Application of portfolio optimization to drug discovery

Yevseyeva, I.; Lenselink, E.B.; Vries, A. de; IJzerman, A.P.; Deutz, A.H.; Emmerich, M.T.M.

Application of portfolio optimization to drug discovery

Iryna Yevseyeva[1,3], Eelke B. Lenselink[2], Alice de Vries[3],
Adriaan P. IJzerman[2], André H. Deutz[3] and Michael T.M. Emmerich[3]

[1]*School of Computer Science and Informatics, Faculty of Technology, De Montfort University, LE1 9BH, Leicester, UK*
[2]*Medicinal Chemistry/Leiden Academic Centre for Drug Research, Leiden University, 2333-CA Leiden, The Netherlands*
[3]*Leiden Institute of Advanced Computer Science, 2333-CA Leiden, The Netherlands*

Abstract

In this work, a problem of selecting a subset of molecules, which are potential lead candidates for drug discovery, is considered. Such molecule subset selection problem is formulated as a portfolio optimization, well known and studied in financial management. The financial return, more precisely the return rate, is interpreted as return rate from a potential lead and calculated as a product of gain and probability of success (probability that a selected molecule becomes a lead), which is related to performance of the molecule, in particular, its (bio-)activity. The risk is associated with not finding active molecules and is related to the level of diversity of the molecules selected in portfolio. It is due to potential of some molecules to contribute to the diversity of the set of molecules selected in portfolio and hence decreasing risk of portfolio as a whole. Even though such molecules considered in isolation look inefficient, they are located in sparsely sampled regions of chemical space and are different from more promising molecules. One way of computing diversity

of a set is associated with a covariance matrix, and here it is represented by the Solow-Polasky measure. Several formulations of molecule portfolio optimization are considered taking into account the limited budget provided for buying molecules and the fixed size of the portfolio. The proposed approach is tested in experimental settings for three molecules datasets using exact and/or evolutionary approaches. The results obtained for these datasets look promising and encouraging for application of the proposed portfolio-based approach for molecule subset selection in real settings.

Keywords: Portfolio Approach, Multicriteria optimization, Decision Support, Drug Discovery

1. Introduction

When searching for the most promising drug like molecules for a drug discovery project, usually, de novo drug discovery uses *in vitro* experiments (colloquially called "test-tube experiments"). For this, circa 100 promising molecules are selected from a database and typically only circa 1 percent of the molecules are tested successfully *in vitro*, that is they become so called *lead* molecules [4]. High-throughput screening (HTS) allows for testing a large number of molecules by robotized machines using advanced laboratory equipment. However, testing *in vitro* is an expensive process and cannot always be applied to all possible projects, even though in industry millions of molecules can be screened if the target is interesting enough. To reduce costs, HTS can be complemented by preliminary *in silico* (performed via computer

simulation) virtual screening (VS). VS approaches [27] are used to pre-select molecules from virtual libraries or large databases of commercially available molecules (e.g. ZINC [13]) based on their chemical properties. Selection is typically done based on the assessment of the success probability of candidate molecules using either a compound-based method or a target-based method or a combination of both. Typically, no explicit economical information is taken into account and simple methods like clustering are applied. The success probability is not given directly, but in the form of a score corresponding to (bio-)activity that is proportional to it.

A typical scenario in a pharmaceutical research laboratory is that a chemist selects a subset from a large vendor database of molecules (e.g. ZINC) and orders these. Each molecule has a price, and the budget of the chemist is limited, but it has to be allocated for buying molecules. That is, money not spent cannot be used for another purpose (and, thus, will be lost). Note that here we do not look into experimental planning and drug production, see e.g. [1], which is a separate subject of research.

Classical approaches for selecting promising molecules are based on predicted activity score. However, selecting molecules based on their performance (success probabilities / activity scores) only is not enough and even risky. It is due to a high probability of selecting well-performing, but similar molecules, which all might be unsuccessful for the same reason.

An alternative approach is to take into account diversity of selected subsets of molecules. However, selection purely based on diversity will neglect

the information about activity given to the chemists by VS models. Moreover, price might also play a role in the choice as it influences the number of molecules that can be bought given a limited budget. Hence, existing, just clustering- or just scoring-based selection models are not sufficient for handling these problems.

In this work, we consider the first stage of drug discovery of identifying lead candidates with an approach which takes into account performance of molecules according to their predicted (bio-) activity and diversity of selected molecules simultaneously. Similar approach was taken in [17], where activity score and diversity were maximized at the same time. Here, the molecule subset selection problem is considered and modeled by analogy with a well-known financial portfolio selection problem, see e.g. [5]. A similar binary problem of finding an optimal combination of items subject to constraints is known in operations research as the knapsack problem, see e.g. [16]. According to the portfolio optimization approach when selecting a subset of molecules to be tested *in vitro*, in addition to choosing molecules with highest performance values (and maximizing the average quality of the selected subset of molecules), the molecules with the most dissimilar structures should be considered. The former aspect contributes to maximizing the quality of the selected subset of molecules and the latter one corresponds to maximizing the diversity of such a subset.

The financial portfolio return, more precisely the return rate, is interpreted as the return rate from a potential lead and calculated as a prod-

uct of the gain and the probability of success (probability that a selected molecule becomes a drug in the end), which is related to the performance of the molecule, in particular, its (bio-)activity. The risk is associated with not finding active molecules when choosing a portfolio and is related to the level of diversity of the molecules in the portfolio. The diversity can be expressed as a covariance matrix used by Solow and Polasky [23] for measuring diversity of a biological population. Interestingly, as an example of a utilitarian approach to the biological diversity preservation, Solow and Polasky indicated the potential utility in future from one of the preserved species as a cure of some yet unknown disease (see [23]).

Some molecules, when considered in isolation look inefficient, but as part of a portfolio may contribute to the decreasing risk of a portfolio as a whole and may be included in a portfolio as they are located in sparsely sampled regions of chemical space and are different from more promising molecules. In addition, the limited budget provided for buying molecules and the fixed size of the portfolio are taken into account in the introduced drug portfolio model as constraints.

This article is structured as follows: In the next section 2, we consider the general (multiobjective) formulation of the (financial) portfolio selection problem and, then, in section 3, we model the lead subset selection problem as portfolio optimization. In section 4, we propose algorithms to solve portfolio selection formulations, and in section 5, we discuss results obtained for three molecule datasets. Finally, in section 6, we draw conclusions and indicate

directions for future research.

2. Related Work

2.1. Portfolio selection as a multi-objective optimization problem

The most-widely used formulation of portfolio selection problem was developed by Markowitz early in the 50s [15]. It addresses a way of selecting a combination of several assets called *portfolio* that collectively would be of the best quality and be as diverse as possible. Hence, portfolio optimization should simultaneously satisfy two conflicting goals, minimizing risk and maximizing expected return of the portfolio, that is formally:

$$\begin{aligned} \min \sigma^2(\mathbf{x}) = \sum_{i=1}^{N_{Total}} \sum_{j=1}^{N_{Total}} q_{ij} x_i x_j &= \mathbf{x}^\top \mathbf{Q} \mathbf{x}; \\ \max E(\mathbf{x}) = \sum_{i=1}^{N_{Total}} r_i x_i &= \mathbf{r}^\top \mathbf{x}; \\ \text{s.t.} \sum_{i=1}^{N_{Total}} x_i &= 1; \\ x_i \in [0,1], i &= 1, \ldots, N_{Total}, \end{aligned} \tag{1}$$

where N_{Total} is the number of assets; x_i is the proportion of money invested in the asset i; r_i is the expected return (per period) of the asset i; and q_{ij} is the real-valued covariance of expected returns of the assets i and j.

As a result of optimizing this problem not a single portfolio but a set of portfolios are selected that are optimal with respect to the two specified objectives.

For this problem the search space of portfolios $\mathcal{S}$ is $[0,1]^{N_{Total}}$. The set of feasible portfolios $\mathcal{F}$ is the subset of portfolios in $\mathcal{S}$ with $\sum_{i=}^{N_{Total}} x_i = 1$. We consider two real valued objective functions defined on $\mathcal{S}$, $\sigma^2(\mathbf{x}) = \mathbf{x}^\top \mathbf{Q}\mathbf{x}$ and $E(\mathbf{x}) = \mathbf{r}^\top \mathbf{x}$. Each portfolio $\mathbf{x}$ is associated with a 2-dimensional evaluation vector in the objective space, $(\sigma^2(\mathbf{x}), E(\mathbf{x}))^\top$, where the risk objective is to be minimized and the return objective is to be maximized.

