

MULTIOBJECTIVE RANKING AND SELECTION BASED ON HYPERVOLUME

Juergen Branke

Warwick Business School
The University of Warwick
Coventry, CV4 7AL, UK

Wen Zhang

Warwick Business School
The University of Warwick
Coventry, CV4 7AL, UK

Yang Tao

School of Engineering and Computer Science
The University of Durham
Durham, DH1 3LE, UK

ABSTRACT

In this paper, we propose a myopic ranking and selection procedures for the multi-objective case. Whereas most publications for multi-objective problems aim at maximizing the probability of correctly selecting all Pareto optimal solutions, we suggest minimizing the difference in hypervolume between the observed means of the perceived Pareto front and the true Pareto front as a new performance measure. We argue that this hypervolume difference is often more relevant for a decision maker. Empirical tests show that the proposed method performs well with respect to the stated hypervolume objective.

1 INTRODUCTION

Ranking and Selection (R&S) aims at efficiently identifying the best out of a given set of alternatives, where best is defined by expected performance, and performance can only be estimated by sampling. However, many practical real-world optimization problems require the consideration of multiple contradicting objectives (Ponweiser et al. 2008), which means that there usually does not exist a single solution that is best in all objectives. Instead, there is a set of so-called Pareto optimal solutions with different trade-offs of the objectives. A solution is called Pareto optimal if it is not dominated by any other solution. A solution is said to dominate another solution if it is at least as good in each objective, and strictly better in at least one objective. In the absence of additional preference information, it is not possible to rank Pareto optimal solutions. Thus, in multi-objective ranking and selection (MORS), it is usually the goal to identify all Pareto optimal solutions so that they can be presented to the decision maker (DM) to choose from.

Most MORS procedures aim at maximizing the probability of correct selection (PCS), which in this case means exactly identifying the set of Pareto optimal solutions (i.e., correctly classifying each solution as either Pareto optimal or dominated). Examples include the multi-objective optimal computing budget allocation (MOBCA) proposed by Lee et al. (2010), which is the multi-objective version of the Optimal Computing Budget Allocation (OCBA) algorithm (Chen et al. 2000); the approach by Hunter and Feldman (2015) and Feldman, Hunter, and Pasupathy (2015) which allocates samples to maximize the rate of decay, is asymptotically optimal, and can take into account correlation between objectives; and our own myopic strategy M-MOBA (Branke and Zhang 2015). MOCBA has also been extended to allow for an indifference zone (Teng, Lee, and Chew 2010) and other measures of selection quality such as expected opportunity cost (EOC) (Lee, Chew, and Teng 2007; Lee, Chew, and Teng 2010; He, Chick, and Chen

2007). Branke and Gamer (2007) and Frazier and Kazachkov (2011) use expected utility assuming a linear utility function model. Other related methods include approaches based on the idea of racing, such as S-Race proposed by Zhang, Georgiopoulos, and Anagnostopoulos (2013) or the racing algorithm for use inside an evolutionary algorithm presented by Marceau-Caron and Schoenauer (2014), and a trust-region based method for approximating the Pareto front of a bi-objective stochastic optimization problem (Kim and Ryu 2011). Also, there are related methods that look at stochastic constraints in addition to a single performance criterion (e.g., Andradottir and Kim 2010; Pasupathy et al. 2014; Hu and Andradottir 2014)

In this paper, we propose a new performance measure for MORS, the hypervolume (HV) difference, based on the HV measure that is commonly used to evaluate results in multiobjective optimization (Emmerich and Klinkenberg 2008). As far as we know, this paper is the first attempt to use hypervolume as the performance measure to the MORS problem. We then continue to derive a new myopic ranking and selection procedure similar to our M-MOBA (Branke and Zhang 2015) but based on the new HV difference criterion rather than the probability of correct selection.

