



# Finding Near-Optimal Portfolios with Quality-Diversity

Bruno Gašperov<sup>(✉)</sup>, Marko Đurasević, and Domagoj Jakobovic

Faculty of Electrical Engineering and Computing, University of Zagreb,  
Zagreb, Croatia

{bruno.gasperov,marko.durasevic,domagoj.jakobovic}@fer.hr

**Abstract.** The majority of standard approaches to financial portfolio optimization (PO) are based on the mean-variance (MV) framework. Given a risk aversion coefficient, the MV procedure yields a single portfolio that represents the optimal trade-off between risk and return. However, the resulting optimal portfolio is known to be highly sensitive to the input parameters, i.e., the estimates of the return covariance matrix and the mean return vector. It has been shown that a more robust and flexible alternative lies in determining the entire region of near-optimal portfolios. In this paper, we present a novel approach for finding a diverse set of such portfolios based on quality-diversity (QD) optimization. More specifically, we employ the CVT-MAP-Elites algorithm, which is scalable to high-dimensional settings with potentially hundreds of behavioral descriptors and/or assets. The results highlight the promising features of QD as a novel tool in PO.

**Keywords:** quality-diversity · illumination algorithm · MAP-Elites · portfolio optimization · near-optimal portfolios

## 1 Introduction

### 1.1 Portfolio Optimization and Near-Optimal Portfolios

Portfolio optimization (PO) entails finding the optimal allocation of limited capital across a range of available financial assets within a specified timeframe. The optimality is typically defined with respect to a risk-adjusted return metric, such as the Sharpe ratio [1], or other metrics that account for the investor's specific risk preferences like the CARA utility function [2]. Classical approaches to PO are deeply rooted in the mean-variance (MV) framework, also referred to as modern portfolio theory (MPT) [3]. MV optimization provides a systematic methodology for constructing optimal portfolios that leverage the principle of diversification. By distributing investments among assets with dissimilar risk-return profiles (e.g., assets with mutually uncorrelated or negatively correlated returns), individual risks offset each other, thereby reducing the total portfolio risk. Given a risk aversion coefficient, the MV optimization process diversifies

assets to produce a single optimal portfolio representing the best risk-return trade-off. The set of optimal portfolios for different risk preferences constitutes a Pareto front called the efficient frontier.

However, the MV framework is predicated upon a number of simplified and unrealistic assumptions, including the normality of returns and the stationarity of the return covariance matrix. It also tends to produce portfolios highly concentrated in only a few assets, jeopardizing diversification. Furthermore, MV optimizers are particularly sensitive to estimation errors, with small changes in input parameters, especially expected return estimates, noticeably affecting the resulting optimal portfolio weights [4]. Such notorious issues have spurred researchers to propose novel PO approaches focusing on improving robustness. Various attempts have been made in this direction, including methods based on shrinking the covariance matrix [5] or expected returns [6], introducing constraints to the original MV framework [7], resampling the efficient frontier [8], or directly applying techniques from the area of robust optimization [9].

Another alternative, the focus of our work, involves identifying portfolios that are not strictly optimal but rather near-optimal (with respect to MV optimality). In the first phase (the optimization process), an entire subspace of mutually diverse<sup>1</sup> near-optimal portfolios is determined without focusing on any specific portfolio. This offers enhanced robustness to estimation errors [10], as it sidesteps the complex issues stemming from the interplay between the objective and the inputs, which are inherent to the MV optimization. In the second phase (the *a posteriori* analysis), the investor selects the final portfolio from the subspace of near-optimal portfolios. This decision-making scheme allows investors to incorporate their expert opinions, subjective views, or any other soft factors that may be challenging to integrate directly into an optimization problem formulation. Simultaneously, it permits the *ad-hoc* consideration of market frictions (e.g., transaction costs) and other regulatory, liquidity, and risk concerns, which are crucial in real-life PO but were not necessarily applied as constraints during the first phase. For instance, when rebalancing portfolio weights, investors may prefer a near-optimal portfolio similar to the current one over the MV optimal portfolio due to the former’s lower turnover, and hence also lower transaction costs. Therefore the consideration of near-optimal portfolios presents a fruitful opportunity for improving the process of PO.

## 1.2 Prior Research

**Near-Optimal Portfolios.** The problem of identifying and analyzing near-optimal portfolios has been scrutinized in several studies. Van der Schans and de Graaf [11] pioneered a novel methodology for constructing such portfolios, outlined as follows. Let  $\mathcal{F} : \mathcal{W} \rightarrow \mathcal{Z}$  be the mapping from the set of admissible portfolio weights  $\mathcal{W}$  to the risk-return space  $\mathcal{Z}$ . First, an optimal portfolio  $\mathbf{w}_0$  (satisfying  $\mathcal{F}(\mathbf{w}_0) \in \mathcal{E}$  where  $\mathcal{E} \subset \mathcal{Z}$  is the efficient frontier) is selected as a

---

<sup>1</sup> In terms of their distance in the space of admissible portfolio weights or, more generally, some behavior space.

