A Portfolio Optimization Approach to Selection in Multiobjective Evolutionary Algorithms

Iryna Yevseyeva¹, Andreia P. Guerreiro², Michael T. M. Emmerich³, and Carlos M. Fonseca²

 ¹ Centre for Cybercrime and Computer Security School of Computing Science, Newcastle University Newcastle upon Tyne, NE1 7RU, UK iryna.yevseyeva@ncl.ac.uk
 ² CISUC, Department of Informatics Engineering, University of Coimbra, Pólo II, Pinhal de Marrocos, 3030-290 Coimbra, Portugal {apg,cmfonsec}@dei.uc.pt
 ³ Leiden Institute of Advanced Computer Science, Leiden University, Niels Bohrweg 1, 2333 CA, Leiden, The Netherlands emmerich@liacs.nl

Abstract. In this work, a new approach to selection in multiobjective evolutionary algorithms (MOEAs) is proposed. It is based on the portfolio selection problem, which is well known in financial management. The idea of optimizing a portfolio of investments according to both expected return and risk is transferred to evolutionary selection, and fitness assignment is reinterpreted as the allocation of capital to the individuals in the population, while taking into account both individual quality and population diversity. The resulting selection procedure, which unifies parental and environmental selection, is instantiated by defining a suitable notion of (random) return for multiobjective optimization. Preliminary experiments on multiobjective multidimensional knapsack problem instances show that such a procedure is able to preserve diversity while promoting convergence towards the Pareto-optimal front.

Keywords: Fitness assignment, portfolio selection, Sharpe ratio, evolutionary algorithms, multiobjective knapsack problem.

1 Introduction

In evolutionary algorithms (EAs), selection shapes the direction in which the search is performed by dictating which individuals are allowed to reproduce. Typically, better individuals are assigned higher fitness, and are, therefore, selected for breeding. Carrying out selection based exclusively on individual performance (e.g., proportionally to a global objective value or individual rank) may work well when a single best solution is sought, but it typically leads to an undesired loss of population diversity when searching for multiple optimal solutions to multimodal problems. For this reason, techniques such as crowding and fitness sharing [17] are used in several multiobjective EAs (MOEAs) [6,7],

T. Bartz-Beielstein et al. (Eds.): PPSN XIII 2014, LNCS 8672, pp. 672–681, 2014.

[©] Springer International Publishing Switzerland 2014

to promote a good coverage of the Pareto-optimal front. Individuals located in more crowded regions are penalized, and greater chance of reproduction is given to individuals in less crowded regions.

Another approach is based on the use of quality indicators [20]. It consists of defining a notion of population quality and then inferring how much each individual contributes to the quality of the population. Algorithms such as LAHC [12], SMS-EMOA [4] and HyPE [1] instantiate this idea by using the hypervolume indicator as the measure of population quality, and clearly define what the contribution of each individual to that value is, albeit in different ways. In this work, the opposite view is adopted. An interpretation of fitness assignment as a (financial) portfolio selection problem (PSP) is proposed, where individuals are seen as assets with given (randomly distributed) monetary return values, and the fitness assigned to each individual represents an investment in that individual. In this case, it is the quality of the population that is inferred from the quality of the individuals that compose it, represented by the corresponding return distributions.

It is known from portfolio selection theory that investing only in one asset, or in similar assets, carries a risk associated with the variability of the individual returns, which is directly reflected in the variability of the overall return. Similarly, it is well known that selecting only a few of the best individuals in an EA population may lead to loss of population diversity and even to premature convergence. Therefore, the proposed analogy is completed by associating lack of population diversity with risk in the financial sense.

This paper is organized as follows. The classical PSP formulation is reviewed in the next section, leading to the proposed interpretation of fitness assignment as a portfolio selection problem. In Section 3, a new fitness assignment strategy for MOEAs is developed based on the classical PSP formulation, by specifying suitable notions of expected return and risk. This strategy is then extended to encompass solution archiving as well, allowing parental and environmental selection to be unified into a single selection problem. Preliminary experimental results on multiobjective multidimensional knapsack problem instances are presented in Section 4. The paper concludes with a discussion of the proposed approach.