Optimizing two or more conflicting objectives simultaneously is referred to as Multiobjective Optimization (MOO). The portfolio selection problem formulated as in (1) is bi-objective: Minimizing the risk and maximizing the expected return should be taken into consideration and optimized at the same time. These objectives are generally in conflict with each other and finding a portfolio with minimal risk and maximal return simultaneously is infeasible. Hence, decreasing risk for a portfolio can be obtained at the cost of lowering its return only.

Interestingly, including some assets, which look inefficient when considered in isolation, may benefit the portfolio as a whole, since they contribute to decreasing the risk of a portfolio when considered in combination with other assets. This is due to their location in sparsely sampled regions of search space and their difference from more promising assets. Cost of assets may also be taken into account as a separate objective, but we included it in the return (which is reduced by the costs invested in initial assets) and in the budget constraint.

Recently, the principles of portfolio optimization have been successfully

applied not only for optimizing financial portfolio selection [24], but also in other domains, such as strategic decision making [14] (for instance, team management), projects selection [11], IT project portfolio management [3], and evolutionary algorithms selection [29]. For instance, for evolutionary algorithms selection it is important to keep good, but different individuals, which should avoid fast convergence of the population to a single individual or few similar individuals. Hence, the selection procedure should simultaneously optimize quality and diversity of population. In [29], a multiobjective evolutionary algorithm based on the portfolio selection idea was introduced and results comparable to the results of the state-of-the-art algorithms were obtained.

2.2. A posteriori Markowitz model

The general idea of the a posteriori approach to solving MOO problems rephrased in terms of portfolios is: first, to compute the set of efficient (or non-dominated) portfolios and, then, to select a single portfolio from it. The selection of a final portfolio can be done by the decision maker or expert, e.g. with the help of multi-criteria decision aiding approaches, see e. g. [2]. Given two objective functions, in our case $\sigma^2(\mathbf{x}) = \mathbf{x}^\top \mathbf{Q}\mathbf{x}$ and $E(\mathbf{x}) = \mathbf{r}^\top x$, one can associate to each solution $\mathbf{x}$ a 2-dimensional evaluation vector in the objective space, $(\sigma^2(\mathbf{x}), E(\mathbf{x}))^\top$, where the risk objective is to be minimized and the return objective is to be maximized; $\mathbf{r}$ and $\mathbf{Q}$ are defined as before in (1).

A portfolio $\mathbf{x}^{(1)}$ dominates a portfolio $\mathbf{x}^{(2)}$ (in symbols $\mathbf{x}^{(1)} \prec \mathbf{x}^{(2)}$), if and only if $E(\mathbf{x}^{(1)}) \geq E(\mathbf{x}^{(2)})$ and ($\sigma^2(\mathbf{x}^{(1)}) < \sigma^2(\mathbf{x}^{(2)})$ or $E(\mathbf{x}^{(1)}) > E(\mathbf{x}^{(2)})$) and $\sigma^2(\mathbf{x}^{(1)}) \leq \sigma^2(\mathbf{x}^{(2)})$. The efficient set X_E (of portfolios) is given by the portfolios that are not dominated by any other portfolio. The image of this set is called the Pareto front PF, i. e.

$$PF = \{(y_1, y_2)^\top \in \mathbb{R}^2 \mid \exists \mathbf{x} \in X_E : y_1 = \sigma^2(\mathbf{x}) \text{ and } y_2 = E(\mathbf{x})\}.$$

An example of Pareto fronts of optimal portfolios can be seen in Figure 4. Note that we chose the first coordinate (y_1) for the risk objective (or variance), and the second coordinate (y_2) for the expected return objective, thereby following the convention in portfolio optimization.

It should be noted that here the formulation (1) is adapted from the continuous version to a discrete, in particular an *integer* one. In integer adaptation an asset is either taken or not at a fixed price. The search space of the problem $\mathcal{S}$ is $\{0, 1\}^{N_{Total}}$.

The Pareto front will be obtained at the upper left boundary of the set of attainable solutions $Y = \{(y_1, y_2)^\top \in \mathbb{R}^2 \mid \exists \mathbf{x} \in \{0, 1\}^{N_{Total}} : y_1 = \sigma^2(\mathbf{x}) \text{ and } y_2 = E(\mathbf{x})\}$. It can, for instance, be obtained by a series of constrained single objective optimization problems. Moreover, a fixed budget B is allocated for buying assets of a portfolio, which in research projects is lost if not spent. Hence, the *integer adaptation of the Markowitz model* is as

follows:

$$\begin{aligned} E(\mathbf{x}) &\rightarrow \max; \\ \sigma^2(\mathbf{x}) &\rightarrow \min; \\ \text{s.t.} \sum_{i=1}^{N_{Total}} c_i x_i = \mathbf{x}^\top \cdot \mathbf{c} &\leq B; \\ x_i &\in \{0,1\}, i = 1, \ldots, N_{Total}, \end{aligned} \tag{2}$$

where c_i refers to the cost of an asset, which can be different for different assets.

The set of feasible portfolios $\mathcal{F}$ is now the subset of portfolios in $\mathcal{S}$ with $\sum_{i=1}^{N_{Total}} c_i x_i = \mathbf{x}^\top \cdot \mathbf{c} \leq B$.

Earlier VS for drug discovery was formulated as a multiobjective optimization problem in [17], where both activity and diversity were maximized simultaneously. Our portfolio-based formulation is similar, however, different diversity measure based on Solow-Polasky diversity [23] is used, see section 3.2, and expected return based on activity is computed instead of activity score maximization.

2.3. A priori Sharpe ratio model

All portfolios belonging to the efficient set present tradeoffs between return and risk. Eventually however, from the set of efficient portfolios a single one should be chosen. Instead of letting the decision maker make a subjective decision by viewing solutions on the Pareto front (a posteriori decision

making) one could also establish beforehand a criterion by which the best solution on the Pareto front is selected (a priori decision making).

The investment management suggests a large number of measures to evaluate return-to-risk ratios of portfolios, relatively to time period (e. g., standard deviation), to market behavior (e. g., beta ratio), to benchmark asset (e. g., tracking error, excess return, Sharpe ratio). The Sharpe ratio, also called reward-to-volatility ratio, is the most widely used risk-adjusted performance index [5] and will be used here.

The Sharpe ratio can be defined with the help of the capital allocation line (CAL). It is a straight line on the return-risk graph (see Figure 1) that shows all possible combinations of risky portfolios with the risk-free asset $r_f \geq 0$. The risk-free asset, r_f, has a return that is smaller than the minimal expected return of an efficient portfolio $r_f < r_{min}$, and it assumes risk-free investment. The optimal CAL corresponds to the portfolios with lowest risk for any given value of return $r > r_f$. The slope of the optimal CAL is a sub-derivative of the function that defines the Pareto front of efficient portfolios. The point at which the CAL touches the front of efficient portfolios corresponds to the Sharpe ratio that provides an optimal risky portfolio.

Here, the risk free investment is chosen to be $r_f = -B$, as this will be the exact return if we do not invest in the research. Then the Sharpe ratio is defined as

$$Sh(\mathbf{x}) = \frac{E(\mathbf{x}) - r_f}{\sigma(\mathbf{x})}.$$

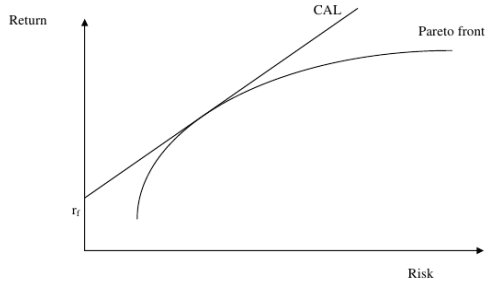

Figure 1: Sharpe ratio on intersection of CAL and Pareto front

The Sharpe ratio characterizes how well the return of a portfolio compensates the risk taken, and it measures excess of return per unit of risk. When comparing two portfolios, the one with the higher Sharpe ratio gives more return per risk. Finding the portfolio with maximal *Sharpe ratio* yields the following nonlinear integer programming problem:

$$\begin{aligned} \frac{E(\mathbf{x}) - r_f}{\sigma(\mathbf{x})} &\rightarrow \max; \\ \text{s.t. } \mathbf{x}^\top \cdot \mathbf{c} &\leq B; \\ x_i &\in \{0,1\}, i = 1, \ldots, N_{Total}, \end{aligned} \tag{3}$$

where B refers to the budget, which in research projects if not spent is lost.