reference portfolio in accordance with the investor’s risk preferences. A small region  $\mathcal{R} \subset \mathcal{Z}$  around  $\mathcal{F}(\mathbf{w}_0)$  is then defined and the portfolio  $\mathbf{w}_1$ , s.t.  $\mathcal{F}(\mathbf{w}_1) \in \mathcal{R}$ , located furthest away<sup>2</sup> from  $\mathbf{w}_0$ , is found. Subsequently, the portfolio  $\mathbf{w}_2$ , again s.t.  $\mathcal{F}(\mathbf{w}_2) \in \mathcal{R}$ , positioned furthest away from the convex hull  $\mathcal{H}$  spanned by previously found portfolios  $\{\mathbf{w}_0, \mathbf{w}_1\}$  is itself added to  $\mathcal{H}$ . Generally, in step  $i$ :

$$\mathbf{w}_i = \arg \max_{\mathcal{F}(\mathbf{w}) \in \mathcal{R}} d(\mathbf{w}, \mathcal{H}), \quad (1)$$

with  $d(\mathbf{w}, \mathcal{H}) = \min_{\mathbf{w}' \in \mathcal{H}} \|\mathbf{w} - \mathbf{w}'\|$ . This process iterates until  $\mathcal{S} \subset \mathcal{Z}$ , the convex hull in the risk-return space corresponding to the portfolios in  $\mathcal{H}$ , covers  $\mathcal{R}$  to the required level of precision  $\epsilon$ . Finally, a set of  $K$  diverse near-optimal portfolios  $\{\mathbf{w}_0, \mathbf{w}_1, \dots, \mathbf{w}_{K-1}\}$  is obtained. Since a convex combination of near-optimal portfolios is also near-optimal [10], any portfolio in  $\mathcal{H}$  acts as a viable option for the investor. The authors finally show that this region of near-optimal portfolios is more robust to input estimate uncertainties than a single optimal portfolio. However, this approach exhibits several limitations. Firstly, it involves solving a difficult non-convex optimization problem of finding the portfolio furthest from the convex hull, leading the authors to use a somewhat *ad hoc* combination of a support vector machine [12] and a basin-hopping algorithm [13]. Secondly, it becomes unfeasible in high-dimensionality settings, i.e., when faced with a large number of assets ( $N$ ), which is typically the case in modern PO<sup>3</sup>. Thirdly, it is restricted to finding near-optimal portfolios directly in the space of portfolio weights and is not easily applicable to other possible behavior (feature) spaces, which may comprise variables such as fundamentals, technicals, and risk factors. Put differently, it does not tackle the generalized variant where in step  $i$ :

$$\mathbf{w}_i = \arg \max_{\mathcal{F}(\mathbf{w}) \in \mathcal{R}} d(\phi(\mathbf{w}), \mathcal{H}) \quad (2)$$

where  $\mathcal{H} = \{\phi(\mathbf{w}_0), \dots, \phi(\mathbf{w}_{i-1})\}$  for some function  $\phi$ . To ameliorate some of these problems, Cajas [15] proposed the near-optimal centering (NOC) method, based on finding the analytic centers of near-optimal regions. The method can be used in conjunction with any convex risk measure and has been empirically demonstrated to lead to improved diversification and robustness when compared to traditional PO methods. Van Eeghen [14] expands on [11] by using polytope theory to inspect the structure and robustness of near-optimal regions. Moreover, the author proposes a new implementation of the method from [11] that reduces computation time while maintaining accuracy.

The topic of near-optimal portfolios has been investigated or touched upon in several other works as well. Chopra [16] uses a grid search to discover a subset of near-optimal portfolios and analyzes their composition. Benita *et al.* [17] emphasize that near-optimal portfolios might provide a higher degree of robustness to various scenarios due to significant differences in their makeups and provide an

<sup>2</sup> In the sense of the Euclidean distance between the portfolio weight vectors  $\|\mathbf{w}_1 - \mathbf{w}_0\|$ .

<sup>3</sup> Van Eeghen [14] reports computation times of around 2 hours and more per run already for  $N > 20$ .

illustrative example. Lastly, Fagerström and Oddshammar demonstrate that the Conditional Value-at-Risk (CVaR) optimization model tends to produce portfolios that are near-optimal under the MV framework [18], i.e., located very close to the efficient frontier.

**Evolutionary Computation and Quality-Diversity Optimization.** Evolutionary computation (EC) methods have a rich history of successful applications in PO [19], partly owing to their ability to handle non-convex search spaces that arise when real-world constraints (e.g., buy-in thresholds, turnover constraints, other regulatory and risk constraints) are imposed into the problem setting [20, 21]. However, there is a significant lack of research regarding the applications of exploration algorithms, such as quality-diversity (QD) [22] and novelty search (NS) [23] in PO and quantitative finance generally, despite their significant potential and their successes in other domains [24, 25]. The links between portfolio diversification and the divergent search paradigm, while arguably natural, remain understudied. Several somewhat related works nevertheless exist. Zhang *et al.* [26] address the specific problem of finding formulaic alpha factors that can predict and explain asset returns, making them suitable for use with multi-factor asset pricing models. To efficiently explore the space of formulaic alphas, with a focus on less frequently visited regions, they propose a search that combines QD and principal component analysis (PCA). This PCA-enhanced search is then used as an integral part of their `AutoAlpha` hierarchical evolutionary algorithm. Another approach is put forth by Yuksel [27], in which meta-learning QD optimization is employed to tackle the problem of large-scale sparse index tracking. It is concluded that the proposed method can be utilized in other scenarios where diversity among co-optimized solutions is needed, as well as in the presence of noisy reinforcement learning rewards.

### 1.3 Objectives and Contributions

In this paper, we tackle the problem of identifying a diverse set of near-optimal portfolios (i.e., portfolios with risk-return profiles located close to that of the reference optimal portfolio), as part of the broader PO problem. Specifically, we aim to answer the following research question (**Q**): *How to obtain a wide range of portfolios that are all near-optimal but mutually diverse in the portfolio weight space or the otherwise defined behavior space (BS)?* While previous research has found that, in some tasks, elite solutions are concentrated within a small part of the genotypic space (“the elite hypervolume”) [28], our task runs in the opposite direction, as it involves finding genotypically different solutions (portfolios) that are all elite. We set out to test the following hypothesis (**H**): *The combination of convergent and divergent search provided through QD algorithms can be leveraged to obtain a set of diverse near-optimal portfolios.* The hypothesis stems from the observation that QD algorithms provide a natural choice for the underlying problem due to their ability to yield a range of diverse yet high-performing solutions. While some related approaches exist [26, 27], to the best of our knowledge, this paper is

the first to approach MV-based PO via QD optimization. To ensure the scalability to high-dimensional behavioral and/or asset spaces, which are ubiquitous in modern finance, the approach is powered by the CVT-MAP-Elites algorithm, as vanilla MAP-Elites faces the curse of dimensionality. We first use a toy example with only three assets to show that the approach is competitive against a similar approach based on the construction of convex hulls [11], and later extend our investigation to a higher-dimensional setting. As will be shown, the experimental results collectively clearly point to the promising capabilities of QD as a novel tool in the arsenal of modern PO practitioners.