2 Background

2.1 Portfolio Selection

In the classical Markowitz formulation [16] of the portfolio selection problem, asset returns are modeled as random variables, the expected values of which can usually be estimated from historical data. Risk is assessed as the variance of the overall portfolio return, and depends not only on how much individual asset returns vary, but also on how they vary in relation to one another. Thus, the covariance matrix of the joint asset return distribution is considered in addition to the expected values. A financial portfolio should optimize two conflicting objectives: maximizing the expected portfolio return and minimizing portfolio return variance. Formally:

maximize
$$\sum_{i=1}^{n} r_i x_i = r^T x$$
 (1)

minimize
$$\sum_{i=1}^{n} \sum_{j=1}^{n} q_{ij} x_i x_j = x^T Q x$$
(2)

subject to
$$\sum_{i=1}^{n} x_i = 1, \quad x_i \in [0,1], \ i = 1, \dots, n$$
 (3)

where n is the number of assets, r_i is the expected return of asset i, and q_{ij} is the covariance of the returns of assets i and j. The unknown solution is represented by $x = (x_1, \ldots, x_n)^T$, where each x_i denotes the proportion of capital to be invested in asset i.

Sharpe Ratio. The solution to the PSP defined by expressions (1–3) is a set of Pareto-optimal portfolios. The portfolios in this set are efficient with respect to expected return and return variance, and reflect different investor behavior: portfolios composed mostly of high-return assets are usually riskier, but simply avoiding risk is seldom profitable. Risk is usually reduced by combining assets with negatively correlated returns, although the expected return of the portfolio will necessarily decrease due to the inclusion of lower-return assets.

Several notions of an optimal return-to-risk trade-off have been proposed in the literature. Among them, the most widely used risk-adjusted performance index is the Sharpe ratio [5], also called reward-to-volatility ratio. The Sharpe ratio assesses how well the expected return of a given portfolio compensates the risk taken by measuring the excess return per unit of deviation from the mean with respect to a baseline, risk-free investment. The portfolio with the maximum Sharpe ratio, or *optimal risky portfolio*, x^* , is the solution of the following nonlinear programming problem:

maximize
$$\frac{r^T x - r_f}{\sqrt{x^T Q x}}$$
 (4)

subject to
$$\sum_{i=1}^{n} x_i = 1, \quad x_i \in [0,1], \ i = 1, \dots, n$$
 (5)

where r_f is the (deterministic) return of a reference, *riskless asset*. Naturally, the expected return of an efficient portfolio should be at least r_f .

This non-convex problem can be transformed into the, easier to solve, convex quadratic programming problem:

minimize
$$y^T Q y$$
 (6)

subject to
$$\sum_{i=1}^{n} (r_i - r_f) y_i = 1, \quad y_i \ge 0, \ i = 1, \dots, n$$
 (7)

by homogenizing the objective function (4), as detailed in [5]. A standard quadratic programming solver may then be used to determine the optimal risky portfolio, $x^* = y^*/k$, where $k = \sum_{i=1}^{n} y_i^*$.

2.2 Fitness Assignment as a Portfolio Selection Problem

In order to express fitness assignment in EAs as a portfolio selection problem, suitable analogues of individual expected return, r, and return covariance, Q, must be considered. Having established the values of these parameters, the desired fitness assignment, x, can be obtained by solving the resulting PSP, e.g. for the maximum Sharpe ratio portfolio.

In particular, when the covariance matrix Q is a scalar matrix and the expected return of each individual is set to the corresponding (single) objective value to be maximized, maximizing the Sharpe ratio will assign fitness proportionally to the difference between the individual objective values and the reference return value considered, if this difference is positive, and zero otherwise.