2.4. Portfolios with fixed size

The problem with the Markowitz (2) and optimal Sharpe ratio (3) formulations is that they both favor selection of empty portfolios as they may be best at minimizing risk of any losses. One way to neutralize this effect is to require a fixed number of assets to be selected into the portfolio. This

problem formulation is referred to as *fixed size portfolio selection* and it assumes that the number of assets to be selected is limited to a specific number $N_{Portfolio}$. Then, in addition to the formulation (2) or (3), a constraint of the following form is assumed:

$$\mathbf{x}^\top \cdot \mathbf{e} = N_{Portfolio},$$

where $\mathbf{e}$ is in $\{0,1\}^{N_{Total}}$; each coordinate is either 0 or 1, summing up to portfolio of $N_{Portfolio}$ size (with $N_{Portfolio} << N_{Total}$ not all molecules being selected in portfolio out of N_{Total}).

This formulation is equivalent to the 0-1 quadratic knapsack problem. Problems of this form were intensively studied in the literature due to their simple and practical formulation, but there are difficulties in finding exact solutions for them (as indicated in [16] and [20]).

3. Drug subset selection as portfolio optimization

Several formulations from the previous section can be used for selecting portfolio of molecules that are potential drugs. For formulating such problems the following model variables are considered:

1. A fixed budget B is available and has to be spent. Money that is not used will be lost.
2. Each successful molecule is associated with a gain G, which is the value (expressed in monetary units) gained if the molecule becomes a drug.

The gain is the same for each successful molecule $G_i = G$ and is zero for unsuccessful molecules.

3. For each available molecule $i = 1, \ldots, N_{Total}$ a probability of success p_i is given or obtained a priori.
4. For each candidate molecule $i = 1, \ldots, N_{Total}$ the cost c_i for buying and testing it is known. The cost of different molecules may vary significantly, and this cost does not involve indirect costs, e. g. costs of the *in vitro* testing.
5. From a given set of N_{Total} candidates, a subset of $N_{Portfolio}$ molecules is selected such that
   (a) the budget B is not exceeded.
   (b) The expected return E is to be maximized, where the expected return is given by the expected value of the random variable of the return R of a portfolio of molecules selected for testing.
   (c) The risk σ associated with the expected return is to be minimized.

3.1. A posteriori Markowitz model with fixed size portfolio

The problem corresponding to a posteriori Markowitz model with limited budget and fixed size of portfolio constitutes a two-objective optimization problem that is formulated as follows:

$$
\begin{aligned}
E(\mathbf{x}) &\rightarrow \max; \\
\sigma^2(\mathbf{x}) &\rightarrow \min; \\
\text{s.t. } \mathbf{x}^\top \cdot \mathbf{c} &\leq B; \\
\mathbf{x}^\top \cdot \mathbf{e} &= N_{Portfolio}; \\
x_i &\in \{0,1\}, i = 1, \ldots, N_{Total},
\end{aligned}
\tag{4}
$$

and is referred in the text as the *Markowitz model with fixed size portfolios.*

Here, $x_i,\ i = 1, \ldots, N_{Total}$, denote the decision variables; $x_i = 1$ means that the i-th molecule is selected and $x_i = 0$ means that it is not selected; N_{Total} is the number of available molecules. The search space of the problem $\mathcal{S}$ is $\{0,1\}^{N_{Total}}$. The set of feasible portfolios $\mathcal{F}$ is now the subset of portfolios in $\mathcal{S}$ with $\sum_{i=1}^{N_{Total}} x_i = N_{Portfolio}$, where $N_{Portfolio}$ is the size of the portfolio. We consider two real valued objective functions defined on $\mathcal{S}$, $\sigma^2(\mathbf{x}) = \mathbf{x}^\top \mathbf{Q} \mathbf{x}$ and $E(\mathbf{x}) = \mathbf{r}^\top \mathbf{x}$. Each portfolio $\mathbf{x}$ is associated with a 2-dimensional evaluation vector in the objective space, $(\sigma^2(\mathbf{x}), E(\mathbf{x}))^T$, where the risk objective is to be minimized and the return objective is to be maximized; $\mathbf{r}$ and $\mathbf{Q}$ are defined as before in (1).

The computation of return $E(\mathbf{x})$ and risk $\sigma^2(\mathbf{x})$ is discussed next. The return $E(\mathbf{x})$ is defined as the gains minus the losses. For the expected return it is important to realize that money from the budget that is not invested in molecules is lost. Therefore, the losses will be B and the gains will be the

cumulated gains from molecules that become successful drugs. Hence,

$$E(\mathbf{x}) = G\,\mathbf{p}\cdot\mathbf{x} - B = \left(\sum_{i=1}^{N_{Total}} Gp_i x_i\right) - B.$$

Due to the probabilistic nature of the return (we get it only in case of successful drug(s)), it can be modeled as a random variable. Let $\tilde{x_i}$ denote a random variable of Bernoulli type that models the uncertain return on investment in a molecule i:

$$\tilde{x_i} = \begin{cases} \dfrac{G - c_i}{c_i}, \text{ return rate with probability } p_i \text{ in case of success;} \\ \dfrac{0 - c_i}{c_i} = -1, \text{ return rate with probability } 1 - p_i \text{ in case of no success.} \end{cases}$$

Then, the expected return of a molecule i is defined by:

$$E(\tilde{x_i}) = \frac{G - c_i}{c_i}\cdot p_i + \frac{-c_i}{c_i}\cdot(1 - p_i).$$

Following the classical model of Markowitz, the risk $\sigma^2(\mathbf{x})$ can be expressed by means of a covariance matrix $\mathbf{Q}$ as follows:

$$\sigma^2(\mathbf{x}) = \mathbf{x}^\top\mathbf{Q}\mathbf{x} = \sum_{i=1}^{N_{Total}}\sum_{j=1}^{N_{Total}} x_i q_{ij} x_j,$$

where q_{ij} is a correlation between the return from the i-th molecule r_i and the return from the j-th molecule r_j. The computation of the covariance on the basis of a distance matrix will be derived from the Solow-Polasky model

as discussed in the next section.

3.2. Solow-Polasky diversity measure

One possible interpretation of the covariance can be done using a measure for estimating diversity of a (biological) population introduced by Solow and Polasky (see [23]). Originally, they were searching for a measure that can be used for evaluating population diversity rigorously, assuming some particular properties for this measure are respected. A measure which counts essentially different species and is used in the context of species preservation. Within the utilitarian model of Solow and Polasky the more species is considered to be more useful because of e.g. their potential future medical benefits. In general there are other reasons for species preservation, e.g. for stability of eco-system or ethical reasons. But in our context utilitarian motivation for species preservation fits well.

Hence, they suggested a diversity function:

$$D(\mathbf{s}) = \mathbf{e}^\top F(\mathbf{s})^{-1}\mathbf{e},$$

where $\mathbf{e}$ is an N_{Total}-vector of 1's and $F(\mathbf{s})$ is a non-singular N_{Total}-by-N_{Total} distance matrix $F(\mathbf{s}) = [f(d(s_i, s_j))]$, with a distance function $f(d_{ij})$ taken for each pair of species $d(s_i, s_j)$. Each entry of the Solow-Polasky matrix indicates distance between species s_i and s_j, where $i = 1, \ldots, N_{Total}$ and $j = 1, \ldots, N_{Total}$.

When compared to other diversity measures, e.g. proposed in [28], Solow-

Polasky distance takes into account not only the distance between species in the population but also provides a measure for the number of different species in it. This model is inspired by probabilistic modeling of a set of species, but can be adapted to drug discovery.

Let $S = \{s_1, s_2, \ldots, s_{N_{Total}}\}$ be a set of molecules, $|S| = N_{Total}$. Let S' be any subset of S, then $B(S')$ denotes the composite event that at least one molecule in S' is successful. By $Pr(B(S'))$ we denote the probability of this composite event. The expected benefit of S' can be measured by the product $Pr(B(S')) \cdot V$, where V is a fixed unit value of benefit. Based on this benefit measure different subsets of S can be compared.

Knowing a priori information on the performance of different molecules with respect to the specified goal(s), the probability of their benefit can be defined as $Pr(B_i) = p_i$, where B_i denotes the event that the i-th molecule is successful. Otherwise, if probabilities are unknown, they may be considered as equal $Pr(B_i) = p$ for all B_i, $i = 1, \ldots, N_{Total}$. For the event B_j being successful, the conditional probability for the event B_i is defined in [23] as $Pr(B_i|B_j) = p+(1-p)f(d_{ij})$, where f is a function selected with the following properties: $f(0) = 1$, $f(\infty) = 0, f' \leq 0$. Here, as remarked by Solow and Polasky, f can be interpreted as a correlation function.

