A Learnheuristic Approach to A Constrained Multi-Objective Portfolio Optimisation Problem

SONIA BULLAH, Computer Science and Applied Mathematics, University of the Witwatersrand, South Africa TERENCE L. VAN ZYL, Institute for Intelligent Systems, University of Johannesburg, South Africa

Multi-objective portfolio optimisation is a critical problem researched across various fields of study as it achieves the objective of maximising the expected return while minimising the risk of a given portfolio at the same time. However, many studies fail to include realistic constraints in the model, which limits practical trading strategies. This study introduces realistic constraints, such as transaction and holding costs, into an optimisation model. Due to the non-convex nature of this problem, metaheuristic algorithms, such as NSGA-II, R-NSGA-II, NSGA-III and U-NSGA-III, will play a vital role in solving the problem. Furthermore, a learnheuristic approach is taken as surrogate models enhance the metaheuristics employed. These algorithms are then compared to the baseline metaheuristic algorithms, which solve a constrained, multi-objective optimisation problem without using learnheuristics. The results of this study show that, despite taking significantly longer to run to completion, the learnheuristic algorithms outperform the baseline algorithms in terms of hypervolume and rate of convergence. Furthermore, the backtesting results indicate that utilising learnheuristics to generate weights for asset allocation leads to a lower risk percentage, higher expected return and higher Sharpe ratio than backtesting without using learnheuristics. This leads us to conclude that using learnheuristics to solve a constrained, multi-objective portfolio optimisation problem produces superior and preferable results than solving the problem without using learnheuristics to solve a constrained, multi-objective portfolio optimisation problem produces superior and preferable results than solving the problem without using learnheuristics.

Additional Key Words and Phrases: multi-objective portfolio optimisation, learnheuristics, metaheuristics, genetic algorithms, surrogate-assisted algorithms, machine learning, deep learning, constraints

ACM Reference Format:

Sonia Bullah and Terence L. van Zyl. 2023. A Learnheuristic Approach to A Constrained Multi-Objective Portfolio Optimisation Problem. 1, 1 (April 2023), 12 pages. https://doi.org/10.1145/nnnnnnnnnnnn

1 INTRODUCTION

This study explores a multi-objective portfolio optimisation problem by including numerous factors that produce above-average returns while attempting to minimise the investor's overall risk [15].

This is conducted by introducing realistic constraints into the optimisation model alongside machine learning, deep learning [10] and metaheuristics to solve non-convex problems. The metaheuristic algorithms explored are Evolutionary Algorithms (EA), namely NSGA-II, R-NSGA-II, NSGA-III and U-NSGA-III. In contrast, the machine learning models used to enhance these metaheuristic algorithms will be that of the surrogate-assisted variants of each metaheuristic algorithm used [16, 20].

In the study by Ma *et al.* [12], the researchers compare machine learning and deep learning models with portfolio optimisation models to determine the best combination of investment strategies. Ultimately, it is concluded that the Random Forest Model + Mean-Variance with Forecasting (MVF) Model produces the most efficient and

Authors' addresses: Sonia Bullah, 2107762@students.wits.ac.za, Computer Science and Applied Mathematics, University of the Witwatersrand, Johannesburg, South Africa; Terence L. van Zyl, Institute for Intelligent Systems, University of Johannesburg, Johannesburg, South Africa, tvanzyl@gmail.com.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

^{© 2023} Association for Computing Machinery.

XXXX-XXXX/2023/4-ART \$15.00

profitable results for practical investments. However, the only issue is that a high turnover results in a huge challenge for these models to overcome. A high turnover implies that the transaction costs or transaction fees (both terms can be used interchangeably) will be quite high. Since these costs are very much present in realistic trading scenarios, they should not be overlooked when analysing the portfolio optimisation problem. Hence, this study will introduce them and other realistic factors as constraints in the portfolio optimisation problem. Adding these additional constraints can be done by employing a hybrid approach, which includes merging metaheuristics with machine learning and deep learning techniques. This approach can be called 'learnheuristics' [7].

Overall, this study will focus on how realistic constraints, such as transaction costs and holding costs, can be introduced into the portfolio optimisation problem in such a way as to account for their presence in realistic trading scenarios as well as how learnheuristics can be used to assist in solving a constrained, multi-objective and non-convex portfolio optimisation problem. By introducing a unique, hybrid approach, this study will solve the current issue of unrealistic trading scenarios in portfolio optimisation problems while producing superior results compared to those produced by alternative methods and benefiting the investor simultaneously.