PSPs corresponding to other traditional fitness assignment strategies, such as linear ranking [3] and sigma-scaling [10], can also be formulated by considering ranks instead of objective values and/or appropriately selecting the value of the reference return. Consequently, the portfolio selection interpretation of fitness assignment *extends* conventional proportional fitness assignment by making the notion of *risk* explicit in the form of a covariance matrix.

The question remains of how to design the covariance matrix in order to control the loss of population diversity due to selection, while maintaining an appropriate level of selective pressure towards better solutions. One possible answer in the context of multiobjective optimization is proposed next.

3 A New Approach to Multiobjective Selection

Most current MOEAs attempt to drive the individuals in the population towards, and to distribute them across, the Pareto front of the problem, so that the final solution may be selected by a Decision Maker (DM) in an *a posteriori* fashion. In this way, modeling the (subjective) preferences of the DM is avoided, but the whole Pareto front must be approximated as well as possible in order to maximize the chance that at least one of the solutions found satisfies the DM.

3.1 Fitness Assignment

In such an *a posteriori* setting, the *return* of a given candidate solution may be seen as a random variable modeling the uncertainty associated with the unknown preferences of the DM. A simple scenario will be considered:

- Without loss of generality, the problem to be solved consists of the minimization of a d-dimensional objective function, f.
- There are n > 1 individuals in the population. To each individual $i, i = 1, \ldots, n$, corresponds an objective vector $f_i \in \mathbb{R}^d$.
- Preferences are expressed by the Decision Maker in terms of a single goal vector, drawn from some probability distribution over the objective space. In particular, a uniform distribution on a given orthogonal range [l, u] of the objective space will be assumed, where $l, u \in \mathbb{R}^d$.

- Individuals are either "acceptable" or "not acceptable" depending on whether or not they weakly dominate such a random goal vector, respectively. Therefore, each individual will be deemed acceptable with a certain probability, and the corresponding return (or *acceptability*) is a Bernoulli random variable taking a value of 1 if the solution is acceptable and 0 otherwise.

Under these conditions, the expected return r_i of an individual *i* is equal to the proportion of the (given orthogonal range of the) objective space which f_i dominates, i.e.,

$$r_i = p_i = \frac{\lambda([f_i, \infty[\cap [l, u]))}{\lambda([l, u])}$$
(8)

where $\lambda(\cdot)$ denotes the Lebesgue measure (or hypervolume) of the given region, and $[f_i, \infty]$ is the region dominated by f_i . Since returns are Bernoulli distributed, the return covariance for a pair of individuals is

$$q_{ij} = p_{ij} - p_i p_j \tag{9}$$

$$= \frac{\lambda([f_i, \infty[\cap [f_j, \infty[\cap [l, u])])}{\lambda([l, u])} - p_i p_j$$
(10)

$$=\frac{\lambda([(f_i \vee f_j), \infty[\cap [l, u]))}{\lambda([l, u])} - p_i p_j$$
(11)

where $f_i \vee f_j$ denotes the *join*, or componentwise maximum, of objective vectors f_i and f_j , for i, j = 1, ..., n. Note that $p_{ii} = p_i$, and that $q_{ii} = p_i - p_i^2$ is simply the variance of the return of individual *i*. As a consequence, the return of a riskless asset must be zero $(r_f = 0)$ under this model.

The above expressions show that the return covariance relates the size of the region simultaneously attained by two individuals to the sizes of the regions attained individually by each one, an idea which is also at the heart of the definition of an extended dominance relation known as volume dominance [13], although the details of the two methods are considerably different. Whereas the aim of volume dominance is to establish whether an individual should be considered better than another, here the aim is to gauge the (dis)similarity between individuals. Positive covariance indicates that two individuals attain much the same region of the objective space, whereas negative covariance is a sign that the regions attained do not overlap much. Since portfolio variance is reduced by combining negatively correlated assets, and greater returns are preferred, a risk-adjusted portfolio should consist of a diverse set of individuals along the non-dominated front.