Finding the N_{Total}-variate distribution $Pr(B(S))$ from univariate and bivariate probabilities is not possible. However, the lower bound on it was defined in [10]. One example of the distance function for computing this distance matrix is provided in [23]: $f(d) = e^{-\theta d(s_i,s_j)}$, and it will be used

here.

3.3. A posteriori Markowitz model with Solow-Polasky diversity

In addition to difficulties with computing an exact solution for a fixed size portfolio, in some cases the tendency of selecting the cheapest solutions in the portfolio may be observed if enough diversity is reached at the cost of cheapest assets. Hence, relaxing the constraint on the number of assets in the portfolio may be beneficial.

Fortunately, Solow and Polasky specified a set of requirements which a biological diversity measure should satisfy; see [23] for more details. One of the requirements is *monotonicity in species*, which suggests that the diversity of a set increases with adding new elements to it and decreases with removing elements. This property is taken into account in the next portfolio optimization model.

Since minimizing risk of selecting similar assets into a portfolio can also be interpreted as maximizing diversity of selected portfolio of assets, different formulations of diversity can be taken in the portfolio selection problem. Here, we propose to use Solow-Polasky diversity as a second objective instead of the risk measure calculated as a variance of the returns.

The Solow-Polasky diversity measure is calculated as the sum of the entries of the inverse of the correlation matrix for selected assets:

$$D(\mathbf{x}) = \mathbf{e}^\top F(\mathbf{x})^{-1}\mathbf{e} = \sum_{i=1}^{N_{Total}} \sum_{j=1}^{N_{Total}} F(\mathbf{x})^{-1}_{ij},$$

where $F(\mathbf{x})_{ij}^{-1}$ is the inverse of the correlation matrix for all selected assets.

Then, the two objectives to be optimized are: the return and the diversity of the portfolio, which can be presented in the following model:

$$\begin{aligned} E(\mathbf{x}) &\rightarrow \max; \\ D(\mathbf{x}) &\rightarrow \max; \\ \text{s.t. } \mathbf{x}^{\top} \cdot \mathbf{c} &\leq B; \\ x_i &\in \{0,1\}, i = 1, \ldots, N_{Total}. \end{aligned} \tag{5}$$

Even though both a posteriori approaches use the correlation function suggested by Solow and Polasky, the Markowitz model minimizes the sum of the correlation matrix entries, while the Solow-Polasky diversity model maximizes the sum of the entries of the inverse of the correlation matrix. Hence, the former model favors smaller size portfolios, while the latter one gives preference to larger portfolios.

4. Solution algorithms

Different methods can be used to compute efficient portfolios to the given portfolio selection problem. In this section, the methods that proved to be robust solvers are presented. In general, the difficulty of finding efficient portfolios depends on the number of candidate molecules N_{Total} and the size of the subset that is selected $N_{Portfolio}$.

Portfolio optimization problems belong to the class of NP hard problems

and, under the P $\neq$ NP assumption, the effort needed to solve them exactly is growing exponentially with increasing N_{Total}. Portfolio optimization problems can be formulated either as discrete or continuous/parametric optimization problems. The former presentation is more common due to faster performance on small and medium size problems (with up to 500 assets) with interior-point optimizers. However, in [24], it was shown that for large-scale problems (in the range of 1,000 to 3,000 assets) continuous formulation may be computationally more efficient when solved with some optimizers. Recently, new exact solvers such as Gurobi (see [12]) show fast performance for large instances (at least with datasets with up to 5,000 assets considered in this work) with branch and bound method.

However, for finding the Pareto fronts in Markowitz models and computing Sharpe ratio, some adaptations to the formulations presented earlier need to be performed before applying exact solvers. This will be discussed next, first for the Pareto front computation in the Markowitz model with fixed size portfolio and then for the Sharpe ratio maximization. For the case of the Markowitz model with Solow-Polasky diversity optimization instead of the original risk objective, exact solvers cannot be applied due to the complexity of the risk objective function. But approximate algorithms, such as meta-heuristics and multiobjective evolutionary algorithms in particular, could and will be applied to find approximate solutions.

4.1. Markowitz model with fixed size portfolio computation using ϵ-constraint method

To find the Pareto front and efficient set of the problem, it is proposed to use the ϵ-constraint method, see e.g. [18]. This is done by formulating a series of single objective constrained optimization problems (SOCOPs) with moving constraint on one of the objective function values. Then one objective is optimized subject to the other objective fixed and expressed as a constraint. To obtain, say N_{Pareto}, points on the Pareto front, we solve the following series of N_{Pareto} SOCOPs for ascending expected returns E_j, $j = 1, \ldots, N_{Pareto}$:

$$\begin{aligned} \sigma^2(\mathbf{x}) &\rightarrow \min; \\ \text{s.t.} \quad E(\mathbf{x}) &\geq E_j; \\ \mathbf{x}^\top \cdot \mathbf{c} &\leq B; \\ \mathbf{x}^\top &= N_{Portfolio}; \\ x_i &\in \{0,1\}, i = 1, \ldots, N_{Total}. \end{aligned} \tag{6}$$

The resulting optima will be called $\mathbf{x}_j^*$, and their risk $\sigma_j^{2^*}$ and return E_j^* values. The values of E_j^* are taken evenly spaced between lower bound $E^{\min}$ and upper bound $E^{\max}$. The computation of the lower and upper bounds,

$E^{\min}$ and $E^{\max}$, is done by solving the SOCOPs, respectively:

$$\begin{aligned} E(\mathbf{x}) &\rightarrow \min; \\ \text{s.t. } \mathbf{x}^\top \cdot \mathbf{c} &\leq B; \\ x_i &\in \{0,1\}, i = 1, \ldots, N_{Total}, \end{aligned} \tag{7}$$

and

$$\begin{aligned} E(\mathbf{x}) &\rightarrow \max; \\ \text{s.t. } \mathbf{x}^\top \cdot \mathbf{c} &\leq B; \\ x_i &\in \{0,1\}, i = 1, \ldots, N_{Total}. \end{aligned} \tag{8}$$

Let $\mathbf{x}^{\min}$ denote the solution obtained for the first problem (7) and $\mathbf{x}^{\max}$ denote the solution obtained for the second problem (8). Then, the lower bound for the return is $E^{\min} = E(\mathbf{x}^{\min})$ and the upper bound for the return is $E^{\max} = E(\mathbf{x}^{\max})$.

4.2. Sharpe ratio with fixed size portfolio computation using quadratic programming

In order to maximize the Sharpe ratio, it would be beneficial to get rid of the nonlinear and non-quadratic term $\frac{E(\mathbf{x})-r_f}{\sigma(\mathbf{x})}$ in the problem formulation (3), and then use a quadratic solver. For this, homogenization has been suggested in [5]. However, our experience was that the resulting mixed integer quadratic programming (QP) problem was difficult to solve due to resulting covariance

matrix being not of a semidefinite type.

Alternatively, it is also possible to compute the Pareto front with (1) and find the point on the Pareto front that maximizes the Sharpe ratio computed with (3). Given a sufficiently dense approximation of the Pareto front this is accomplished by evaluating the Sharpe ratio of all points on the Pareto front i.e.:

$$\mathbf{x}^{sharpe} = \arg\max\{Sh(\mathbf{x}_1), \ldots, Sh(\mathbf{x}_{N_{Pareto}})\}.$$

It is important in this context that points that maximize the Sharpe ratio are part of the efficient set.

4.3. Markowitz model with Solow-Polasky diversity computation using multi-objective genetic algorithms

In case of Markowitz model with Solow-Polasky diversity considered as a risk objective (5), the need of obtaining the inverse of the distance matrix makes the application of quadratic programming difficult. An alternative approach is to use approximate methods, for instance, meta-heuristics. While meta-heuristics do not guarantee reaching an optimal solution, they can typically obtain good approximations to optima fairly quickly even for NP hard combinatorial problems, which is the case of knapsack / portfolio optimization problems considered in this work.

Among many meta-heuristics developed so far, multiobjective evolutionary algorithms (MOEAs) are particularly common for solving multi-objective optimization problems. In this study, two common MOEAs are considered:

NSGA-II (see [6]) and SMS-EMOA (see [7]). Using otherwise standard implementations of these meta-heuristic solvers, we introduce two problem specific adaptations. These are the mutation and the recombination operators, which were specifically designed for the subset selection problem.

