* to the short survey, the investor's environmental, social, and governance weights are all taken as 0, since the investor has indicated that they do not have any ESG preferences.

The scoring for each environmental, social, or governance questions from the medium and long surveys that take the form of a ranking is derived using the number of criteria that are ranked over the total number of criteria. For example, if an investor ranks $8/16$ as the environmentally focused question (the first question of the long survey), that will add 0.5 to their environmental score. The medium survey ranking question (the second question in the medium survey) contributes to all three scores in order to account for an investor that has a general preference in ESG but no distinguishing difference between environmental, social, or governance. The three medium survey questions asking for numbers are directly added as percentages to the score given that they are provided as valid percentages (between 0 and 100, then divided by 100).

3.3.4 Converting a Collection of Survey Results into Weights

Because no two survey results are alike and because defining whether a certain score in a certain metric is high or low begs the question "compared to what?", it is necessary to define how a survey score relates to a weight. Given a collection of survey results, each weight (corresponding to an objective) is normalized to its peers using a simple formula.

$$w_{oi} = \frac{s_{oi} - min_o}{max_o - min_o} \tag{7}$$

where s_{oi} is the ith investor's score for objective o, max_o and min_o are the global maximum (minimum) of the scores for objective o. Note that the risk weight is evenly split between the VaR and CVaR objectives. In addition, as there is no return objective deduced from the survey, the return takes up the rest of the weighting.

4 EXPERIMENTATION

4.1 Inputs and Configurations

4.1.1 Price

For consistency and comparability, the closing price for the previous day is used. One of the advantages of doing the optimization as it is done in Architecture 1 is that live pricing data or snapshots taken close to the time of optimization (and therefore close to the time of investing). While this is preferable for this architecture, it is not a possibility for Architecture 2 and so in order to preserve the comparability of the two architectures, the same pricing data are used for both. This can be considered a snapshot taken closest to the optimization when the optimization is run between a market's close and open.

In addition, all investment vehicles with a price of less that \$1.00 per share have been excluded from the data. This is to mitigate the problem-specific over weighting of penny stocks, which are known to have extremely high volatility. The whole list of TSX listings as well as the data for the included stocks are available in Appendix 5

4.1.2 Risk

VaR is calculated using the parametric method (also known as the variance-covariance method) of computing daily VaR. Although this carries the normal distribution assumption, it is extremely computationally efficient and is used widely when historical data is either difficult to come by or is not available at all (for example, for newly listed companies). Here, it is configured to use confidence 95%. VaR is also used in the computation of CVaR, so this normal distribution assumption is also present there.

CVaR (also called expected shortfall) is calculated according to the standardized formula

$$CVaR = \frac{1}{1-c}\int_{-1}^{VaR} x \tag{8}$$

where c refers to the confidence mentioned above (again, 95%) and x is the normally distributed returns over the period.

4.1.3 Return and CAPM

The expected return is calculated according to the capital asset pricing model (CAPM). The standard formula for CAPM is

$$CAPM = R_f + \beta_i(ER_m - R_f) \tag{9}$$

where R_f is the risk-free rate of the market, β_i is the beta value of the stock, and ER_m is the expected return of the market as a whole. The expression $ER_m - R_f$ can be considered the amount that one potentially gains given that they invest in the market as opposed to a risk-free asset and is often called the market risk premium.

For a risk-free asset, the Bank of Canada CORRA (Canadian Overnight Repo Rate Average) rate is used. This is Canada's free rate and is sourced from the Bank of Canada (Bank of Canada,). For the expected return of the market, the S&P TSX index is used. The important thing to choose this rate is that it mirrors the benchmark used to compare the performance of an investment vehicle with the market as a whole. S&P TSX is a broad-reaching index also used elsewhere in the literature when it comes to Canadian markets.

4.1.4 Benchmark

S&P TSX is a commonly used benchmark for the Canadian stock market, so to preserve future comparability, it is used here. It is also useful to use the same benchmark as is used for the calculation of the expected return (CAPM), as this allows for a more accurate comparison between the benchmark and the performance of a portfolio whose returns are already calculated compared to that benchmark.

4.1.5 Morningstar and ESG Data

The TSX and ESG data used in this implementation are sourced from Yahoo Finance. Yahoo Finance aggregates data from several sources: sustainability data is provided by Sustainalytics and Morningstar and corporate governance scoring is provided by Institutional Shareholder Services. Morningstar also provides historical chart data and financial statements. Company profile data comes from S&P Global Market Intelligence. The TSX pricing data are sourced in real time from ICE Data Services. (Finance,)

4.2 Algorithm Configurations

4.2.1 MOEA/D Configuration

For the experiment, configurations that allow this recommender system to be compared to other implementations of MOEA/D. The configurations given in Table 1 are also consistent between the two architectures.