The paper is organized as follows. Section 2 formalizes the problem and describes the assumptions. Section 3 describes the proposed myopic hypervolume-based MORS procedure. Section 4 presents empirical simulation results, and the paper concludes in Section 5 with a summary and suggestions for future work.

2 PROBLEM FORMULATION

2.1 Notation and problem context

Throughout this paper, we assume the goal is to minimize all objectives. Given H objectives and a set of m designs with the true unknown performance of each design i in objective h being denoted by w_{ih} . A design i is said to dominate design j ($i \prec j$) if and only if $w_{ih} \leq w_{jh}$ for all objectives and $w_{ih} < w_{jh}$ for at least one objective. A design that is not dominated by any other design is called Pareto optimal.

The performances of each design in each objective need to be estimated via sampling. Let \mathbf{X}_i be a matrix that contains the simulation output for design i . Then $\mathbf{X}_i = (X_{ihn})$, where X_{ihn} is the h -th objective of design i for simulation replication n . Let furthermore w_{ih} and σ_{ih}^2 be the unknown mean and variance of alternative i , which can only be estimated using the simulation outputs X_{ihn} . We assume that

$$\{X_{ihn} : n = 1, 2, \dots\} \stackrel{iid}{\sim} \mathcal{N}(w_{ih}, \sigma_{ih}^2), \text{ for } i = 1, 2, \dots, m \text{ and } h = 1, 2, \dots, H.$$

Let n_i be the number of samples taken for alternative i so far, \bar{x}_{ih} be the sample mean and $\hat{\sigma}_{ih}^2$ be the sample variance. Then, we will get an observed Pareto set based on the $n = \sum_i n_i$ simulations so far. As n_i increases, \bar{x}_{ih} and $\hat{\sigma}_{ih}^2$ will be updated and the observed Pareto front may change accordingly. If alternative i is to receive another τ_i samples, and \bar{y}_{ih} is the average of the new samples in objective h , then the new overall sample mean in each objective can be calculated as

$$z_{ih} = \frac{n_i \bar{x}_{ih} + \tau_i \bar{y}_{ih}}{n_i + \tau_i}. \quad (1)$$

Before the new samples are observed, the sample average that will arise after sampling, denoted as Z_{ih} , is a random variable, and we can use the predictive distribution for the new samples (DeGroot 2005) and get

$$Z_{ih} \sim \phi(\bar{x}_{ih}, n_i * (n_i + \tau_i) / (\tau_i * \hat{\sigma}_{ih}^2), n_i - 1)$$

where $\phi(\mu, \kappa, \nu)$ denotes the student distribution with mean μ , precision κ and ν degrees of freedom.

2.2 Hypervolume difference

In multiobjective optimization, the goal is usually, to find a set of solutions that approximate, as closely as possible, the set of truly Pareto-optimal solutions. One of the most common measures for the quality of a Pareto front approximation is the hypervolume (HV) measure. The HV measure rewards finding solutions

close to the true Pareto front, as well as a good spread of solutions along the true Pareto front. (Beume, Naujoks, and Emmerich 2007). Let Λ denote the Lebesgue measure, then the HV metric is defined as

$$HV(B, y_{ref}) := \Lambda\left(\bigcup_{y \in B} \{y' \mid y \prec y' \prec y_{ref}\}\right), \quad B \subseteq \mathbb{R}^m \quad (2)$$

where B is a set of solutions and $y_{ref} \in \mathbb{R}^m$ denotes a reference point that is usually user defined and should be dominated by all solutions the user might possibly be interested in. In the context of ranking and selection, the reference point could also be determined based on the initial observations for each system.