2 BACKGROUND

2.1 Portfolio Optimisation

Portfolio optimisation forms the basis of this study and is a fundamental problem faced by many investors [17]. The portfolio optimisation process comprises selecting the optimal portfolio from all possible portfolios to maximise investor wealth while minimising risk. A popular way of doing this includes using the Sharpe ratio, which can be directly optimised [23].

One of the major contributions to the portfolio selection problem is Markowitz's theory of portfolio selection [13]. Markowitz derives a model that assumes that investors aim to maximise their returns for a particular level of risk, or the converse. Many other studies have extended Markowitz's theory to include constraints. Kizys *et al.* [9] acknowledge that the complexity of this model increases significantly when further constraints are added. However, metaheuristics can be used to overcome such obstacles.

Doering *et al.* [7] extend Streichert *et al.* [21] by formulating an unconstrained portfolio optimisation problem as per Markowitz's mean-variance theory in such a way that minimises risk as follows:

$$Min\sum_{i=1}^{N}\sum_{j=1}^{N}\omega_{i}\omega_{j}\sigma_{ij}$$

and maximises profit or returns in the following manner:

$$Max\sum_{i=1}^N \omega_i \mu_i$$

subject to:

$$\sum_{i=1}^{N} \omega_i = 1, \quad \text{where } 0 \le \omega_i \le 1, \forall i = 1, 2, ...N,$$

and N represents the available assets, μ_i represents the expected return of asset *i*, σ_{ij} represents the covariance between two assets *i* and *j*, and ω_i represents the decision variables of the portfolio. This study will adjust such models to include realistic constraints such as transaction costs.

Furthermore, portfolio optimisation can be classified as a multi-objective optimisation problem [22] with the target of finding solutions that can illustrate the trade-off present for different objectives [17]. A multi-objective portfolio optimisation problem requires two objectives to be met [21].

[,] Vol. 1, No. 1, Article . Publication date: April 2023.

Machine learning and deep learning models can be combined with traditional optimisation models [12] to find the best combination for practical investments. These ideas will be explored in the sections below.

2.2 Machine Learning Models and Deep Learning Models for Multi-Objective Portfolio Optimisation

Zhang *et al.* [23] focuses primarily on using deep learning models to evaluate portfolio optimisation problems by first providing an objective function in which the Sharpe ratio is defined and maximised using gradient ascent.

Ma *et al.* [12] make use of popular machine learning models, such as Support Vector Regression (SVR) and Random Forest (RF), as well as deep learning models, such as Deep Multilayer Perceptron (DMLP), LSTM Neural Network and Convolutional Neural Network (CNN). Furthermore, they pair each model with a traditional portfolio optimisation model, namely the Mean-Variance and Omega Models.

In addition to using machine learning and deep learning models, many studies also introduce hybrid algorithms. This can be seen in Kizys *et al.* [9], in which a new algorithm, *ARPO*, is used to consider inevitable, real-world constraints that many researchers refrain from analysing. Similarly, van Zyl *et al.* [22] implement *ParDen* to include machine learning models that are either generative or discriminative and use them as a surrogate to reduce costs that arise from backtesting procedures.

Optimisation problems depend on time-intensive functions due to the problem's computationally expensive nature [14]. Implementing such functions influences the computational overhead of the algorithm, which is used to discover new solutions every time the function is called. Therefore, a *surrogate* model can be incorporated to accelerate the convergence and approximate the time-intensive functions of the models and algorithms currently being used to solve these optimisation problems.

Surrogate-assisted optimisation satisfies the objective of producing efficient approximate solution evaluations (ASEs), which are then used to minimise expensive solution evaluations (ESEs), all while providing a minimal approximation error. Overall, the convergence of the optimisation algorithm being utilised will improve [14].

2.3 Transaction Costs And Holding Costs

Transaction costs arise from turnover; the higher the turnover, the more transaction costs will increase. This suggests that the performance of a model will differ with and without factoring in the relevant transaction costs [12]. Ma *et al.* [12] extends [7, 9, 11, 22] and [1], in which the concept of additional costs is discussed briefly and where transaction costs are recognised as a major constraint in the realistic portfolio optimisation problem.

Likewise, holding costs are also important to consider in practical investment and trading scenarios. The simplest holding cost model accounts for charges incurred for borrowing assets when taking a short position [5]. However, these models can become significantly more complex under certain conditions. Similar to transaction costs, holding costs generally consist of non-negative values. However, negative values are not unlikely.