MOEAs maintain a population (multi-set) of individuals that is changing over time due to the application of variation and selection operators. From a given population $P(t)$ at time t pairs of parents are selected – in the so-called mating selection step – and offspring are then generated by recombination and mutation based on these parents. Then from the offspring and the individuals of previous population $P(t)$ a set of individuals is selected – in the so-called environmental selection step – that forms the next population $P(t+1)$. While the two selection steps are based on choosing individuals with the best objective function values, the two variation steps – recombination and mutation – seek to generate new individuals that resemble some of the traits of their parents. Recombination combines the information of parents, and mutation does a small random modification of a solution.

The NSGA-II and SMS-EMOA algorithms differ in their selection steps: In NSGA-II, a new offspring population of the same size as the population $P(t)$ is generated, and, subsequently, the new population $P(t+1)$ is selected based on so-called non-dominated sorting and crowding-distance. In SMS-EMOA, only one offspring is generated based on $P(t)$ and the next population $P(t+1)$ is obtained by non-dominated sorting and selecting the subset that maximizes the hypervolume indicator. Here, the hypervolume in-

dicator, which the SMS-EMOA seeks to maximize, is a measure computed to show how well a population serves to mark the boundary between the dominated and non-dominated spaces, and, thus, how well it serves to represent the true Pareto front. The MOEA for the portfolio subset selection problems represents individuals (that are portfolios) as *sorted index lists*. For instance, the sequence $(1, 4, 6, 29)$ represents the portfolio that selects the 1st, the 4th, the 6th and the 29th molecules.

The mutation is done by (1) deleting a single randomly chosen molecule from the portfolio, (2) adding a randomly chosen new molecule, and (3) replacing a molecule inside the portfolio by a molecule outside the portfolio. Each of these mutation operators is applied with a certain probability for each molecule, which is denoted by p_{MD}, p_{MA}, and respectively p_{MR}. In case of a fixed number of molecules in the portfolio, only replacement is used. While p_{MD} and p_{MA} determine probability of adding and deleting a single molecule *per portfolio*, the replacement probability p_{MR} is defined *per molecule in the portfolio*.

As a recombination operator m-point crossover is applied. This means we randomly select m points for the number of molecules. After each point we change the parent we use to copy from. To make it applicable for subsets, the subset membership is interpreted as a bit-string (one means a molecule is a member of portfolio, zero means a molecule is not a member of portfolio), and the crossover determines membership based on either one of the two parents selected randomly. The probability of crossover is p_{CO}. If crossover

is not applied, then one of the two parents chosen randomly is copied and will serve as offspring (before mutation).

5. Experimental results

5.1. Molecular portfolio selection model assumptions

First, the information on a covariance matrix Q needs to be formulated. In chemistry, the distance between molecules can be defined by evaluating similarities/differences in the structure of two molecules. Being able to measure the distance $d(x_i, x_j)$ between each pair of molecules i and j, provides means for defining the matrix $F(\mathbf{x})$ of N_{Total}-by-N_{Total} size, e.g. as suggested in [23] with elements $f_{ij} = e^{-\theta d(x_i,x_j)}$, where θ is set to $\theta = 0.5$. The distance between molecules can be computed based on their similarity, e.g. according to the Tanimoto similarity, see [25] also used in [22].

Tanimoto similarity Sim_T is a measure of similarity between two bit vectors A and B. The bit-vectors used here are the molecular fingerprints. A molecular fingerprint is a bit vector, where each bit represents whether a chemical substructure is part of the molecule (1) or not (0). The Tanimoto similarity can be defined as:

$$Sim_T(A, B) = \frac{\sum_z A_z \wedge B_z}{\sum_z A_z \vee B_z},$$

where the index z corresponds to a particular property of molecule structure.

In this study, circular fingerprints ($FCFP_4$) calculated with Pipeline

Pilot 9.0.2.1 were used [21]. To predict activity of molecules, we used a Proteochemometric model as published by van Westen et al. in [26]. The molecules selected here originated from the Enamine building blocks [8] with prices defined per 100mg.

Then, we can calculate the distance between two molecules as a dissimilarity measure, which is diversity:

$$d(x_i, x_j) = 1 - Sim_T, \tag{9}$$

where Sim_T is Tanimoto similarity.

Second, the information on (bio-)activities of the candidate molecules needs to be translated into success probabilities. Activity a_i is normally given as logarithmized activity l_i; in this case, we can use $a_i = e^{l_i}$.

Moreover, from experience chemists know an average probability of success $\bar{p}$, for the sake of the argument estimated as $\bar{p} = 1/100$. Let us consider a vector of N_{Total} activities (exponentiated) $\mathcal{A} = \{a_1, \ldots, a_{N_{Total}}\}$ and let $\mathcal{P} = \{p_1, \ldots, p_{N_{Total}}\}$ denote the success probabilities. Then, the average probability of success can be calculated as:

$$\bar{p} = \frac{1}{N_{Total}} * \sum_{i=1}^{N_{Total}} p_i, \tag{10}$$

and we know that activities are proportional to success probability. Hence,

for some constant k it holds:

$$p_i = k * a_i, \ \forall i = 1, \ldots, N_{Total}. \tag{11}$$

By substituting p_i in (10) as defined in (11), we can obtain k:

$$k = \bar{p} * \frac{N_{Total}}{\sum_{i=1}^{N_{Total}} a_i}. \tag{12}$$

Combining 11 and 12 we get:

$$p_i = a_i * \bar{p} * \frac{N_{Total}}{\sum_{i=1}^{N_{Total}} a_i}. \tag{13}$$

Third, the gain from a new lead compound (i.e. a molecule that may lead to a new drug) may vary between, e.g. $G_L = 10,000$ and $G_U = 100,000$ USD.

Fourth, several findings for the current drug portfolio selection model are based on the analysis of these model assumptions. Since the return of each molecule, which is equal to the product of gain and probability of success (for G_L $r_i = 10,000 * 0.0001 = 1$ or for G_U $r_i = 100,000 * 0.0001 = 10$), is very small, it turns out that it is not profitable to invest into molecules in the early stages of drug discovery. Besides having economical profitability, it is often the case that a budget for drug discovery is made available in research projects to stimulate medical innovation.

Fifth, for the fixed size portfolio model, we assume that 100 molecules

need to be selected out of each dataset into the portfolio of molecules to be tested *in vitro*: $N_{Portfolio} = 100$.

Sixth, the budget to be spent and to be taken into account as a constant in the model is calculated assuming a fixed number of molecules $N_{Portfolio} = 100$ will be bought. Hence, the budget can be obtained as an average cost multiplied by the number of molecules to be bought: $B = N_{Portfolio} * \sum_{i=1}^{N_{Total}} c_i / N_{Total}$.

Seventh, the budget is set to a hundred times the average cost of molecules in the dataset: $B = 100 * \bar{c}$. For the dataset of 1000 molecules this yields $B = 34,502$USD, for the dataset of 2500 molecules this yields $B = 34,400$USD and for the dataset of 5000 molecules this results in $B = 34,622$USD.

Eighth, based on comparison of performance of the algorithms on all three datasets, it was observed that larger datasets perform better, when compared to smaller datasets, assuming the same fixed number of 100 molecules is selected from all three datasets. One could argue that this may be the result of applying more iterations in the bigger datasets. However, this is not the case as all datasets converge after $100,000$ iterations, which means that running the algorithms for more iterations will not be effective.

The reason for this behavior is the way success probabilities of molecules are computed: The success probability of a molecule is calculated inversely proportional to average activity of all molecules belonging to the dataset. It would be a correct approach if the datasets would be uniformly selected from the vendor database. However, in our case the datasets were sorted by

activity before selection and from the same sorted dataset top 1000, 2500 and 5000 most active molecules were selected. The datasets are sorted based on activity because chemists usually do not consider molecules below the cut-off activity. Thus, larger datasets, e.g. with 2500 and 5000 molecules, contain molecules with lower activity on average, when compared to smaller datasets with higher average activity, e.g. with 1000 molecules.

To avoid the situation when the average success probability of a given molecule is lower in a bigger dataset when compared to the average success probability of the same molecule in a smaller dataset, its calculation is adjusted. In particular, the average probability of success of a molecule is computed in such a way that it is independent of the size of the considered dataset, and in such a way it suits better to datasets with a non-uniform distribution of activities.