Figure 1 shows a set of 5 alternatives in 2-objective space. Three of the solutions are Pareto-optimal, and the HV is the shaded area, defined by the Pareto optimal solutions and the reference point R . The dominated solutions do not contribute to the HV. Figure 2 provides an example for the proposed performance measure, the HV difference (HVD). Given two sets of Pareto-optimal solutions A and B ,

$$HVD(A, B, y_{ref}) := HV(A, y_{ref}) + HV(B, y_{ref}) - 2 * \Lambda(HV(A, y_{ref}) \cap HV(B, y_{ref})). \quad (3)$$

We believe that the HV difference is more relevant to a decision maker than the percentage of correctly identified Pareto-optimal solutions. Similar to EOC, this is a smooth error function, and not step-wise as PCS. If a solution that should be dominated is estimated as just marginally Pareto-optimal, this error has less weight than if it was perceived as clearly Pareto-optimal. It will also emphasize the correct estimation of the performance of alternatives that are clearly Pareto-optimal and far away from other solutions, compared to Pareto-optimal solutions that are only marginally Pareto-optimal, and that have close neighbors. This makes a lot of sense, since the former are much more likely to be selected by a decision maker.




Figure 1: Hypervolume of a set of solutions.




Figure 2: Hypervolume difference of two sets of solutions.

3 MYOPIC BI-OBJECTIVE BUDGET ALLOCATION EOC PROCEDURE

Our derivation of the HVD is partly based on the expected hypervolume computation in (Emmerich and Klinkenberg 2008). For the sake of convenience, we restrict our consideration in the remainder of this paper to only two objectives, which has the advantage of allowing visualization. However, extension to more than two objectives should be possible, analogous to (Emmerich and Klinkenberg 2008). We start with a few illustrative examples of the HV change in Section 3.1. Then we discuss how to approach this computationally in Section 3.2, and provide the mathematical derivation of a closed formula in Section 3.3.

3.1 Determining the hypervolume change

After $n = \sum_i n_i$ samples, we can estimate the mean performance for each alternative and each objective. Based on the current estimates, we obtain the current Pareto front, which will be denoted f_1 in this paper.

An example is shown in Figure 3, where the Pareto front consists of points A , B , M and C . The reference point R is not a data point but is needed for hypervolume calculation. If we take another sample of a particular design, say M , then the estimated objective values for M and thus the location of M will change to M' , with corresponding changes to the Pareto front and the HV. We have to integrate over our predictive distribution for M' to calculate the expected HV change. Depending on the region of the new M' , the calculation of the HV change will differ.




Figure 3: An observed Pareto front f_1 .

3.1.1 For a currently dominated alternative

Let us first consider sampling point G , which is not on f_1 (i.e., dominated). If it is updated due to a new sample but remains dominated, it doesn't change the hypervolume. If it becomes non-dominated, for example, if it moves to a new position G' in Figure 4, it will cause the HV to increase and the increased part is the shaded.




Figure 4: Increased part lead by G' .




Figure 5: Decreased part when M moves to M_0 .

3.1.2 For a currently non-dominated alternative

For point M , which is currently non-dominated, the situation is more complex. If the alternative becomes dominated by allocating a new sample to it, such as if it moves to M_0 in Figure 5, the HV will decrease by the shaded area between M and the new Pareto optimal solutions D, E, F . This area is constant as long as M_0 becomes dominated, and can be easily calculated.

If it shifts to a new position M_1 above and left of its left neighbor on the Pareto front, say B in Figure 6 (or M_2 below and right of the right neighbor C), it will cause an increase in HV in one area (near M_1 or near M_2), and a decrease in another (the part dominated by the original M).

If M moves to a location that dominates its previous location, such as M_3 in Figure 7, this change will only increase the HV by the shaded area, there is no reduction of HV.




Figure 6: Change lead by M to M_1 .




Figure 7: Increased part when M moves to M_3 .

If M' moves into the area originally dominated by M , but remains non-dominated, there is only a reduction of HV, an example is provided in Figure 8, which is a zoom-in of the particular area of interest.

Finally, in the remaining area, there is again a part that is increased and a part that is decreased, see Figure 9 for an example.




Figure 8: Increased part when M moves to M_6 .




Figure 9: Change lead by M to M_4 .