Trading costs and holding costs can be combined to formulate a multi-objective optimisation problem with the inclusion of a leverage constraint as done by van Zyl *et al.* [22]. This is shown by:

Maximise:
$$w'\mu - \frac{\gamma}{2}w'\sum w - \gamma_t \phi^{trade}(w - w_0) - \gamma_h \phi^{hold}(w)$$

Subject to: $||w||_1 \le L^{max}$

where w_0 denotes the initial portfolio, $w - w_0$ represents the trades needed to secure w, ϕ^{trade} is the trading cost function with the trade-off parameter γ_t , $\phi^{hold}(w)$ is the trading cost function with the trade-off parameter γ_h and, finally, L^{max} denotes the maximum leverage.

It can be noted that a fully-invested portfolio, in which a long position is maintained, has a leverage of 1.

2.4 Introducing Portfolio Constraints

Most portfolio optimisation problems consider some form of risk management, and many research papers fall short of considering stochastic optimisation problems, which include using constraints in optimisation models [7].

Abate *et al.* [1] consider portfolio constraints, as considering mean-variance optimisation alone leads to unfair asset allocations. They also implement a list of constrained strategies. These strategies include classical constraints, flexible portfolio constraints, norm-based constraints, variance-based constraints, tracking error volatility constraints and beta constraints. They note that tighter constraints can produce diversified portfolios and that choosing the proper constraints plays a vital role in the evaluation process. Investors should typically choose constraints that align with their target or goal.

Kizys *et al.* [9] extend the work of Doering *et al.* [7] and Abate *et al.* [1], in which both articles consider constraints that are unavoidable for investors to ignore in realistic investment scenarios.

The transaction costs play a significant role when choosing constraints. Constrained portfolio optimisation models are favourable to investors and can be seen as an effective alternative to classic investment approaches [7].

2.5 Metaheuristics for Solving Optimisation Problems

Metaheuristic algorithms can provide solutions to portfolio optimisation problems that are non-convex. In general, the single-period optimisation problem (SOP) can be efficiently solved given that the problem is convex [5]. However, once this is adjusted to become a multi-objective optimisation problem (MOP) with non-convex constraints, the problem becomes NP-hard and is more complicated to solve. Metaheuristics play a vital role in this case as they can be used to solve NP-hard optimisation problems and provide exceptional solutions [7], while considering the realistic, non-convex constraints introduced into the problem.

Doering *et al.* [7] classify metaheuristics into two groups: nature-inspired and non-nature metaheuristics. This study will primarily focus on a sub-class of nature-inspired metaheuristics, evolutionary algorithms. In particular, it will focus on the Genetic Algorithm (GA). However, this study focuses on multi-objective problems, so multi-objective metaheuristics must be incorporated. Therefore, the various forms of genetic algorithms, namely, NSGA-II, R-NSGA-III and U-NSGA-III [22], will be explored.

NSGA-II is a genetic algorithm based on non-domination and is popular for solving multi-objective optimisation problems [19]. First, the algorithm initialises the population and considers any constraints that may be specified. Next, the population goes through a non-dominated sort. Each individual is then given a rank or fitness value based on the front to which they belong. Furthermore, a crowding distance is assigned to each individual in the population, quantifying how close the relevant individual is to their neighbours. Next, a binary tournament selection process is used to select parents from the population, and offspring are produced from the selected population. The relevant population and offspring are again sorted using a non-dominated sort procedure, and only the best individuals are chosen.

The R-NSGA-II algorithm is an extension of the NSGA-II algorithm with an improved survival selection process [22]. In this algorithm, a smaller Euclidean distance between parents and the reference point is favoured or preferred [8]. The algorithm begins the same way as NSGA-II, resulting in a single, merged population. The merged population is then segmented into various non-dominated levels using non-dominated sorting. Then each front of the population is clustered to ensure that the solutions are diverse. The whole population size is then diminished to match that of the parent population ¹ [8]. The difference between this algorithm and NSGA-II is that solutions are based on the rank or fitness value while segmenting the front [22].

¹In the case of not all members being selected, the ones with the smaller Euclidean distances between the member itself and the reference point will be conserved.

[,] Vol. 1, No. 1, Article . Publication date: April 2023.

Unlike NSGA-II, which uses selection based on rank or fitness and crowding distance, NSGA-III uses reference directions to ensure diversity amongst the solutions [18]. The NSGA-III algorithm begins with a random population and disperse reference points on a hyper-plane and aims to find only one population member per reference point. The population is segregated into various non-dominated levels, and then non-dominated sorting is performed, just as the NSGA-II algorithm does. After this, an offspring population is produced from which a combined population is formed. Next, points are iteratively selected until a point is reached at which all solutions from a complete front cannot be selected. The final front will not be able to be fully selected. However, only a few solutions from this front need to be selected in this case. Then, each population member and member of the final front is linked to a reference point obtained from the minimum perpendicular distance of each member. Next, a niching strategy is implemented and followed to obtain members of the final front corresponding to the population's marginalised reference points. Population members associated with these points are preferred. The algorithm attempts to find one population member corresponding to each reference point close to the Pareto-optimal front [18].