4.2.2 Problem Specific Mutation Mechanism

An undersold influence on the performance of the solutions is the mutation operation in the genetic algorithm. There is a problem non-specific method of

Table 1: MOEA/D Configurations.

No. runs	10
No. individuals	500
No. generations	10
No. nearest neighbours	50
Percentage of genes mutating	20%
Budget	\$100,000

combining two parents into a child: simply swap a random collection of half the genes from each parent. This is common in the literature; however, when this was attempted, often the swap would be reversed because the combination violated the budgetary constraint. Thus, the swap has been swapped for this problem specific mutation. It takes half the value of each of the investment vehicles from each parent (and therefore half the used budget) rounded down. For example, given a mother that has options 30 of investment vehicle A, 20 of B, and 10 of C and a father with 10 of A, 21 of B, and 10 of D, the child will have $A = \lfloor \frac{30+10}{2} \rfloor = 20$, $B = \lfloor \frac{20+21}{2} \rfloor = 20$, $C = \lfloor \frac{10+0}{2} \rfloor = 5$, and $D = \lfloor \frac{0+10}{2} \rfloor = 5$.

Notice that the sum of units from before the mutation is less than the exact average number of units between both parents. Because the child will tend to have a lower budget utilization than its parents (the constraint is kept by the fact that the floor of the average of two numbers is less than or equal to the exact average, so this difference would in effect turn to unused "cash"), there is then a refill phase. Each option is checked in a random order to see whether the additional budget space can be used: if so, one additional unit is added. In this way, the mutation tends towards higher budget utilization over time.

4.3 Results

4.3.1 Example Investors - Sam, Jars, and Alice

Sam is very socially aware and has deep rooted preferences. They want their portfolio to reflect their values as much as it can. Sam is also gearing up for retirement - given that they don't have a long time horizon for this investment, preserving the initial capital is high on their priorities as well.

Sam took the survey and their weights are given in Table 2.

Table 2: Sam's weights.

Risk	0.1
Environmental	0.3
Social	0.3
Governance	0.2

Jars is just starting their career - they have finally paid back their student debt and now want their money to work as hard as they are. Jars does not particularly trust the ESG rating system and so would rather take their chances with casting a wider net. Jars took the survey and their weights are given in Table 3.

Table 3: Jars' weights.

Risk	0.9
Environmental	0.0
Social	0.0
Governance	0.0

Alice has been acutely aware of all things Corporate Governance for as long as they can remember. They feel their best impact can stay in their expertise. Alice has a little money for investing, but is also planning a down payment on a house so would like to keep a balanced portfolio in order to be able to jump on any housing opportunity they see. Alice took the survey and their weights are given in Table 4.

Table 4: Alice's weights.

Risk	0.3
Environmental	0.0
Social	0.0
Governance	0.6

4.3.2 Architecture 1 Results

The three portfolio examples of investors are given in the Appendix 5.

Given Architecture 1 as defined, these would be their results. The results are the best taken from the 10 runs. Sam's portfolio metrics are in Table 5, Jars' portfolio metrics are in Table 6, and Alice's portfolio metrics are in Table 7

Table 5: Sam's Architecture 1 portfolio.

	Portfolio	Benchmark (S&P TSX)
Return	0.8750	0.0112
VaR	0.9901	0.9999
CVaR	0.9894	1.000
Environmental	0.8681	N/A
Social	1.6878	N/A
Governance	2.9182	N/A

4.3.3 Architecture 2 Results

Architecture 2 gives a Pareto front of potential portfolios. All of these portfolios are available in Appendix 5. The metrics for the best of the 10 runs are presented here. Sam's portfolio metrics are in Table 8, Jars' portfolio metrics are in Table 9, and Alice's portfolio metrics are in Table 10.

Table 6: Jars' Architecture 1 portfolio.

	Portfolio	Benchmark (S&P TSX)
Return	0.7760	0.0112
VaR	0.7868	0.9999
CVaR	0.7818	1.000
Environmental	0.0000	N/A
Social	0.0000	N/A
Governance	0.0000	N/A

Table 7: Alice's Architecture 1 portfolio.