Illustrative examples are presented in Fig. 1. When all individuals are nondominated and evenly distributed on a linear front (left), the fitness assigned to each one by maximizing the Sharpe ratio is the same. On the other hand, if they are not evenly distributed (center), isolated individuals are assigned greater fitness values than those in crowded regions of the front. In the more general case (right), dominated individuals, as well as some non-dominated individuals, are assigned zero fitness, whereas the remaining non-dominated individuals are



Fig. 1. Maximum Sharpe ratio fitness assignment. Left and center: 30 non-dominated individuals. Right: an arbitrary 30-individual population. Circle area is proportional to assigned fitness. l = (0, 0), u = (1, 1).

assigned high fitness values. Therefore, the optimal risky portfolio may not maximize the hypervolume indicator on an arbitrary front, although it would seem to correctly identify at least some interesting non-dominated solutions.

3.2 Environmental Selection and Archiving

The portfolio selection model described so far is aimed at *parental* selection, i.e., at the selection of individuals for breeding. Once fitness has been assigned, this is typically achieved with a sampling procedure such as roulette-wheel selection (RWS) [8] or stochastic universal sampling (SUS) [2].

Environmental selection, on the other hand, is aimed at replacing old individuals with new ones, and may be performed based on fitness, the number of iterations an individual has survived in the population, or even randomly. Also, offspring may replace the parents unconditionally or depending on whether they outperform them. A replacement strategy where the best individuals always survive is known as elitist.

In contrast to the single objective case, the implementation of elitist environmental selection in the multiobjective case must deal with the possible incomparability between individuals. If only dominated individuals are ever replaced by new ones, either the population is allowed to grow indefinitely or the algorithm will terminate as soon as all individuals in the population are non-dominated. For this reason, alternative environmental selection strategies where non-dominated individuals may be replaced by new ones in order to keep the size of the population constant have been proposed and extensively studied [12,11,15]. Because the population acts as an archive of non-dominated solutions, such strategies have become known as *bounded archiving* strategies.

MOEAs typically implement environmental selection separately from parental selection. However, the two can be meaningfully combined into a single portfolio selection problem with a cardinality constraint. Assuming a parental population size of n and the production of m offspring at each iteration, the new problem consists of assigning non-zero fitness to at most n individuals from the n + m

parents and offspring currently available. In this way, environmental selection is performed so as to maximize the Sharpe ratio of the resulting portfolio. It is not difficult to see that this approach guarantees that the Sharpe ratio may never decrease from one iteration to the next. Therefore, this combined environmental selection and fitness assignment mechanism implements a monotonic bounded archiving strategy [15].

Unfortunately, portfolio selection with cardinality constraints is no longer a convex optimization problem, and may be difficult to solve exactly (it is generally NP-hard). In practice, however, this will depend on how tightly constrained a particular instance turns out to be. Indeed, solving the relaxed problem will, in many cases, lead to a solution that satisfies the cardinality constraint, unless the population is so well distributed that more than n individuals would, in principle, be assigned non-zero fitness. This simply requires that m is chosen to be sufficiently small in comparison to n.

4 Experimental Results

The proposed approach to multiobjective selection was implemented in an otherwise conventional mutation-selection evolutionary algorithm with population size n = 200. The parental population was sampled proportionally to assigned fitness using SUS [2] to select m = 50 parents that were mutated to produce moffspring. For simplicity, no recombination operator was used.

The algorithm was applied to multiobjective knapsack instances from [19] with 100 and 500 items, with 2, 3 and 4 objectives, and as many capacity constraints as objectives. For the purpose of constructing the portfolio selection problems, knapsack values and weights were taken to range from zero to their maximum values (when all items are included). The return distribution of each individual was computed as described in section 3.1, but using the preferability relation [7] instead of dominance, in order to accommodate constraints and objectives in the same formulation. Additional (linear) constraints on fitness were imposed so that no individual could expect to reproduce more than once in each generation.