The population will be segmented into non-dominated fronts using a non-dominated sorting process for a multi-objective optimisation problem. U-NSGA-III uses a niched tournament selection process as this has a superior performance compared to random selection [22]. This places weight on non-dominated solutions over dominated solutions and solutions nearer to the reference directions than those farther away. Thereafter, the algorithm follows the same process as that of the NSGA-II algorithm [18].

2.6 The Learnheuristic Approach

'*Learnheuristics*' is a hybrid approach that combines metaheuristics with machine learning or statistical learning models [7]. This study will use a learnheuristic approach to pair each of the four multi-objective genetic algorithms with a surrogate model. The surrogate-assisted algorithms will then be used to evaluate the portfolio optimisation problem.

Calvet *et al.* [6] explain that the hybridisation of machine learning and metaheuristics can be split into two classes, namely one in which machine learning is used to enhance metaheuristics and the other class being one where metaheuristics are used to improve the machine learning models or techniques used.

Using machine learning to enhance metaheuristics can be split into specifically-located and global-level hybridisations. Specifically-located hybridisations take on approaches such as fine-tuning parameters, quick initialisation, inexpensive evaluation techniques, population management, introducing the knowledge in operators and, lastly, using machine learning techniques as local searches. On the other hand, global hybridisations focus on reducing the search space, algorithm selection, using hyperheuristics (which are search methods used for choosing heuristics that will provide solutions to computational search problems), cooperative strategies and, finally, using new types of metaheuristics.

Conversely, using metaheuristics to improve machine learning techniques spans topics such as classification, regression, clustering and rule mining.

Taking on a learnheuristic approach implies that either one or both of the classes mentioned above be implemented. Table 1 indicates that while the related work either acknowledges or implements methods that consider some form of realistic constraints, only the survey presented by Doering *et al.* [7] focuses on metaheuristics to solve multi-objective optimisation problems. Therefore, this study will combine techniques outlined in the related work and attempt to enhance metaheuristics by using machine learning techniques and producing improved results.

Author	Optimisation	Constraints	Algorithms	Advantages
Ma et al.	Single- & Multi- objective	With con- straints	Machine learning & Deep learning	Provides solutions to NP-hard optimisation problems; produces high-quality solutions.
Zhang <i>et al.</i>	Single- objective	With con- straints	Deep learning	Improved performance; method is rational and practical.
Doering <i>et al</i> .	Single- objective	With and without constraints	Metaheuristic	Yields a high turnover.

Table 1. A summary of the work related to this study.

3 METHODOLOGY

To test whether or not a learnheuristic approach to a multi-objective portfolio optimisation problem will produce higher returns while minimising the risk of an investment, we generate a hybrid learnheuristic algorithm and compare the results to those of the baseline algorithms. In the past, research surrounding multi-objective portfolio optimisation, such as the research conducted by Ma *et al.* [12], has been explored without realistic constraints, such as transaction costs. Therefore, while attempting to maximise returns and minimise risk concurrently, our study will factor in these constraints in a learnheuristic approach similar to the process outlined in the study by Calvet *et al.* [6].

The implementation of this study can be separated into five main interdependent sections. The sections will all be described in detail below and are summarised in Fig. 1.

3.1 Data Cleaning And Exploration

The S&P 500 data was pre-processed to ensure correct datatypes and eliminate redundant attributes. Duplicates, missing values and null values were also dropped. We used only the top ten stocks regarding expected returns as our portfolio to conduct the rest of the study.

3.2 Forecasting

To forecast the data, some form of noise needs to be added. This noise represents an error that will increase or decrease predictions. This makes it difficult to make accurate predictions about the data. We use a mean of 0 and a standard deviation of 0.1 to acquire random samples from a normal distribution. This, in turn, outputs an array containing the noise, which is now added to the data, resulting in forecasted data values. Lastly, any null values need to be dropped before the portfolio optimisation.

3.3 Genetic Algorithms

In this study, the four genetic algorithms are needed to satisfy two main objectives, namely, optimising the portfolio and then finding the optimal constraint value simultaneously and, secondly, producing the best set of weights that can be used to perform backtesting. All genetic algorithms are implemented using the *Pymoo* [3] library, and the portfolio optimisation method followed is dependent on *PyPortfolioOpt* [2].