As before it is assumed that success probability is proportional to the activity. However, now the average probability $\bar{p}_{1000}$ is fixed to be proportional to the average activity $\bar{a}_{1000}$ of the 1000 molecules dataset:

$$\bar{p}_{1000} = k * \bar{a}_{1000}, \ \forall i = \{1, \ldots, N_{Total}\}. \tag{14}$$

which leads to the k computed as:

$$k_{1000} = \bar{p}_{1000} * \frac{1}{\bar{a}_{1000}}, \tag{15}$$

Thus, the probability of success of each molecule can be computed as:

$$p_i = a_i * \bar{p}_{1000} * \frac{1}{\bar{a}_{1000}}. \tag{16}$$

Hence, this fixed average probability $\bar{p}_{1000}$ of the 1000 molecules dataset will be used for computing the probabilities of success of molecules p_i in the datasets with 2500 and 5000 molecules.

5.2. Molecular compounds datasets

For testing efficiency of the proposed models for molecule subset selection, we have used 3 datasets of 1000, 2500 and 5000 molecules taken from the ZINC database of molecular compounds (see [13]), as available at vendor Enamine. Each molecule was provided with its known structure and its cost per 100mg. The Tanimoto similarity was calculated for each pair of molecules.

These three datasets are demonstrated in Figure 2 (a) with activity and cost of the 1000 molecules set depicted in (dark) green, the 2500 molecules set depicted in green and (light) pink, and the 5000 molecules set depicted in green, pink and blue.

5.3. Experimental settings for MOEAs

In the experiments of this study, the following settings are used: $p_{MA} = 0.5$ (per portfolio), $p_{MD} = 0.1$ (per portfolio), and $p_{MR} = 0.01$ (per molecule). The number of crossover points was set to 1, and the probability of crossover

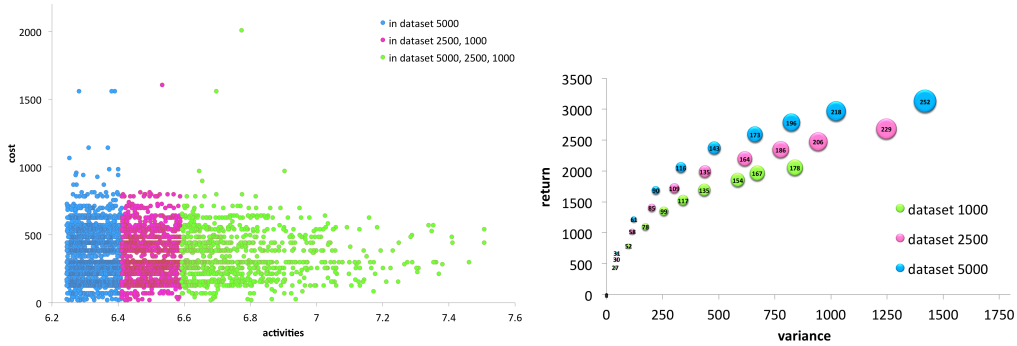

(a) (b)

Figure 2: (a) Cost and activity of molecules in three data sets. (b) Size of the population of portfolios in a single run (depicted in a circle) of SMS-EMOA for three data sets

was set to $p_{CO} = 0.2$. In other words, for every 10 offspring there are 2 that have been created using 2 parents, while the other 8 offspring are copies of some parents. Replacing a molecule in the portfolio with $p_{MR} = 0.01$ means replacing one molecule per offspring on average. A molecule is added to half of the offspring on average: $p_{MA} = 0.5$, and is removed from a portfolio once per 10 offspring on average: $p_{MD} = 0.1$. The size of the population of portfolios $P(t), t = 1, 2, \ldots$ was set to 10 in order to conform with the setting we used to sample the Pareto front by means of quadratic programming.

For fair comparison of MOEAs, in all experiments we run NSGA-II for $10,000$ iterations and SMS-EMOA for $100,000$ for the dataset of 1000 molecules. This is due to the fact that SMS-EMOA creates 1 offspring at each iteration, whereas NSGA-II creates 100 offspring at each iteration. For the larger datasets, we increased the number of iterations with the same factor as the dataset size. That is, for the 2500 molecules dataset we ran NSGA-

II for 25,000 iterations and SMS-EMOA for 250,000 iterations, whereas for the 5000 molecules dataset, we ran NSGAII for 50,000 iterations and SMS-EMOA for 500,000 iterations.

Due to the design of mutation operator used in this work, which allows not only replacing molecules in portfolio, but also adding or removing portfolios in the population, the population size varies. Figure 2 (b) gives insight into the cardinality of the sets of portfolios in the population obtained after a typical run of MOEAs (SMS-EMOA in this case, but similar results were obtained for NSGA-II).

A problem with the model (2) that minimizes risk without a cardinality constraint can be observed. In particular, this model allows selection of very small subsets of portfolios, and even the empty set of portfolios, as a part of the optimal front. Given the model this makes sense as there is no subset of portfolios with a higher return other than the one with a variance of 0USD. However, in practice this is undesirable.

5.4. Discussion of the experimental results

We tested the portfolio selection problem models formulated in sections 2 and 3 with the algorithms presented in section 4 on all three datasets discussed above.

All experiments were performed on a desktop PC with an i5 core 3.2 GHz processor and 4 GB memory under Windows XP operating system. Gurobi MIP solver version 4.0 was used and MOEAs were encoded in Python version

3.3.

5.4.1. Markowitz model with fixed portfolio size

Gurobi MIP results In the first experiment we computed Pareto front of portfolios optimal from the point of view of their return and risk according to the Markowitz model with fixed size portfolios ($N_{Portfolio} = 100$) using the formulation (6) as discussed in section 2.2. Here it is assumed that probability of success for each molecule is proportional to its activity and is computed by (11), and covariance between molecules is computed based on a distance as defined in (9).

The ϵ-constraint approach to MOO was used. In particular, the return objective was set to a constraint (computed for 15 different points between lower and upper return bounds, $E^{\min}$ and $E^{\max}$, respectively) and Gurobi MIP solver utilizing branch and bound method was applied to the three datasets with 1000, 2500 and 5000 molecules. The results of runs for all three datasets are presented in Figures 4 (a), (b) and (c), respectively, in black color.

Next, we analyze the content of portfolios belonging to the Pareto front of optimal portfolios with 100 molecules selected in each portfolio using the dataset with 2500 molecules as an example. In particular, we show four heat-maps indicating the similarity of selected molecules for portfolios with three different return values equal to 0USD, 1104USD and 1449USD and one randomly selected portfolio of 100 molecules demonstrated in Figure 3

(a), (b), (c) and (d), respectively. (Random selection was performed using a random percent filter in Pipeline Pilot using seed 333.) The darker the color the more similar the molecules are: the blue color gradient corresponds to a similarity equal to 1, dodger-blue to a similarity equal to 0.5, and white to a similarity equal to 0.

When compared to the baseline portfolio with 100 randomly selected molecules depicted in Figure 3 (d), the portfolio with 0USD return value depicted in Figure 3 (a) looks much more diverse, the portfolio with 1104USD return value depicted in Figure 3 (b) is slightly more diverse, and the portfolio with 1449USD return value depicted in Figure 3 (c) is much less diverse. The portfolio with 1104USD return value shown in Figure 3 (b) shows better diversity when compared to the baseline and relatively high return portfolios, being either close to or exactly the portfolio with optimal Sharpe ratio. This output is in line with the portfolio selection theory, according to which higher return portfolios are less diverse, since they also have higher risk, and the lower return portfolios are more diverse and have lower risk.

MOEAs results Comparison of MOEAs results is not trivial: On the one hand, due to the randomness of the population initialization and of the application of the crossover and mutation operators for individuals of the population, not a single run, but some averaged performance of MOEAs' several runs should be compared for evaluating performance of each MOEA. On the other hand, comparison of the convergence of each algorithm is difficult due to the fact that no true Pareto front is known. Therefore, only

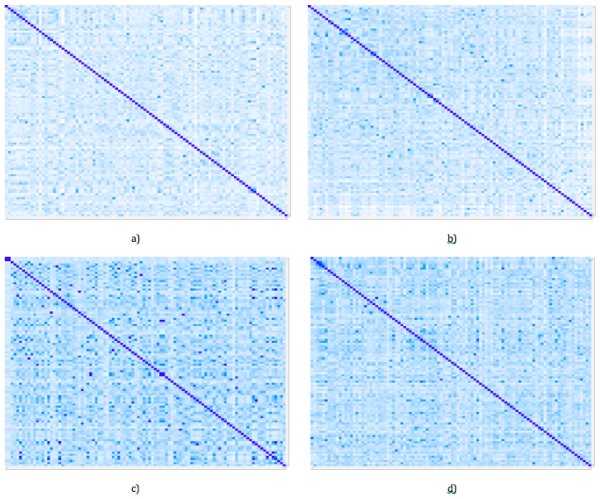

Figure 3: Similarity of 100 molecules portfolios belonging to the Pareto front and selected from the set of 2500 molecules with (a) 0USD return, (b) 1104USD return, (c) 1449USD return and (d) portfolio with 100 randomly selected molecules.

a visual comparison can be made based on the attainment surfaces [9] (or attainment curves for bi-objective optimization, which is our case) covered by each MOEA. This approach allows the comparison of lowest and highest Pareto front solutions achieved by each algorithm as well as their average performances. For computing the attainment surface a generalization of the median as an average is used, which is robust against outliers. In the next experiments only best front is taken from all runs of an algorithm for comparing to other algorithms performance.