## 2 Methodology

### 2.1 Problem Formulation

Let  $\mathcal{W}$  be the set of all admissible portfolio weights:

$$\mathcal{W} = \left\{ (w_1, \dots, w_i, \dots, w_N) \mid w_i \geq 0 \quad \forall i, \quad \sum_{i=1}^N w_i = 1 \right\}, \quad (3)$$

where  $N \geq 2$  is the total number of assets and  $w_i$  denotes the portfolio weight in the  $i$ -th asset. For simplicity, short selling is not permitted, but it can be easily integrated into the framework if needed, by relaxing the  $w_i \geq 0$  constraint. Also, let  $b : \mathcal{W} \rightarrow \mathcal{B}$  be a behavior function, mapping  $\mathcal{W}$  to a BS denoted by  $\mathcal{B}$ . Assume that  $\mathcal{B}$  is split into  $M$  niches (regions), i.e.,  $\mathcal{B} = \mathcal{N}_1 \cup \mathcal{N}_2 \cup \dots \cup \mathcal{N}_M$ . Finally, let  $f : \mathcal{W} \rightarrow \mathbb{R}$  be a fitness function. Each candidate portfolio  $\mathbf{w}$  is then associated with its behavioral descriptor (BD)  $\mathbf{b}_{\mathbf{w}} \in \mathcal{B}$  and fitness value  $f(\mathbf{w}) \in \mathbb{R}$ . The goal is to find:

$$\forall \mathcal{N}_i \quad \mathbf{w}^* = \arg \max_{\mathbf{w}, \mathbf{b}_{\mathbf{w}} \in \mathcal{N}_i} f(\mathbf{w}). \quad (4)$$

### 2.2 Behavior Function and Space

We explore two different behavior functions and BS designs. In both cases, centroidal Voronoi tessellation (CVT) is used to partition the BS into niches.

**Portfolio Weights.** In the simplest variant, the behavior function  $b_1$  is set as an identity function, i.e.,  $\forall \mathbf{w}$ ,  $b_1(\mathbf{w}) = \mathbf{w}$ , with  $\mathcal{B}_1 = \mathcal{W}$ . Portfolio weight vectors  $\mathbf{w}$  simultaneously serve as both genotypes and phenotypes. Similar can be seen in some of the approaches used to tackle the Rastrigin function benchmark [29] through QD algorithms [30, 31].

**Asset's Fundamentals.** In another variant, we separate the two spaces (genotypic and phenotypic) and set  $b_2(\mathbf{w}) = \mathbf{p}_{\mathbf{w}}$ , where  $\mathbf{p}_{\mathbf{w}} \in \mathcal{B}_2$  is a vector that describes the fundamental properties of the assets that dominate  $\mathbf{w}$  weight-wise. More specifically:

$$\mathcal{B}_2 = \left\{ (s_1, \dots, s_i, \dots, s_L, c) \mid s_i \geq 0 \quad \forall i, \quad c > 0, \quad \sum_{i=1}^L s_i = 1 \right\}, \quad (5)$$

where  $s_i$  denotes the sector exposure of a portfolio to sector  $i$ , and  $L$  is the number of sectors. Sector exposure is defined as the sum of portfolio weights assigned to assets belonging to the respective sector, i.e.,  $s_i = \sum_{j \in S_i} w_j$ , with  $S_i$  denoting the set of indices of assets belonging to sector  $i$ . The variable  $c$  denotes normalized market capitalization. Note that, unlike in the previous case, it is not possible to uniformly sample from the BS directly. Also, any other asset characteristics or factors of importance to the investor (e.g., ESG<sup>4</sup> factors, geographical diversification, etc.) might be used instead.

### 2.3 Fitness Functions

**Fitness1** - The fitness (quality) function can be given by the negative distance between the risk profile of the candidate portfolio  $\mathbf{w}$  and that of the reference optimal portfolio  $\mathbf{w}_0$  (obtained via MV optimization):

$$f_1(\mathbf{w}) = -\|\mathcal{F}(\mathbf{w}_0) - \mathcal{F}(\mathbf{w})\| = -\|(\mu_0, \sigma_0) - (\mu, \sigma)\|, \quad (6)$$

where  $(\mu, \sigma) \in \mathcal{Z}$  is a vector consisting of the expected return and the volatility of the portfolio  $\mathbf{w}$  (and equivalently for  $\mathbf{w}_0$ ). This fitness function will be used to take the convex hull approximation method to the near-optimal region.

**Fitness2** - Alternatively, a modification of **Fitness1** is proposed:

$$f_2(\mathbf{w}) = \begin{cases} -\|(\mu_0, \sigma_0) - (\mu, \sigma)\|, & \text{if } (\mu, \sigma) \text{ not in } \mathcal{R} \\ \|\mathbf{w} - \mathbf{w}_0\|, & \text{otherwise} \end{cases} \quad (7)$$

where  $\mathcal{R}$  is the region of near-optimality around  $(\mu_0, \sigma_0)$ ,  $\mathcal{R} \subset \mathcal{Z}$ , defined as:

$$\mathcal{R} = \{(\mu, \sigma) \mid \mu \geq (1 - c)\mu_0, \sigma \leq (1 + c)\sigma_0\} \quad (8)$$

for some constant  $c \in (0, 1)$ . **Fitness2** promotes solutions in each niche whose risk profiles are as close as possible to the risk profile of  $\mathbf{w}_0$  until close proximity is reached. Once inside the near-optimality region  $\mathcal{R}$ , solutions that are furthest weight-wise from  $\mathbf{w}_0$  are preferred to ensure more genotypic diversity through an additional mechanism. Ideally, the distance from the convex hull of the entire archive of near-optimal solutions (instead of only from  $\mathbf{w}_0$ ) should be used as the measure. However, for computational efficiency, we employ this simple heuristic which will be shown to demonstrate strong performance in lower-dimensional settings.