First, an objective function is adjusted to perform portfolio optimisation on the data. This is done by obtaining the forecasted dataset's expected returns and sample covariance. The expected returns are computed using the Capital Asset Pricing Model

$$R_i = R_f + \beta_i (E(R_m) - R_f),$$

A Learnheuristic Approach to A Constrained Multi-Objective Portfolio Optimisation Problem • 7



Fig. 1. A summary of the methodology followed in this study.

where R_f represents the risk-free rate, β_i is the beta (or sensitivity) of asset *i* and $(E(R_m) - R_f)$ the market risk premium. In addition, the sample covariance is obtained using the Ledoit-Wolf covariance shrinkage estimator

$$\hat{\Sigma} = \alpha F + (1 - \alpha)S_{f}$$

where α represents the shrinkage constant, *F* the a structured estimator and *S* a sample covariance matrix.

This, in turn, allows us to find optimal portfolios and produce an efficient frontier once the results are correctly plotted. Furthermore, it should be noted that this method will solve an unconstrained portfolio. However, the focus of this study is on constrained portfolios. As such, we add a constraint to the optimisation problem.

Once an efficient frontier has been produced, we proceed by adding a constraint to the efficient frontier. Instead of adding an arbitrary constraint with a random value, we search for the best possible parameter. Once this has been found, the Sharpe ratio is maximised, and the overall portfolio performance is returned.

The genetic algorithms are executed by randomly generating an initial population and assigning a fitness score to each individual. Next, individuals go through the selection and reproduction stages before crossover and mutation occur, in which new genetic information is added to the child population. Finally, generational replacement occurs, and then the algorithm terminates based on a stopping condition.

Once each genetic algorithm returns the array of portfolio performance results obtained, we adjust the array by removing the Sharpe Ratio values from each sub-array as we do not want a three-dimensional plot. This leaves us with sub-arrays of size two, which contain the optimal portfolio's return and volatility (or risk). These results are then plotted.

This process is run ten times on a population of 12 over 30 generations. We compare results by plotting all four genetic algorithms on the same axis. To compare the success of each algorithm, we note the average time taken for each algorithm to run to completion and the average hypervolume produced by the algorithm. Firstly, we begin a timer each time an algorithm starts and end it once the algorithm completes its last generation. The algorithms' average time (in seconds) is returned at the end. After each algorithm runs, we use the hypervolume function from the *Pymoo* [3] library. A hypervolume is a performance indicator which uses a reference point to calculate the area influenced by the solution set concerning a specific reference point [3]. It should be noted that higher hypervolumes suggest better-performing algorithms. Once we have the relevant hypervolumes for each generation of each algorithm, we calculate the average hypervolume from all ten runs and compare the algorithms to each other to determine how well they perform over the total number of generations.

Within the objective function, we also add an objective to the efficient frontier, which aims to adjust the gamma value. This parameter for L2 regularisation can be increased when requiring further non-negligible weights. The genetic algorithms search for the best possible value of gamma, which, in turn, produces the optimal set of weights that are used to perform backtesting. This process is elaborated on in the next section. Like the first case, once the set of weights for each stock is found, the portfolio is optimised, and the maximum Sharpe ratio is found. The final weights of each stock of the optimised portfolio are then added to an array, which is utilised in the backtesting procedure.

3.4 Learnheuristics: Surrogate-Assisted Genetic Algorithms

Once our genetic algorithms run, we adjust each function to include a surrogate model. We call the GPSAF implementation from the *Pysamoo* [4] framework to do this. Furthermore, we manually fine-tune the *alpha, beta,* $n_max_doe, n_max_infills, n_offsprings$ and *seed* parameters to produce an optimal convergence rate for each algorithm. We continue adjusting these parameters until the surrogate-assisted genetic algorithms are 'learning' throughout the iterations. We eventually produce an improved solution compared to the non-surrogate-assisted genetic algorithms. Again, these learnheuristic algorithms are run ten times on a population size of 12 over 30 generations, with the average time taken and average hypervolume being recorded at the end. We compare results by plotting each genetic algorithm alongside the surrogate-assisted version of the algorithm on the same axis and take note of the difference in convergence rate, average hypervolume and the average time taken for the algorithms to run to completion.

3.5 Backtesting

We run simulations or backtests to help optimise our strategy. To do this, we call the *Backtest* function from the *portfolio-backtest*² library. The weights obtained by the genetic algorithm are substituted into the function and the target return, target variance, start period and end period. We first run a backtest using the function's built-in method of determining the weights for each asset and then use each of the four surrogate-assisted genetic algorithms to search for the optimal set of weights that can then be passed into the backtesting function. Lastly, we record the data, namely the annual volatility, expected annual return and Sharpe ratio outputted from each backtest and compare the results.