Individuals were represented as binary strings, and mutation consisted of either flipping one bit at random or exchanging two randomly-selected bits of different value, so as to uniformly sample the resulting 1-flip-exchange neighborhood [14]. In order to study the long-run behavior of the algorithm, here referred to as Portfolio Optimization Selection Evolutionary Algorithm (POSEA), and account for its stochasticity, 13 long runs with 1 million function evaluations each were performed for each instance. For comparison purposes, equally long runs were performed using SPEA2 [18], NSGA-II [6] and SMS-EMOA [4] with the same population size (n = 200) and mutation operator. As before, no recombination was used. The number of offspring per generation, m, was set to the default in each case: m = n in SPEA2 and NSGA-II and m = 1 in SMS-EMOA.

Experimental results are presented in Figs. 2 and 3, where the maximum, median and minimum of 13 runs are shown for each algorithm on 2-, 3- and 4-objective instances. SPEA2 and NSGA-II do not use hypervolume information, and tend to perform worse or, at most, slighly better than SMS-EMOA



Fig. 2. Performance on 100-item knapsack instances: 2 objectives (left), 3 objectives (center) and 4 objectives (right)



Fig. 3. Performance on 500-item knapsack instances: 2 objectives (left), 3 objectives (center) and 4 objectives (right)

and POSEA with respect to both the hypervolume indicator and the Sharpe ratio. SMS-EMOA can be seen to perform best with respect to the hypervolume indicator, whereas POSEA clearly does not cover the Pareto-front as well as SMS-EMOA. On the other hand, SMS-EMOA achieves lower or, at most, similar values of Sharpe ratio, indicating that POSEA focused on the most relevant non-dominated solutions according to the DM model adopted.

5 Conclusions

In this work, a fitness assignment approach based on (financial) portfolio optimization was proposed. By modeling the uncertainty associated with Decision Maker preferences probabilistically, the quality (or return) of each individual solution becomes a random variable, and fitness assignment consists of forming a portfolio of individuals balancing overall expected return against return variance, e.g., based on the Sharpe ratio.

Although the probabilistic Decision Maker model adopted in this work is rather simplistic, empirical evidence suggests that it possesses some interesting properties, such as not favoring dominated solutions over non-dominated ones and promoting diversity in the population, even if it does not maximize the hypervolume indicator in the general case. A theoretical study of these and other properties is currently under way [9], but more experimentation is required to evaluate the performance of POSEA and how the number of offspring, m, may influence it. Furthermore, since the size of the quadratic programming problem to be solved at each generation is independent of the number of objectives, the method is potentially much faster than hypervolume-based selection (depending on m), especially as the number of objectives grows.

The proposed approach establishes a bridge between multiobjective selection and optimization under uncertainty. By considering alternative probabilistic DM models based on other indicators and/or preference articulation strategies, the portfolio optimization paradigm should contribute to unifying solution-oriented preferences and set-oriented preferences under a common framework.

Acknowledgments. I. Yevseyeva's research was funded by the Academy of Finland (grant 126476) and the EC Erasmus Mundus ECW Lot 6 project at the Faculty of Science and Technology of the University of Algarve, Portugal. A. P. Guerreiro acknowledges Fundação para a Ciência e a Tecnologia (FCT) for Ph.D. studentship SFHR/BD/77725/2011, co-funded by the European Social Fund and by the State Budget of the Portuguese Ministry of Education and Science in the scope of NSRF– HPOP–Type 4.1–Advanced Training.