### 2.4 Recombination Operator

Given the conditions in Eq. 3, a suitable constraint-preserving reproduction operator is required. To this end, a recombination operator (with mutation) that also includes clipping and normalization is used. It is described in great detail in Algorithm 1. Additional constraints, such as cardinality restrictions or buy-in thresholds, can be incorporated as needed.

<sup>4</sup> Environment, social and governance.

---

**Algorithm 1:** Recombination operator

---

**Inputs:** Parent portfolios  $\mathbf{w}_1, \mathbf{w}_2$ , mutation rate  $m$ **Result:** Child portfolio  $\mathbf{w}_3$ 

```

1  $\lambda \leftarrow \text{UniformRandom}[0, 1]$  ; // Randomly generated weight parameter
2  $\delta \leftarrow \text{UniformRandom}[-m, m, \text{len}(\mathbf{w}_1)]$  ; // Mutation vector
3  $\mathbf{w}_3 \leftarrow \lambda \mathbf{w}_1 + (1 - \lambda) \mathbf{w}_2 + \delta$  ; // Child portfolio
4  $\mathbf{w}_3 \leftarrow \text{Clip}(\mathbf{w}_3, 0, 1)$  ; // Clipping to ensure non-negative weights
5  $\mathbf{w}_3 \leftarrow \frac{\mathbf{w}_3}{\sum_{i=1}^{\text{len}(\mathbf{w}_3)} \mathbf{w}_{3,i}}$  ; // Normalization of the child portfolio
6 return  $\mathbf{w}_3$ 

```

---

## 2.5 Algorithm

Due to the continual increase in the number of investable securities in financial markets, modern PO typically operates in high-dimensional settings, potentially encompassing hundreds or even thousands of financial assets. The ensuing high-dimensionality can be tackled with the **CVT-MAP-Elites** algorithm, where the number of niches is constant and independent of the dimensionality of the BS [32]. It is a variant of the vanilla **MAP-Elites** algorithm, which is itself based on maintaining an archive of elite solutions, one for each niche in the BS. The underlying idea is to find a plethora of behaviorally diverse yet high-performing solutions. Unless noted otherwise, unstructured archive  $\mathcal{A}$  is assumed. For smaller BS dimensionalities (typically 2D to 6D [32]), the basic **MAP-Elites** algorithm [33] can be used instead. We use **CVT-MAP-Elites** in all of the performed experiments. The used hyperparameter values are provided separately for each of the experiments.

## 3 Experimental Results

### 3.1 Toy Example

We begin by considering the toy example first introduced in [16] and later revisited in [10], in which only three ( $N = 3$ ) asset classes - stocks, bonds, and treasury bills are considered<sup>5</sup>. The mean return, standard deviation, and correlation estimates for the three asset classes (given in Table 1) are used to construct the expected return vector  $\hat{\boldsymbol{\mu}}$  and the return covariance matrix  $\hat{\mathbf{S}}$  estimates. The resulting MV efficient frontier is shown in Fig. 1, accompanied by additional elements, including the near-optimality region  $\mathcal{R}$ . To ensure consistency with [10], the same reference portfolio  $\mathbf{w}_0$  is adopted as well as the identical value of  $c = 0.1$ . After running the QD algorithm  $R = 100$  times<sup>6</sup> for each

<sup>5</sup> More precisely, the assets include the S&P 500 market index, Lehman Brothers Long Term Government Bond Index, and one-month Treasury bills. The original data is presented monthly and spans the period from 1980 to 1990, but the estimates are transformed into annual values in our work.

<sup>6</sup> At the start of each QD run, niches are recalculated, discarding the old CVT results.

fitness function, while using  $\mathcal{B}_1$  and with  $M = 200$ , the convex hulls depicting regions of near-optimality in the portfolio weights space are obtained. Each run takes less than 100s even without any parallelization. The maximum number of evaluations is set to  $N_{max} = 250\,000$  and the number of CVT samples equals  $N_{CVT} = 10\,000$ . Random initialization is done until  $P_{init} = 10\%$  of niches are filled with solutions. Figure 2 presents the results for four different methods (**M1** - Chopra [16], **M2** - van der Schans and de Graaf [10, 11], **M3** - QD with *Fitness1*, and **M4** - QD with *Fitness2*), with more details given in Table 2. Larger hull surfaces (volumes or hypervolumes in higher-dimensional spaces) are desirable, as they indicate greater compositional diversity within the set of found near-optimal portfolios, allowing investors more freedom to accommodate their specific preferences. The best results overall are obtained by **M4**, the modified version of **M3**, as it benefits from the additional diversification mechanism. The dominance of **M2** over **M3** with respect to the convex hull surface is not a matter of concern but rather an anticipated outcome. Namely, **M2** selects portfolios (exclusively from  $\mathcal{R}$ ) by directly maximizing distances from the convex hull of previously found solutions, whereas **M3** seeks solutions within each niche whose risk profiles are strictly closest to  $\mathcal{F}(\omega_0)$ . As a result, **M3** pushes strongly towards  $\mathcal{F}(\omega_0)$  (lying on the efficient frontier) regardless of whether exploring inside or outside of  $\mathcal{R}$  and hence shrinks  $\mathcal{S}$ , the corresponding region in the risk-return space. We therefore suspect **M3** to yield portfolios with better risk-reward profiles (measured by the Sharpe ratio) than **M2** and also a lower  $\mathcal{S}$  surface, all of which is indeed confirmed by the results laid out in Table 1. It is also noteworthy that **M4** achieves an even higher value of the Sharpe ratio than **M3**, despite its emphasis on genotypic diversity. The likely reason is that, given that  $\omega_0$  is not the maximum Sharpe portfolio, it is possible for the neighboring risk profiles in  $\mathcal{R}$  to dominate over it.