Comparison of Gurobi MIP solver and MOEAs result: We compared results obtained by the Gurobi MIP solver using branch and bound method to the results obtained by two MOEAs, NSGA-II and SMS-EMOA. To make comparison fair we run all algorithms for circa 10 minutes for the

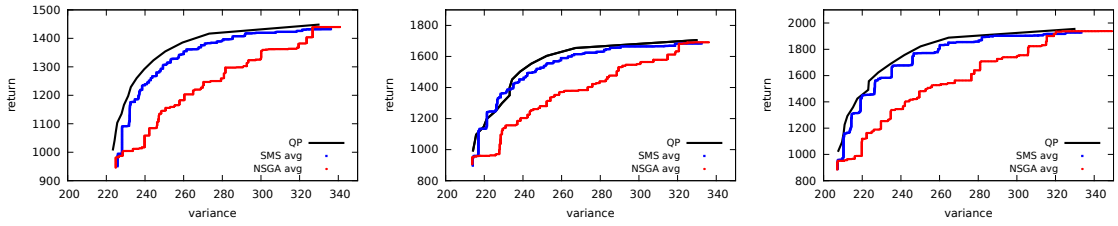

Figure 4: Comparison of the performance of the Gurobi MIP solver and the two MOEAs: NSGA-II and SMS-EMOA (on the plot denoted as QP, NSGA avg and SMS avg, respectively) for the 1000, 2500, 5000 molecules dataset, (a) (b) and (c), respectively.

dataset with 1000 molecules, for circa 20 minutes for the dataset with 2500 molecules, and for circa 30 minutes for the dataset with 5000 molecules. The results of this comparison presented in Figures 4 (a), (b) and (c) show the best performance of the exact Gurobi MIP solver for the datasets with 1000 and 5000 molecules, and better performance of the SMS-EMOA when compared to NSGA-II on all three datasets. As can be seen from Figure 4 (b) in some concave regions of the Pareto front SMS-EMOA outperformed Gurobi MIP solver, which means that specified time limit was not sufficient for branch and bound method of Gurobi MIP solver to find optimal solution.

5.4.2. Sharpe ratio with fixed portfolio size

We now show and discuss the results obtained for model (3). Figure 5 (a) demonstrates values of Sharpe ratio computed for 100 molecules selected from the 1000-molecule dataset at each of the 15th iterations of the ϵ-constraint method. These portfolios belong to the Pareto front of optimal portfolios and are obtained with the Gurobi MIP solver. In this case the portfolio

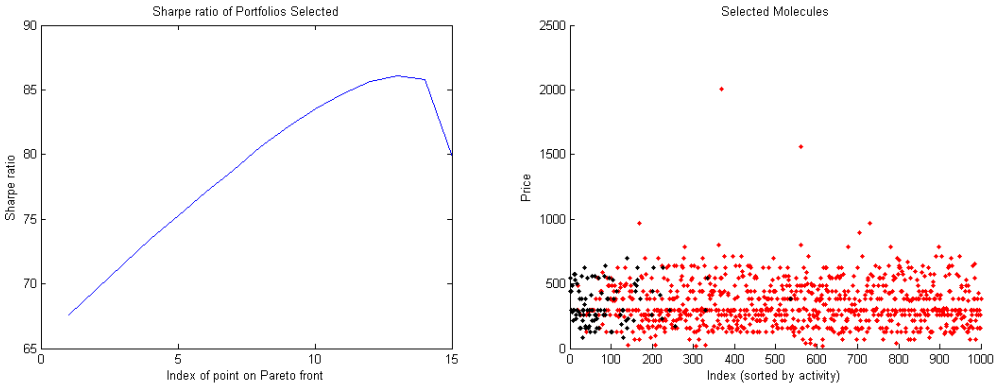

(a) (b)

Figure 5: (a) Sharpe ratios of 15 portfolios of 100 molecules belonging to the Pareto front and selected from the set of 1000 molecules. (b) Prices and activities of the 1000 dataset molecules (in red) and of the Sharpe optimal portfolio molecules (in black).

obtained at the 13th iteration has the highest Sharpe ratio value and should be selected as the most promising one for potential drug discovery. Next, we will analyze the content of this portfolio.

In Figure 5 (b), the molecules of the 1000 dataset are presented. Here, the molecules are allocated according to their activity (see X-axis) and price (see Y-axis), respectively. The molecules selected in the Sharpe optimal portfolio are marked in red and the non-selected molecules are depicted in black. As can be observed from this figure, not only the cheapest molecules are selected and not only the most active ones, but some balance between price and activity is reached for the portfolio of molecules as a whole.

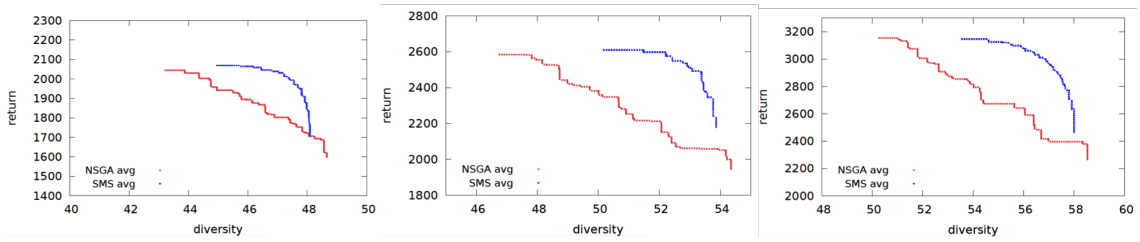

Figure 6: Comparison of NSGA-II and SMS-EMOA for Solow-Polasky diversity model with 1000, 2500 and 5000 molecules dataset, (a), (b), and (c), respectively.

5.4.3. Markowitz model with Solow-Polasky diversity

We now show and discuss the results obtained for model (5). Note that these results are for MOEAs only, as application of the MIP solver is complicated due to the need of obtaining the inverse of distance matrix.

Comparison of Pareto fronts obtained by NSGA-II and SMS-EMOA for Solow-Polasky diversity model provided in Figures 6 (a), (b), and (c) for datasets with 1000, 2500 and 5000 molecules, respectively, show outperformance of SMS-EMOA when compared to NSGA-II.

The formulation of the Solow Polasky diversity measure includes the inversion of a matrix making it difficult to optimize this measure by means of an exact solver, unlike the Sharpe ratio maximization formulation, which can be solved by quadratic programming. However, it might be possible to construct an approximation algorithm with an exact error bound for computing Solow Polasky diversity measure. Based on numerical experiments we conjecture that the Solow Polasky diversity is a submodular set function. If this is true, a greedy subset selection heuristic would yield an approximation

with approximation ratio $(1 - 1/e)$. We were not able to provide a formal proof for submodularity and leave this question to the future work.

6. Conclusion and future research

In this work, we presented a new approach to formulating the selection of molecules for de-novo drug discovery. In particular, the well-known in finance portfolio-based approach was used to model molecular subset selection for drug discovery as a portfolio selection. In addition to taking into account (bio-)activity of the molecules selected in the portfolio, the model considers the diversity of such portfolio. Moreover, it respects the limited budget provided for buying molecules and the fixed size of the portfolio as constraints. Molecules selected in the portfolio are balanced in terms of their price, expected individual performance and diversity.