	Portfolio	Benchmark (S&P TSX)
Return	0.6980	0.0112
VaR	0.9566	0.9999
CVaR	0.9531	1.000
Environmental	0.0000	N/A
Social	0.0000	N/A
Governance	4.9238	N/A

Table 8: Sam's Architecture 2 portfolio.

	Portfolio	Benchmark (S&P TSX)
Return	0.4679	0.0112
VaR	0.9934	0.9999
CVaR	0.9948	1.000
Environmental	0.1350	N/A
Social	0.4011	N/A
Governance	0.1546	N/A

4.3.4 General Portfolio Observations

There appears to be a strong correlation between the three ESG criteria. This provides a reason for these to be combined into one weight for ESG, which is frequently assumed but not explored in the literature.

VaR and CVaR are also strongly correlated across the board. This is unsurprising as CVaR is based on VaR and while the values are not precisely equal, a company that is risky is very likely to be risky based on both metrics.

As assumed in the mean variance model, higher risk appears to be correlated to higher returns. This can be attributed in part to the metric used for the return. CAPM takes risk into account when it attempts to predict the return using the beta of the investment vehicle, so risk can be considered to be doubly counted here. However, the risk measure used is neither VaR nor CVaR, it is the value β, so it is possible that the outliers perform differently. Both risk metrics and return, in these graphs, appear to be inversely related to ESG criteria. That is, with higher ESG ratings come lower risk and lower returns. As mentioned in the review of the literature, the literature differs substantially on this topic.

With all that being said, this does corroborate some conclusions found in the literature. Although

Table 9: Jars' Architecture 2 portfolio.

	Portfolio	Benchmark (S&P TSX)
Return	0.4679	0.0112
VaR	0.9934	0.9999
CVaR	0.9948	1.000
Environmental	0.1350	N/A
Social	0.2310	N/A
Governance	0.1546	N/A

Table 10: Alice's Architecture 2 portfolio.

	Portfolio	Benchmark (S&P TSX)
Return	0.4679	0.0112
VaR	0.9934	0.9999
CVaR	0.9948	1.000
Environmental	0.1350	N/A
Social	0.2310	N/A
Governance	0.1546	N/A

the Architecture 2 portfolios for each of our investors are similar, they do fulfill the needs of that investor remarkably well. This aligns with the result of (Utz et al., 4 16), (De Spiegeleer et al., 0 16) and (Naffa and Fain, 1 01) who concluded that there is no substantial difference in performance whether given ESG constraints/objectives or not. The best portfolio from Architecture 2 was better in most regards than any portfolio from Architecture 1. All Architecture 1 portfolios were more significantly tailored to each of the individuals' preferences.

5 CONCLUSION

Ultimately, if the goal of the system is to tailor and adhere to the preferences of the individuals, the better performer is clearly Architecture 1; if the goal is to produce quality portfolios, the better architecture is Architecture 2. For Architecture 2, a limited number of high performing portfolios was generated: as a result, more investors would potentially be recommended to the same few portfolios as in Architecture 1. It was demonstrated that a system leveraging multi-objective evolutionary computation, specifically MOEA/D, was able to produce highly performant portfolios tailored to an individual given a custom survey. This survey, written using the greater context of other risk and ESG relevant surveys, was conducted and used to construct a weighting to normalize a given investor's own survey responses and allow a single portfolio from the collection of the best portfolios to be matched to that investor.

Two potential architectures were considered. For Architecture 1, the optimization is run for each investor that takes the survey. This allows real-time pricing as well as a more dynamic resource allocation (because for n investors, when n is small, the expensive optimization task is n times, which is also small). However, this has the cost of the inverse: at high volumes, the expensive optimization task is run many times (for the same reason). For Architecture 2, a multi-objective optimization is run less frequently (thereby taking a snapshot of pricing data and losing the advantage of live pricing data) and the investor is given a portfolio from the Pareto front. This subset are all the best dominant portfolios and the matching process is extremely fast. The result of the experiment was that both architectures produced highly performant portfolios that performed comparablely. Each portfolio produced by both architectures significantly outperformed the benchmark portfolio (S&P TSX) significantly. For Architecture 2, the number of portfolios generated was not high for large numbers of generations: as a result, more investors would potentially be recommended to the same few portfolios compared to Architecture 1.