**Table 1.** The estimates from historical data (annual)

	Mean (%)	Std (%)	Corr. stocks	Corr. bonds	Corr. T-Bills	Optimal weights (%) (moderate risk aversion)
Stocks	15.876	16.603	1.000	–	–	58.1
Bonds	12.324	13.801	0.341	1.000	–	22.8
T-Bills	8.748	0.759	–0.081	0.050	1.000	19.1

### 3.2 Higher-Dimensional Setting

In this section, encouraged by the positive results from the toy example, we apply the proposed method to a higher-dimensional setting comprised of a large number of assets. Under such conditions, QD is expected to offer heuristic solutions in a reasonable time. More specifically, we consider the universe of  $N = 105$  different equity assets covering  $L = 11$  different sectors. The Sharpe optimal portfolio is used as the reference portfolio. Moreover, the Ledoit-Wolf shrinkage





**Fig. 1.** The MV efficient frontier with the chosen optimal portfolio and its region of near-optimality.

**Table 2.** Performance evaluation of the four tested methods

Method	Convex hull $\mathcal{H}$ surface	Risk-return subspace $\mathcal{S}$ surface ( $\times 10^3$ )	Sharpe ratio of portfolios	% of niches with opt. port
<b>M1</b> <sup>a</sup>	0.0948	0.1973	$1.1368 \pm 0.0959$	—
<b>M2</b> <sup>b</sup>	0.1724	0.3712	$1.1762 \pm 0.1299$	—
<b>M3</b>	$0.1529 \pm 0.0015$	$0.3652 \pm 0.0076$	$1.2696 \pm 0.0009$	$48.980 \pm 0.003$
<b>M4</b>	<b><math>0.1746 \pm 0.0002</math></b>	$0.3882 \pm 0.0008$	<b><math>1.2824 \pm 0.0009</math></b>	<b><math>49.070 \pm 0.003</math></b>


<sup>a</sup>Results from [16] (deterministic).

<sup>b</sup>Results from [10] (deterministic).

[5] is used<sup>7</sup> for estimating the covariance matrix  $\hat{\mathbf{S}}$ , while CAPM returns [34] are utilized to derive the mean return estimate  $\hat{\boldsymbol{\mu}}$ . The look-back window used for parameter estimation encompasses the period from January 2nd, 2020, to September 1st, 2023, corresponding to exactly  $T = 924$  trading days. We again set  $c = 0.01$ , leading to a very strict criterion of near-optimality (i.e., a small surface of  $\mathcal{R}$ ), and also select  $M = 5000$ . Due to a plethora of assets, investors might prefer to first perform asset allocation at the sector/industry or some other macro level<sup>8</sup> before selecting individual constituent securities. Consequently, we use the BS design  $\mathcal{B}_2$ . As the fitness function,  $f_1$  is selected. The algorithm ultimately returns an archive of solutions  $\mathcal{A}$  comprised of both near-optimal portfolios and non-near-optimal portfolios (for niches in which no near-optimal portfolios have been found).

<sup>7</sup> With constant variance set as the shrinkage target.

<sup>8</sup> Similar approaches are employed in top-down investment strategies such as Tactical Asset Allocation (TAA) [35].



**Fig. 2.** Convex hulls in the genotypic/phenotypic space for different methods for a single QD run. The constraint  $\sum x_i = 1$  introduces coplanarity in the otherwise three-dimensional space.

Considering both the absence of suitable benchmarks and the computational challenges in calculating convex hull volumes in spaces with higher dimensionalities<sup>9</sup>, we draw inspiration from previous research on benchmarking QD algorithms [37] and use it as a starting point in creating appropriate performance metrics. All of the used metrics are described in Table 3. Remark that metrics  $C'$ , QDScore2 and AP2 only consider niches in which near-optimal portfolios have been found, unlike QDScore1 and AP1, which (also) require information on other niches. Following the notation from Table 3, the number of niches with and without near-optimal portfolios is  $C'M$  and  $|\mathcal{A}| - C'M$ , respectively.


**Experimental Results.** Figure 3 presents the main experimental results. The upper subplots show the first three metrics ( $C'$ , QDScore1, QDScoreMOD) plotted against the number of QD evaluations, as well as the archive profiles (AP1 and AP2). The mean values as averaged over  $R = 20$  runs are shown, together with empirical percentiles (the 5-th and the 95-th  $p^{th}$ ). As before, niches are recalculated from scratch at the beginning of each QD run. To ensure the high quality of tessellation, the maximum number of QD evaluations is set to  $N_{max} = 2.2 \times 10^6$ ,  $P_{init} = 10\%$ , and the number of CVT samples equals  $N_{CVT} = 50\,000$ . In terms of the modified coverage, it is evident that the percentage of niches with near-optimal solutions sharply rises, reaching 60% already at a bit over 500 000 evaluations, and finally getting to approximately 75%. Such high percentages clearly indicate that, despite the use of a stringent criterion for near-optimality, the method succeeds in finding a wide range of near-optimal

<sup>9</sup> With the QuickHull algorithm [36], the execution time grows by  $n^{\lfloor \frac{d}{2} \rfloor}$ , where  $n$  is the input size and  $d$  the dimensionality.