Three models were proposed and tested on three molecular compounds datasets, in particular, classical Markowitz portfolio selection model, Sharpe ratio optimization and diversity optimization models. For solving Markowitz model with fixed size portfolio that optimizes return and risk simultaneously, we used ϵ-constrained approach in combination with Gurobi MIP solver and applied approximation approaches, in particular, multiobjective evolutionary algorithms, NSGA-II and SMS-EMOA. As expected QP solver was most efficient in calculating Pareto fronts except for some parts, which is due to a QP's fixed exploration time threshold. SMS-EMOA outperformed NSGA-II for Markowitz portfolio selection model. For the single objective Sharpe ratio

maximization model we adjusted Gurobi MIP solver and analyzed content of the selected optimal portfolios. Finally, for solving diversity optimization model only approximate algorithms, NSGA-II and SMS-EMOA, were used, with SMS-EMOA performing better than NSGA-II on all datasets. Solving this model with a quadratic solver requires obtaining inverse of distance matrix, which is difficult in practice as initial research shown. The presented preliminary test results of these novel formulations obtained for three molecular compounds datasets look promising and encourage us to do future research.

We have also discerned a number of future research topics that could be investigated further. In particular, different formulations of risk could be tested. For instance, other popular risk measures, such as Value-at-Risk (or return-to-standard deviation index) and the diversity inversely proportional to the number of species in the population can be investigated further. It would also be interesting to construct Sharpe ratio as a tangential point to the Pareto front (with CAL) and directly compute the Sharpe ratio optimum via homogenization. As the initial trials indicated the later approach is really challenging, but it might turn out to be easier for alternative risk formulations. Furthermore, alternative diversity measures, e.g. Weitzman diversity [28], can be considered in optimization models. A sensitivity analysis for some parameters of the models (e.g., theta parameter in Solow-Polasky diversity measure) will be of value for the proposed portfolio approach.

Current results show that application of existing QP solvers to large size

problems (bigger than 5000 molecules) is difficult due to large run times. To improve exact solvers' performance for large models, relaxation of integrity constraints can be applied. This would lead to rounding-off running time to polynomial, but will require covariance matrix to be positive definite. Alternatively, a MOEA could be used for preselection and QP solver for the final portfolio selection. It should be noted, however, that it takes approximately 10 minutes for SMS-EMOA to find Pareto front of portfolios for a dataset of 1000 molecules. Hence, in this case either parallelization or fast heuristic filters can be used for preselection as well.

An important task for future work is not only to scale up the models proposed in this work for larger portfolios, but also to further investigate the availability of exact solvers and performance for smaller portfolios. Moreover, experience with actual performance of the models in drug discovery practice needs to be assessed by comparing data of outcomes of a larger number of *in vitro* drug discovery studies with what has been predicted by the models.

In this work, only structural similarity of molecules was taken into account. Recent research [19] has shown that biological similarity plays an important roles in comparison of molecules. Similarly to structural diversity, biological diversity can be maximized as a third objective in the last proposed model. Two other models can also be adjusted to take into account biological similarity in the risk calculation. Moreover, at the later stages of drug discovery process additional objectives can be considered for molecular portfolio selection, such as minimizing side effects of the discovered lead can-

didates. The experimental validation of the discovered molecular portfolio via *in vitro* testing and chemists feedback on the results of such testing will be a natural next stage for the proposed in this work molecular portfolio selection approach.

7. Acknowledgements

Adriaan P. IJzerman and Eelke B. Lenselink thank the Dutch Research Council (NWO) for financial support (NWO-TOP #714.011.001).

References

[1] Allmendinger, R., Simaria, A., Farid, S.. Multiobjective evolutionary optimization in antibody purification process design. Biochemical Engineering Journal 2014;91:250–264.

[2] Belton, V., Stewart, T.. Multiple Criteria Decision Analysis: An Integrated Approach. Dordrecht, Netherlands: Kluwer Academic Publishers, 2001.

[3] Bonham, S.S.. IT Project Portfolio Management. Norwood: Artech House Inc., 2005.

[4] Brown, N.. Chemoinformatics – An introduction for computer scientists. ACM Computing Surveys 2009;41(2):8:1–8:38.

[5] Cornuejols, G., Tutuncu, R.. Optimization Methods in Finance. Cambridge, UK: Cambridge University Press, 2007.

[6] Deb, K., Pratap, A., Agarwal, S., Meyarivan, T.. A fast and elitist multiobjective genetic algorithm: NSGA-II. IEEE Transactions on Evolutionary Computation 2002;6(2):182–197.

[7] Emmerich, M.T.M., Beume, N., B., N.. An EMO algorithm using the hypervolume measure as selection criterion. In: Coello Coello, C., Hernández Aguirre, A., Zitzler, E., editors. Proceedings of the

3d International Conference on Evolutionary Multi-Criterion Optimization (EMO 2005). Berlin Heidelberg, Germany: Springer-Verlag; volume 3410; 2005. p. 62–76.

[8] Enamine Ltd., . Enamine screening compounds and building blocks. 2014. URL: www.enamine.net.

[9] Fonseca, C.M., Grunert da Fonseca, V., Paquete, L.. Exploring the performance of stochastic multiobjective optimisers with the second-order attainment function. In: Coello Coello, C., Hernández Aguirre, A., Zitzler, E., editors. Proceedings of the 3d International Conference on Evolutionary Multi-Criterion Optimization (EMO 2005). Berlin Heidelberg, Germany: Springer-Verlag; volume 3410; 2005. p. 250–264.

[10] Gallot, S.. A bound for the maximum of a number of random variables. Journal of Applied Probability 1966;3(2):556–558.

[11] Ghasemzadeh, F., Archer, N.. Project portfolio selection through decision support. Decision Support Systems 2000;29(1):73–88.

[12] Gurobi Optimization Inc., . Gurobi optimizer reference manual. 2014. URL: http://www.gurobi.com.

[13] Irwin, J.J., Shoichet, B.K.. ZINC - A free database of commercially available compounds for virtual screening. Journal of Chemical Information and Modeling 2005;45(1):177–182.

[14] Kirkwood, C.W.. Strategic Decision Making: Multiobjective Decision Analysis with Spreadsheets. Belmont, USA: Duxbury Press, 1997.

[15] Markowitz, H.. Portfolio selection. Journal of Finance 1952;7(1):77–91.

[16] Martello, S., Toth, P.. Knapsack Problems: Algorithms and Computer Implementations. Chichester, UK: John Wiley & Sons Ltd., 1990.

[17] Meinl, T., Ostermann, C., Berthold, M.R.. Maximum-score diversity selection for early drug discovery. Journal of Chemical Information and Modeling 2011;51(2):237–247.

[18] Miettinen, K.. Nonlinear Multiobjective Optimization. Dordrecht, Netherlands: Kluwer Academic Publishers, 1999.

[19] Paricharak, S., IJzerman, A.P., Bender, A., Nigsch, F.. Analysis of iterative screening with stepwise compound selection based on Novartis in-house HTS data. ACS Chemical Biology 2016;11(5):1255–1264.

[20] Pisinger, D.. The quadratic knapsack problem - A survey. Discrete Applied Mathematics 2007;155(5):623–648.

[21] Rogers, D., Hahn, M.. Extended-connectivity fingerprints. Journal of Chemical Information and Modeling 2010;50(5):742–754.

[22] Rogers, D.J., Tanimoto, T.T.. A computer program for classifying plants. Science 1960;132(3434):1115–1118.

[23] Solow, A.R., Polasky, S.. Measuring biological diversity. Environmental and Ecological Statistics 1994;1(2):95–103.

[24] Steuer, R.E., Qi, Y., Hirschberger, M.. Comparative issues in large-scale mean-variance efficient frontier computation. Decision Support Systems 2011;51(2):250–255.

[25] Tanimoto, T.T.. An Elementary Mathematical theory of Classification and Prediction. Technical Report 8; IBM Internal Report; 1958.

[26] van Westen, G.J.P., van den Hoven, O.O., van der Pijl, R., Mulder-Krieger, T., de Vries, H., Wegner, J.K., IJzerman, A.P., van Vlijmen, H.W.T., Bender, A.. Identifying novel adenosine receptor ligands by simultaneous proteochemometric modeling of rat and human bioactivity data. Journal of Medicinal Chemistry 2012;55(16):7010–7020.

[27] Walters, W.P., Stahl, M.T., Murcko, M.A.. Virtual screening – An overview. Drug Discovery Today 1998;3(4):160–178.

[28] Weitzman, M.L.. On diversity. The Quarterly Journal of Economics 1992;107(2):363–405.

[29] Yevseyeva, I., Guerreiro, A.P., Emmerich, M.T.M., Fonseca, C.M.. A portfolio optimization approach to selection in multiobjective evolutionary algorithms. In: Bartz-Beielstein, T., Branke, J., Filipič, B., Smith, J., editors. Proceedings of the 13th International Conference on Parallel Problem Solving from Nature - PPSN XIII (PPSN 2014)). Switzerland: Springer International Publishing; volume 8672; 2014. p. 672–681.