All algorithms and implementations were run on a Quad-Core Intel Core i5 running macOS 12.3. The complete code, datasets and results can be found on the GitHub repository.

²The documentation can be found on: https://pypi.org/project/portfolio-backtest/

[,] Vol. 1, No. 1, Article . Publication date: April 2023.

4 RESULTS AND DISCUSSION

In this section, we present the results of the four multi-objective genetic algorithms used to optimise the portfolio and then consider the limitations and future work.

4.1 Results

Based on the results presented in Fig. 2(a) below, it is evident that the NSGA-III algorithm runs to completion in the shortest amount of time, followed by the U-NSGA-III algorithm and then the R-NSGA-II algorithm, while the NSGA-II algorithm takes the longest time to run. However, given that the NSGA-II algorithm produces the highest hypervolume followed by the R-NSGA-II, U-NSGA-III, and, lastly, NSGA-III algorithms, it suggests that the best performing algorithm is NSGA-II.

Algorithm	Generations	Population	HV	Time (sec)
NSGA-II	30	12	0.7225	1.0115
R-NSGA-II	30	12	0.5790	0.7941
NSGA-III	30	12	0.4986	0.7557
U-NSGA-III	30	12	0.5126	0.7681



(a) The average hypervolume and average time taken (in seconds) for each algorithm to run over 30 generations with a population size of 12.

(b) A comparison of all four genetic algorithms with a population size of 12 over 30 generations.

Fig. 2. A comparison of the different genetic algorithms.

Based on Fig. 2(b), it can be seen that the NSGA-II algorithm is the first to converge and produces a smoother line plot. The remaining three algorithms contain more spikes in their respective line plots, however, the general outcome suggests that the hypervolume increases as the number of generations increases.

Next, we present the results of each of the four learnheuristic algorithms (or surrogate-assisted genetic algorithms). Fig. 3(a) indicates that the surrogate-assisted R-NSGA-III algorithm runs to completion in the shortest amount of time, followed by the U-NSGA-III algorithm and then the NSGA-III algorithm, while the NSGA-II algorithm takes the longest time to run. However, similar to the non-surrogate variant, the NSGA-II algorithm yields the highest hypervolume, thus making it the most efficient algorithm.

Fig. 4(a), Fig. 4(b), Fig. 4(c) and Fig. 4(d) below each indicate that the learnheuristic algorithms produce higher hypervolumes than the baseline algorithms. By comparing Fig. 2(a) and Fig. 3(a), we observe that while the learnheuristic algorithms take significantly longer to run to completion, they produce much higher hypervolumes and converge faster while also creating smoother plots. As Blank *et al.* [2] mentioned, the learnheuristic algorithms accelerate the convergence of the multi-objective metaheuristic algorithms used to solve the optimisation problem.

Lastly, we run simulations or backtests on our data and record the annual volatility, expected annual return and Sharpe ratio generated for a tangency portfolio of ten assets. Based on Fig. 3(b), we see that the tangency portfolio for each of the four learnheuristic algorithms outperforms the backtest run without a learnheuristic algorithm. The portfolios that run by utilising the surrogate-assisted genetic algorithms to find the optimal weights for asset allocation yield higher returns for a lower risk percentage and produce greater Sharpe ratios.

Algorithm	Generations	Population	HV	Time (sec)
SA_NSGA-II	30	12	0.7543	18.0166
SA_R-NSGA-II	30	12	0.7181	15.2820
SA_NSGA-III	30	12	0.6403	16.6078
SA_U-NSGA-III	30	12	0.7093	16.4171

(a) The average hypervolume and average time taken (in seconds) for each surrogate-assisted metaheuristic algorithm to run over 30 generations with a population size of 12.

Algorithm	Volatility (%)	Return (%)	SR	Time (sec)
No GA	22.4	26.6	1.10	-
SA_NSGA-II	18.2	55.6	2.94	40.7051
SA_R-NSGA-II	18.2	55.6	2.94	43.0664
SA_NSGA-III	18.2	55.6	2.94	29.5306
SA_U-NSGA-III	18.2	55.6	2.94	28.7683

(b) The annual volatility, expected annual return and Sharpe ratio (SR) of a tangency portfolio generated from backtests as well as the time taken (in seconds) for each algorithm to run.

Fig. 3. A comparison of the different surrogate-assisted metaheuristic algorithms.



(a) A comparison of the NSGA-II and surrogate-assisted NSGA-II algorithms.



(c) A comparison of the NSGA-III and surrogate-assisted NSGA-III algorithms.



(b) A comparison of the R-NSGA-II and surrogate-assisted R-NSGA-II algorithms.



(d) A comparison of the U-NSGA-III and surrogate-assisted U-NSGA-III algorithms.