**Table 3.** List of metrics

Metric	Description	Expression
Modified coverage ( $C'$ )	The proportion of niches with near-optimal portfolios	$\frac{\text{No. of niches with n.o. portfolios}}{\text{No. of niches}}$
QD-score (QDScore1)	The cumulative normalized fitness of all portfolios in $\mathcal{A}$	$\sum_{i=1}^{ \mathcal{A} } \frac{f_i - \min_j(f_j)}{\max_j(f_j) - \min_j(f_j)}$
Modified QD-score (QDScoreMOD)	The cumulative normalized fitness of all found near-optimal portfolios	$\sum_{i=1}^{C'M} \frac{f_i - \min_j(f_j)}{\max_j(f_j) - \min_j(f_j)}$
Modified archive profile 1 (AP1)	The proportion of found non-near-optimal portfolios exceeding some threshold value	$\sum_{i=1}^{ \mathcal{A}  - C'M} \mathbb{1}(f_i \geq f_{\text{threshold}})$
Modified archive profile 2 (AP2)	The proportion of found near-optimal portfolios exceeding some threshold value	$\sum_{i=1}^{C'M} \mathbb{1}(f_i \geq f_{\text{threshold}})$

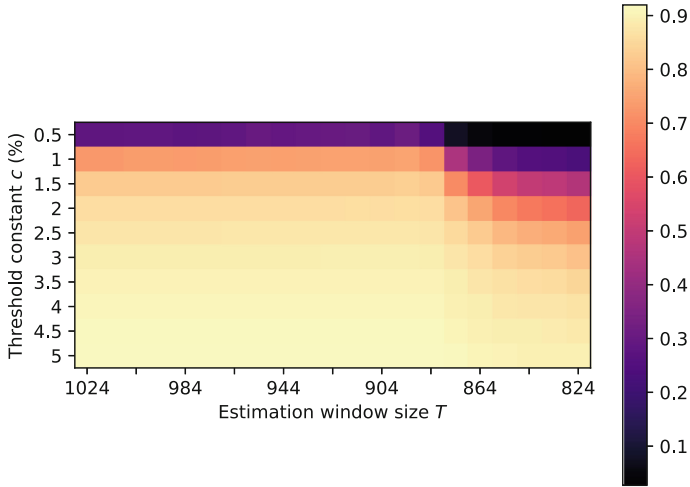
portfolios encompassing various industry compositions and market capitalization values. This is a fortunate conclusion, especially for investors who prefer human-in-the-loop approaches, enabling them to incorporate their own preferences that might be hard to formalize, into the decision-making process. As for the QD-score metrics, we observe that QDScore1 steadily increases over time, indicating an improvement in diversity/performance among the found near-optimal solutions (portfolios). Likewise, the QDScore1 curve shows that the majority of niches are populated relatively quickly (in under 500 000 evaluations). Also note that QDScore1 expectedly converges faster than QDScoreMOD, as filling niches with near-optimal solutions takes more time compared to arbitrary solutions. Modified archive profiles, calculated after performing the maximum number of QD evaluations, are provided at the bottom of Fig. 3. While AP1 and AP2 have similar shapes (under different threshold scales), a small right tail can be seen with AP2, showing a number of "super" near-optimal portfolios with risk profiles extremely similar to that of the reference portfolio. The significant compositional diversity of the obtained portfolios is depicted in Fig. 5, which displays the Sharpe optimal portfolio alongside two mutually diverse near-optimal portfolios generated by the method in the feature (behavior) space. Despite the multitude of potential solutions, in order to finalize the investment decision-making process, it is necessary to select a single portfolio from the set of obtained near-optimal portfolios. To this end, Algorithm 2 delineates the entire end-to-end investment decision-making process, incorporating the proposed QD method.




**Fig. 3.** The mean modified coverage  $C'$  (the upper left subplot), QDScore1, and QDScore2 (the upper right subplot) plotted against the number of evaluations, alongside empirical percentiles. The modified archive profiles are displayed at the bottom, with the set including non-near-optimal portfolios on the left and the set including near-optimal portfolios on the right.

**Robustness.** Generally, the evaluation of QD solutions is exacerbated by the presence of stochasticity in the underlying environment [29]. In our case, the evaluation is fully deterministic once the estimates  $(\hat{\mu}, \hat{S})$  are fixed. However, there is stochasticity involved due to the very fact that “true” parameter values  $(\mu, S)$  are hidden, whereas the estimates (which can be derived in multiple ways) represent random variables. With this in mind, we investigate the robustness of the generated portfolios to a certain type of change in the estimates  $(\hat{\mu}, \hat{S})$ . More specifically, we study whether previously found near-optimal portfolios generally remain near-optimal when re-estimating under different estimation window<sup>10</sup> sizes  $T$ . The results are shown in Fig. 4. As expected, the mean modified coverage  $C'$  remains robust to changes in the estimation window size  $T$  when larger threshold constants  $c$  are employed. On the other hand, the sensitivity of  $C'$  to changes in  $T$ , in particular to its reduction, is much more emphasized for  $c$  values in the range  $[0.5\%, 2.5\%]$ , i.e., for stricter near-optimality criteria. For example, with  $c = 0.5\%$  and  $T = 824$ , on average only 136.15 niches

<sup>10</sup> The choice of the estimation window size is a non-trivial issue that has been studied before [38–40], with larger sizes leading to reduced estimation errors at the price of assuming unrealistically long stationarity periods.



**Fig. 4.** The mean modified coverage  $C'$  for different threshold constants  $c$  and estimation window sizes  $T$ . Observe the relatively high sensitivity of  $C'$  to the shortening of  $T$ , especially for stricter near-optimality criteria (i.e., for smaller  $c$  values).