References

- 1. Bader, J., Zitzler, E.: HypE: An algorithm for fast hypervolume-based manyobjective optimization. Evolutionary Computation 19(1), 45–76 (2011)
- Baker, J.E.: Reducing bias and inefficiency in the selection algorithm. In: Proc. Second International Conference on Genetic Algorithms, pp. 14–21 (1987)

- Baker, J.E.: Adaptive selection methods for genetic algorithms. In: Proc. First International Conference on Genetic Algorithms, pp. 101–111 (1985)
- Beume, N., Naujoks, B., Emmerich, M.: SMS-EMOA: Multiobjective selection based on dominated hypervolume. EJOR 181, 1653–1669 (2007)
- 5. Cornuejols, G., Tutuncu, R.: Optimization Methods in Finance. Cambridge University Press (2007)
- Deb, K., Pratap, A., Agarwal, S., Meyarivan, T.: A fast and elitist multiobjective genetic algorithm: NSGA-II. IEEE Transactions on Evolutionary Computation 6(2), 182–197 (2002)
- Fonseca, C.M., Fleming, P.J.: Multiobjective optimization and multiple constraint handling with evolutionary algorithms—Part I: A unified formulation. IEEE Transactions on Systems, Man and Cybernetics—Part A: Systems and Humans 28(1), 26–37 (1998)
- 8. Goldberg, D.: Genetic Algorithms in Search, Optimization, and Machine Learning. Addison-Wesley, Reading (1989)
- 9. Guerreiro, A.P.: Portfolio Selection in Evolutionary Algorithms. Ph.D. thesis proposal, University of Coimbra, Coimbra, Portugal (2012)
- Hancock, P.J.B.: An empirical comparison of selection methods in evolutionary algorithms. In: Fogarty, T.C. (ed.) AISB-WS 1994. LNCS, vol. 865, pp. 80–94. Springer, Heidelberg (1994)
- Knowles, J., Corne, D.: Properties of an adaptive archiving algorithm for storing nondominated vectors. IEEE Transactions on Evolutionary Computation 7(2), 100–116 (2003)
- Knowles, J., Corne, D., Fleisher, M.: Bounded archiving using the Lebesgue measure. In: Proc. IEEE Congress on Evolutionary Computation (CEC 2003), vol. 4, pp. 2490–2497. IEEE Press, New York (2003)
- Le, K., Landa-Silva, D.: Obtaining better non-dominated sets using volume dominance. In: IEEE Congress on Evolutionary Computation (CEC 2007), pp. 3119–3126 (September 2007)
- Liefooghe, A., Paquete, L., Figueira, J.R.: On local search for bi-objective knapsack problems. Evolutionary Computation 21(1), 179–196 (2013)
- López-Ibáñez, M., Knowles, J., Laumanns, M.: On sequential online archiving of objective vectors. In: Takahashi, R.H.C., Deb, K., Wanner, E.F., Greco, S. (eds.) EMO 2011. LNCS, vol. 6576, pp. 46–60. Springer, Heidelberg (2011)
- 16. Markowitz, H.: Portfolio selection. Journal of Finance 7(1), 77–91 (1952)
- Sareni, B., Krahenbuhl, L.: Fitness sharing and niching methods revisited. IEEE Transactions on Evolutionary Computation 2(3), 97–106 (1998)
- Zitzler, E., Laumanns, M., Thiele, L.: SPEA2: Improving the strength pareto evolutionary algorithm. In: Evolutionary Methods for Design, Optimization and Control with Applications to Industrial Problems (EUROGEN 2001), pp. 95–100 (2002)
- Zitzler, E., Thiele, L.: Multiobjective Optimization Using Evolutionary Algorithms — A Comparative Case Study. In: Eiben, A.E., Bäck, T., Schoenauer, M., Schwefel, H.-P. (eds.) PPSN 1998. LNCS, vol. 1498, pp. 292–301. Springer, Heidelberg (1998)
- Zitzler, E., Thiele, L., Laumanns, M., Fonseca, C., Grunert da Fonseca, V.: Performance assessment of multiobjective optimizers: An analysis and review. IEEE Transactions on Evolutionary Computation 7(2), 117–132 (2003)