Fig. 4. A comparison of the different algorithms.

In these cases, the expected annual returns and Sharpe ratios are more than double compared to the backtest run without using a genetic algorithm to find the optimal asset weights. In the study by Zhang *et al.* [23], the Sharpe ratio is maximised within the objective function before deep learning models evaluate the portfolio optimisation problem. Similarly, in this study, we solve the optimisation problem by maximising the Sharpe ratio in the objective function, which is then called by each metaheuristic algorithm.

Although each of the four learnheuristic algorithms produces optimal results, Fig. 3(b) indicates that the surrogate-assisted U-NSGA-III algorithm runs to completion in the shortest time, followed by NSGA-III, NSGA-II and, finally, R-NSGA-II. Therefore, the surrogate-assisted U-NSGA-III algorithm is the most efficient learnheuristic for backtesting purposes. Overall, it is evident that the learnheuristic algorithms produce superior results compared to solving the portfolio optimisation problem without a learnheuristic approach.

4.2 Limitations

Despite the findings above, it should be noted that the learnheuristic algorithms are sensitive to specific parameters, such as the *n_offsprings*, *alpha*, *beta*, *n_max_doe*, *n_max_infills* and *seed* parameters. If one of these parameters is slightly off, the algorithms may prematurely terminate. Furthermore, the algorithms fail to consider large portfolios as increasing the number of stocks in the portfolio amplifies parameter sensitivity. Hence this study only considers the results produced from a portfolio of ten assets.

4.3 Future Work

4.3.1 Implementing Search Algorithms. Since the parameters of the algorithms explored are pretty sensitive, search algorithms, such as a grid search or a random search, can be utilised to find the most appropriate parameters to pass into the algorithms.

4.3.2 Increasing The Size of The Portfolio. Once stable parameters have been obtained, future research can analyse much larger portfolios, such as the complete S&P 500 index, instead of only the top ten considered in the current study.

4.3.3 Applying Alternative Multi-Objective Metaheuristics. Future research can be extended to the current study by implementing different multi-objective metaheuristic algorithms to solve the same problem identified in this study. Possible algorithms to consider include Particle Swarm Optimisation (PSO), Differential Evolution (DE) or Evolutionary Strategies (ES).

5 CONCLUSION

Portfolio optimisation problems must be solved while considering realistic constraints inevitable in practical trading and investment scenarios. This study introduces constraints, such as transaction costs, into a multi-objective portfolio optimisation problem using taking on a hybrid approach termed learnheuristics, which uses machine learning techniques, such as surrogate models, to enhance the metaheuristic algorithms being used to solve an optimisation problem, which is non-convex.

After implementing each of the four learnheuristic algorithms, it is evident that the NSGA-II learnheuristic algorithm is the best-performing algorithm as it reaches a higher hypervolume than the other three learnheuristic algorithms implemented. Furthermore, the hypervolume of each learnheuristic algorithm surpasses the hypervolume of the equivalent non-surrogate variant, which is our baseline algorithm, despite taking much longer for the learnheuristic algorithms to run to completion.

Based on the findings, we can conclude that the learnheuristic algorithms outperform the baseline algorithms regarding hypervolume and convergence rate. However, they take much longer to run to completion. Overall, using learnheuristics to obtain the weights of each asset for backtesting purposes results in a lower annual risk, higher annual expected return and higher Sharpe ratio compared to backtesting without using learnheuristics. Therefore, it can be concluded that learnheuristics produce superior and preferable results when solving a constrained, multi-objective portfolio optimisation problem.

This study also consists of a few limitations. Firstly, the learnheuristic algorithms are sensitive to many input parameters, and since we aim to reach optimal solutions, finding accurate parameters is vital. Secondly, the

implementation of the study is only run on a portfolio consisting of ten stocks, as adding more stocks to the portfolio does not result in optimal solutions due to the sensitivity of the learnheuristic parameters. Therefore, in future research, implementing a search algorithm, such as a random or grid search, to tune the parameters would be beneficial. This, in turn, may also assist with allowing the learnheuristic algorithms to consider portfolios with more than ten assets and still produce optimal results. Furthermore, alternative multi-objective metaheuristic algorithms can also be considered in future research.