**Fig. 5.** Two of the obtained near-optimal portfolios juxtaposed against the Sharpe optimal portfolio in the BS. The first (second) near-optimal portfolio is highly concentrated in industrials (financial services) stocks and low (high) market cap stocks, while the Sharpe optimal portfolio remains more balanced.

(or 2.7%) contain near-optimal solutions. We leave for further research the study of whether solutions that remain near-optimal under a wider range of estimates present a superior investment choice.

---

**Algorithm 2:** Portfolio selection process

---

**Inputs:** Investor’s risk aversion  $\gamma$  and preferred BD  $\mathbf{b}$ , historical data  $\mathcal{D}$ **Result:** Final portfolio  $\mathbf{w}$ 

```

1  $\hat{\mathbf{S}}, \hat{\boldsymbol{\mu}} \leftarrow \text{EstimateParameters}(\mathcal{D})$ 
2  $\mathbf{w}_0 \leftarrow \text{CalculateEfficientFrontierPortfolio}(\gamma, \hat{\mathbf{S}}, \hat{\boldsymbol{\mu}})$ ; // MV step - calculating
   the required portfolio on the efficient frontier
3  $\mathcal{A} \leftarrow \text{RunQD}$ ; // Run QD to obtain the archive of portfolios
4  $n_b \leftarrow \text{DetermineNicheIndex}(\mathbf{b})$ ; // Determine the niche index for the
   preferred BD
5 if NearOptimalPortfolioExistsIn( $n_b$ ) then
6 |  $\mathbf{w} \leftarrow \text{ElitePortfolioFrom}(n_b)$ 
7 else
8 |  $n'_b \leftarrow \text{ClosestNicheWithNearOptimalPortfolio}(n_b)$ ; // Among all niches
   with a near-optimal portfolio, select the one closest to  $n_b$ 
9 |  $\mathbf{w} \leftarrow \text{ElitePortfolioFrom}(n'_b)$ 
10 end
11 return  $\mathbf{w}$ 

```

---

## 4 Conclusion and Further Work

This paper is concerned with the problem of finding mutually diverse portfolios located in the region of near-optimality of some reference optimal portfolio. We introduce a novel method for discovering a wide spectrum of mutually diverse (in a predefined BS) near-optimal portfolios based on QD optimization. The main findings, pointing to high coverage and QD-scores, underscore the capacity of QD to serve as a novel instrument in the field of PO. In addition to QD, the **knobelt** algorithm [41] might be used to balance compositional diversity (*novelty*) with proximity to the selected optimal portfolio. Similar approaches might also be employed for a somewhat different problem; to approximate the entire efficient frontier of optimal solutions, covering a wide range of different risk preferences and hence catering to a versatile set of investors. These include QD with a fitness function that penalizes distances from the efficient frontier in each niche, as well as vanilla NS formulations in which the risk-reward space is used as the phenotypic space. Other objectives besides MV may be explored in the future as well, together with different BS designs (e.g. factor-based, distance from the equally weighted or currently selected portfolio) and definitions of near-optimality (e.g. those employing soft constraints). Links with sparse PO should also be investigated. Lastly, we anticipate further work to leverage the power of QD in visualizing and illuminating the portfolio search space. More broadly, we hope to see future approaches harnessing the potential of open-endedness-based approaches, including QD and NS, in computational finance.

## References

1. Sharpe, W.F.: The sharpe ratio. *Streetwise-the Best J. Portfolio Manag.* **3**, 169–185 (1998)
2. Babcock, B.A., Choi, E.K., Feinerman, E.: Risk and probability premiums for cara utility functions. *J. Agricult. Resource Econ.* **22**, 17–24 (1993)
3. Markowitz, H.M., Todd, G.P.: *Mean-variance analysis in portfolio choice and capital markets*, vol. 66. John Wiley & Sons (2000)
4. Best, M.J., Grauer, R.R.: On the sensitivity of mean-variance-efficient portfolios to changes in asset means: some analytical and computational results. *Rev. Financial Stud.* **4**(2), 315–342 (1991)
5. Ledoit, O., Wolf, M.: Honey, i shrunk the sample covariance matrix. *UPF economics and business working paper*, vol. (691) (2003)
6. Black, F., Litterman, R.: Asset allocation: combining investor views with market equilibrium. *Goldman Sachs Fixed Income Res.* **115**(1), 7–18 (1990)
7. DeMiguel, V., Garlappi, L., Nogales, F.J., Uppal, R.: A generalized approach to portfolio optimization: improving performance by constraining portfolio norms. *Manag. Sci.* **55**(5), 798–812 (2009)
8. Michaud, R.O., Michaud, R.O.: *Efficient asset management: a practical guide to stock portfolio optimization and asset allocation*. Oxford University Press (2008)
9. Yin, C., Perchet, R., Soupé, F.: A practical guide to robust portfolio optimization. *Quantitative Finance* **21**(6), 911–928 (2021)
10. de Graaf, T.: *Robust Mean-Variance Optimization*. PhD thesis, Master Thesis, Leiden University & Ortec Finance (2016)
11. van der Schans, M., de Graaf, T.: Robust optimization by constructing near-optimal portfolios. Available at SSRN 3057258 (2017)
12. Wang, L.: *Support vector machines: theory and applications*, vol. 177. Springer Science & Business Media (2005). <https://doi.org/10.1007/b95439>
13. Wales, D.J., Doye, J.P.K.: Global optimization by basin-hopping and the lowest energy structures of lennard-jones clusters containing up to 110 atoms. *J. Phys. Chem. A* **101**(28), 5111–5116 (1997)
14. van Eeghen, W.J.B., van Gaans, O.W., van der Schans, M.: *Analysis of near-optimal portfolio regions and polytope theory* (2018)
15. Cajas, D.: Robust portfolio selection with near optimal centering. Available at SSRN 3572435(2019)
16. Vijay Kumar Chopra: Improving optimization. *J. Invest.* **2**(3), 51–59 (1993)
17. Benita, G., Baudot-Trajtenberg, N., Friedman, A.: The challenges of managing large fx reserves: the case of israel. *BIS Paper*, (104m) (2019)
18. Fagerström, S., Oddshammar, G.: *Portfolio optimization-the mean-variance and cvar approach* (2010)
19. Brabazon, A., O’Neill, M., Dempsey, I.: An introduction to evolutionary computation in finance. *IEEE Comput. Intell. Mag.* **3**(4), 42–55 (2008)
20. Branke, J., Scheckenbach, B., Stein, M., Deb, K., Schmeck, H.: Portfolio optimization with an envelope-based multi-objective evolutionary algorithm. *Eur. J. Oper. Res.* **199**(3), 684–693 (2009)
21. Qi, R., Yen, G.G.: Hybrid bi-objective portfolio optimization with pre-selection strategy. *Inform. Sci.* **417**, 401–419 (2017)
22. Chatzilygeroudis, K., Cully, A., Vassiliades, V., Mouret, J.-B.: Quality-diversity optimization: a novel branch of stochastic optimization. In: Pardalos, P.M., Rasskazova, V., Vrahatis, M.N. (eds.) *Black Box Optimization, Machine Learning, and*