The following directions could be pursued in this research:

- exploring more under-served markets, the point at which too many objectives becomes problematic for the portfolios generated by adding more objectives such as transaction cost, liquidity, and media sentiment, and how different ESG vendors/metrics change the portfolios generated, and how different algorithms change the portfolios generated,
- using the survey as an input to a different problem, such as using the open text answers to perform sentiment analysis or any other NLP application,
- and realizing the implementation as a robo-advisor and address the following questions:
  - how well does the robo-advisor perform when compared to a human advisor?
  - do people trust the robo-advisor as much as the human advisor?
  - does the robo-advisor properly take the needs of the individual into account (ie does the survey satisfy this need fully)?

REFERENCES

Bank of Canada. Canadian overnight repo rate average.

Darmstadt, T. F. and Approximity, A. R. (2003). Risk and performance optimization for portfolios of bonds and stocks. In *Proceedings of the International AFIR Colloquium*.

De Spiegeleer, J., Höcht, S., Jakubowski, D., Reyners, S., and Schoutens, W. (2020-10-16). ESG: A new dimension in portfolio allocation.

Finance, Y. Exchanges and data providers on yahoo finance | finance for web help - SLN2310.

Initiative, T. . I. (2022). Draft questionnaire & guidance for client sustainability preferences - 2dii.

Korani, W. and Mouhoub, M. (2021). Review on nature-inspired algorithms. In *Operations research forum*, volume 2, page 36. Springer.

Korani, W. and Mouhoub, M. (2022). Discrete mother tree optimization and swarm intelligence for constraint satisfaction problems. In *Proceedings of the 14th International Conference on Agents and Artificial Intelligence, ICAART 2022, Volume 3, Online Streaming, February 3-5, 2022*, pages 234–241. SCITEPRESS.

Lynn, N. and Suganthan, P. N. (2015-10-01). Heterogeneous comprehensive learning particle swarm optimization with enhanced exploration and exploitation. *Swarm and Evolutionary Computation*, 24:11–24.

Milhomem, D. and Dantas, M. (2022-01-01). Analysis of new approaches used in portfolio optimization: A systematic literature review. In *Evolutionary and Memetic Computing for Project Portfolio Selection and Scheduling*, pages 125–157. Springer.

Morningstar, I. Risk profiling and risk scoring tools.

Naffa, H. and Fain, M. (2022-01-01). A factor approach to the performance of ESG leaders and laggards. *Finance Research Letters*, 44:102073.

Nayar, N., Gautam, S., Singh, P., and Mehta, G. (2021). Ant colony optimization: A review of literature and application in feature selection. In Smys, S., Balas, V. E., Kamel, K. A., and Lafata, P., editors, *Inventive Computation and Information Technologies*, pages 285–297. Springer Nature.

Sadreddin, A., Mouhoub, M., and Sadaoui, S. (2022). Portfolio selection for SAT instances. In *IEEE International Conference on Systems, Man, and Cybernetics, SMC 2022, Prague, Czech Republic, October 9-12, 2022*, pages 2962–2967. IEEE.

Samantha Bastiani, S., Cruz-Reyes, L., Fernandez, E., Gómez, C., and Rivera, G. (2015). An ant colony algorithm for solving the selection portfolio problem, using a quality-assessment model for portfolios of projects expressed by a priority ranking. In Melin, P., Castillo, O., and Kacprzyk, J., editors, *Design of Intelligent Systems Based on Fuzzy Logic, Neural Networks and Nature-Inspired Optimization*, pages 357–373. Springer International Publishing.

Utz, S., Wimmer, M., Hirschberger, M., and Steuer, R. E. (2014-04-16). Tri-criterion inverse portfolio optimization with application to socially responsible mutual funds. *European Journal of Operational Research*, 234(2):491–498.

APPENDIX

Recommender System Code

The following repositories were written in the making of this recommender system, hosted at https://rileyherman.ca/survey:

- Main implementation of algorithms
- Import data from yahoo finance and massage into input
- Match portfolios from outputs to survey results (Architecture 2)
- Website backend, including po, pomatch, and poimport
- Website frontend, including survey and portfolio presentation

Survey

Available as a public link here.

Portfolios

Please note that these portfolios are not intended as investment advice and have absolutely no guarantee of performance: they are provided as an addendum to the recommender system. The author does not hold any responsibility should an individual see these results and decide to invest.

- Graphs, including generational progressions as well as pairwise comparisons between objectives
- All Architecture 2 portfolios across all runs
- All architecture 1 portfolios across all runs
- The beta values for each of the portfolios above
- The snapshot of all the data from TSX used to generate the portfolios above
- A snapshot of companies listed on the TSX. Each company listed here that does not exist in the included data set either had insufficient data to be included or was priced at less than $1.00 per share
- The comparison between the portfolios generated and the benchmark