REFERENCES

- Guido Abate, Tommaso Bonafini, and Pierpaolo Ferrari. 2022. Portfolio Constraints: An Empirical Analysis. International Journal of Financial Studies 10, 1 (2022), 9.
- [2] J. Blank and K. Deb. 2020. pymoo: Multi-Objective Optimization in Python. IEEE Access 8 (2020), 89497–89509.
- [3] Julian Blank and Kalyanmoy Deb. 2022. pysamoo: Surrogate-Assisted Multi-Objective Optimization in Python. arXiv:2204.05855 [cs.NE]
- [4] Julian Blank and Kalyanmoy Deb. 2022. pysamoo: Surrogate-Assisted Multi-Objective Optimization in Python. arXiv preprint arXiv:2204.05855 (2022).
- [5] Stephen Boyd, Enzo Busseti, Steve Diamond, Ronald N Kahn, Kwangmoo Koh, Peter Nystrup, Jan Speth, et al. 2017. Multi-period trading via convex optimization. Foundations and Trends® in Optimization 3, 1 (2017), 1–76.
- [6] Laura Calvet, Jésica de Armas, David Masip, and Angel A Juan. 2017. Learnheuristics: hybridizing metaheuristics with machine learning for optimization with dynamic inputs. Open Mathematics 15, 1 (2017), 261–280.
- [7] Jana Doering, Renatas Kizys, Angel A Juan, Angels Fito, and Onur Polat. 2019. Metaheuristics for rich portfolio optimisation and risk management: Current state and future trends. Operations Research Perspectives 6 (2019), 100121.
- [8] Ernestas Filatovas, Algirdas Lančinskas, Olga Kurasova, and Julius Žilinskas. 2017. A preference-based multi-objective evolutionary algorithm R-NSGA-II with stochastic local search. Central European Journal of Operations Research 25, 4 (2017), 859–878.
- [9] Renatas Kizys, Angel A Juan, Bartosz Sawik, and Laura Calvet. 2019. A biased-randomized iterated local search algorithm for rich portfolio optimization. Applied Sciences 9, 17 (2019), 3509.
- [10] Siddeeq Laher, Andrew Paskaramoorthy, and Terence L Van Zyl. 2021. Deep learning for financial time series forecast fusion and optimal portfolio rebalancing. In 2021 IEEE 24th International Conference on Information Fusion (FUSION). IEEE, 1–8.
- [11] Bin Li and Steven CH Hoi. 2014. Online portfolio selection: A survey. ACM Computing Surveys (CSUR) 46, 3 (2014), 1-36.
- [12] Yilin Ma, Ruizhu Han, and Weizhong Wang. 2021. Portfolio optimization with return prediction using deep learning and machine learning. Expert Systems with Applications 165 (2021), 113973.
- [13] HM Markowitz. [n. d.]. Portfolio Selection (1952) 7 J.
- [14] Robert Andrew Martin. 2021. PyPortfolioOpt: portfolio optimization in Python. Journal of Open Source Software 6, 61 (2021), 3066. https://doi.org/10.21105/joss.03066
- [15] Andrew B Paskaramoorthy, Tim J Gebbie, and Terence L van Zyl. 2020. A framework for online investment decisions. Investment Analysts Journal 49, 3 (2020), 215–231.
- [16] Rylan Perumal and Terence L van Zyl. 2020. Surrogate Assisted Methods for the Parameterisation of Agent-Based Models. In 2020 7th International Conference on Soft Computing & Machine Intelligence (ISCMI).
- [17] BY Qu, Q Zhou, JM Xiao, JJ Liang, and PN Suganthan. 2017. Large-scale portfolio optimization using multiobjective evolutionary algorithms and preselection methods. *Mathematical Problems in Engineering* 2017 (2017).
- [18] Haitham Seada and Kalyanmoy Deb. 2014. U-NSGA-III: A unified evolutionary algorithm for single, multiple, and many-objective optimization. COIN report 2014022 (2014).
- [19] Aravind Seshadri. 2006. A fast elitist multiobjective genetic algorithm: NSGA-II. MATLAB Central 182 (2006).
- [20] L Stander, M Woolway, and TL van Zyl. 2020. Extended surrogate assisted continuous process optimisation. In 2020 7th International Conference on Soft Computing & Machine Intelligence (ISCMI). IEEE, 275–279.
- [21] Felix Streichert, Holger Ulmer, and Andreas Zell. 2004. Evolutionary algorithms and the cardinality constrained portfolio optimization problem. In *Operations Research Proceedings 2003*. Springer, 253–260.
- [22] Terence L van Zyl, Matthew Woolway, and Andrew Paskaramoorthy. 2021. Parden: Surrogate assisted hyper-parameter optimisation for portfolio selection. In 2021 8th international conference on soft computing & machine intelligence (ISCMI). IEEE, 101–107.
- [23] Zihao Zhang, Stefan Zohren, and Stephen Roberts. 2020. Deep learning for portfolio optimization. The Journal of Financial Data Science 2, 4 (2020), 8–20.