- No-Free Lunch Theorems. SOIA, vol. 170, pp. 109–135. Springer, Cham (2021). [https://doi.org/10.1007/978-3-030-66515-9\\_4](https://doi.org/10.1007/978-3-030-66515-9_4)
23. Lehman, J., Stanley, K.O.: Novelty search and the problem with objectives. Genetic programming theory and practice IX, pp. 37–56 (2011)
  24. Gomes, J., Urbano, P., Christensen, A.L.: Evolution of swarm robotics systems with novelty search. *Swarm Intell.* **7**, 115–144 (2013)
  25. Pugh, J.K., Soros, L.B., Stanley, K.O.: Quality diversity: a new frontier for evolutionary computation. *Front. Robot. AI* **3**, 40 (2016)
  26. Zhang, T., Li, Y., Jin, Y., Li, J.: Autoalpha: an efficient hierarchical evolutionary algorithm for mining alpha factors in quantitative investment. arXiv preprint [arXiv:2002.08245](https://arxiv.org/abs/2002.08245) (2020)
  27. Yuksel, K.A.: Generative meta-learning robust quality-diversity portfolio. In: Proceedings of the Companion Conference on Genetic and Evolutionary Computation, pp. 787–790 (2023)
  28. Vassiliades, V., Mouret, J.-P.: Discovering the elite hypervolume by leveraging interspecies correlation. In: Proceedings of the Genetic and Evolutionary Computation Conference, pp. 149–156 (2018)
  29. Digalakis, J.G., Margaritis, K.G.: On benchmarking functions for genetic algorithms. *Inter. J. Comput. Math.* **77**(4), 481–506 (2001)
  30. Bossens, D.M., Tarapore, D.: Quality-diversity meta-evolution: customising behaviour spaces to a meta-objective. arXiv preprint [arXiv:2109.03918](https://arxiv.org/abs/2109.03918) (2021)
  31. Sfikas, K., Liapis, A., Yannakakis, G.N.: Monte carlo elites: Quality-diversity selection as a multi-armed bandit problem. In: Proceedings of the Genetic and Evolutionary Computation Conference, pp. 180–188 (2021)
  32. Vassiliades, V., Chatzilygeroudis, K., Mouret, J.-B.: Using centroidal voronoi tessellations to scale up the multidimensional archive of phenotypic elites algorithm. *IEEE Trans. Evol. Comput.* **22**(4), 623–630 (2017)
  33. Mouret, J.-B., Clune, J.: Illuminating search spaces by mapping elites. arXiv preprint [arXiv:1504.04909](https://arxiv.org/abs/1504.04909) (2015)
  34. Fama, E.F., French, K.R.: The capital asset pricing model: theory and evidence. *J. Econ. Perspect.* **18**(3), 25–46 (2004)
  35. Faber, M.: A quantitative approach to tactical asset allocation. *J. Wealth Manag.* Spring (2007)
  36. Barber, C.B., Dobkin, D.P., Huhdanpaa, H.: Qhull: Quickhull algorithm for computing the convex hull. Astrophysics Source Code Library, pp. ascl-1304 (2013)
  37. Flageat, M., Lim, B., Grillotti, L., Allard, M., Smith, S.C., Cully, A.: Benchmarking quality-diversity algorithms on neuroevolution for reinforcement learning. arXiv preprint [arXiv:2211.02193](https://arxiv.org/abs/2211.02193) (2022)
  38. Gašperov, B., Šarić, F., Begušić, S., Kostanjčar, Z.: Adaptive rolling window selection for minimum variance portfolio estimation based on reinforcement learning. In: 2020 43rd International Convention on Information, Communication and Electronic Technology (MIPRO), pp. 1098–1102. IEEE (2020)
  39. Wang, P.-T., Hsieh, C.-H.: On data-driven log-optimal portfolio: a sliding window approach. *IFAC-PapersOnLine* **55**(30), 474–479 (2022)
  40. Chuanzhen, W.: Window effect with markov-switching garch model in cryptocurrency market. *Chaos, Solitons Fractals* **146**, 110902 (2021)
  41. Kelly, J., Hemberg, E., O’Reilly, U.-M.: Improving genetic programming with novel exploration - exploitation control. In: Sekanina, L., Hu, T., Lourenço, N., Richter, H., García-Sánchez, P. (eds.) EuroGP 2019. LNCS, vol. 11451, pp. 64–80. Springer, Cham (2019). [https://doi.org/10.1007/978-3-030-16670-0\\_5](https://doi.org/10.1007/978-3-030-16670-0_5)