3.2 Calculating the hypervolume change

As we have seen above, drawing an additional sample for a particular alternative and the subsequent change in this alternative's mean value can lead to an increase of HV in some area and a decrease in some other area. As we will see below, calculating the expected HV change requires to break down the calculation into different cells, but for each cell, we can find a closed form expression for the expected HV change. Then, these expected changes need to be added up to result in the the overall expected HV change.

In the following, we will explain the computation for one particular cell, with other cells being computed in a similar fashion.

Consider Figure 10, where all individuals on the Pareto front f_1 are labelled x_1, \dots, x_k , with coordinates x_{ij} for alternative i and objective j , and the x_i sorted in increasing order of objective 1. For technical reasons, let us define $x_{01} = -\infty, x_{02} = x_{r2}, x_{k+1,1} = x_{r1}, x_{k+1,2} = -\infty$. We consider another sample for design x_m , and the calculation for one particular cell that is marked by bold in the figure and defined by upper right corner u with coordinates (x_u, y_u) and lower left corner l with coordinates (x_l, y_l) . Let us assume that these two corners are defined by the Pareto optimal solutions p and q , by $u = (x_{p+1,1}, x_{q-1,2})$ and $l = (x_{p,1}, x_{q,2})$.

Then, the contribution of the cell to the expectation of the HV change when sampling design x_m is

$$\int_{y_l}^{y_u} \int_{x_l}^{x_u} \left[(x_{p+1,1} - x)(x_{p2} - y) + \sum_{p < i < q} (x_{i+1,1} - x_{i1})(x_{i2} - y) \right] \cdot \phi_{m1}(x) \cdot \phi_{m2}(y) dx dy \quad (4)$$

where ϕ_{mi} is the predictive probability distribution of the new location of x_m in dimension i .




Figure 10: Different cells that need to be considered when calculating the expected HV change from re-sampling.

3.3 Mathematical derivation of closed formula

In the following, we derive a closed form for calculating the expected HV change in one cell.

Let $\phi(x; \mu, \kappa, \nu)$ denote the distribution of $\mu + \frac{1}{\sqrt{\kappa}} T_\nu$, where T_ν is a random variable with standard t distribution with ν degrees of freedom, i.e., the t distribution we estimate for the new location of an alternative's mean values after having taken another sample, with mean μ , precision κ and ν degrees of freedom. The cumulative density function is then

$$\Phi(x; \mu, \kappa, \nu) = \Phi_t(\sqrt{\kappa}(x - \mu); \nu). \quad (5)$$

with $\Phi_t(x; \nu)$ the cumulative standard t -distribution, and the probability density function is

$$\phi(x; \mu, \kappa, \nu) = \sqrt{\kappa} \cdot \phi_t(\sqrt{\kappa}(x - \mu); \nu) = \sqrt{\frac{\kappa}{\nu\pi}} \frac{\Gamma(\frac{\nu+1}{2})}{\Gamma(\frac{\nu}{2})} \cdot \left(1 + \frac{\kappa(x - \mu)^2}{\nu} \right)^{-\frac{\nu+1}{2}} \quad (6)$$

with $\phi_t(x; \nu)$ the standard t -distribution. The HV change, due to the point we are considering moving to a new position (x, y) , is always a function in the form $axy + bx + cy + d$. The constant coefficients a, b, c, d are different in different areas and some of the coefficients could be 0 sometimes. The contribution of the area $[x_l, x_u] \times [y_l, y_u]$ (e.g, the small cell discussed in Figure 10) to the expectation of the HV change is

$$\begin{aligned} & \int_{x_l}^{x_u} \int_{y_l}^{y_u} (axy + bx + cy + d) \cdot \phi_{i1}(x) \cdot \phi_{i2}(y) dx dy \\ &= a \int_{x_l}^{x_u} x \phi_{i1}(x) dx \int_{y_l}^{y_u} y \phi_{i2}(y) dy + b \cdot \Phi_{i2}(y)|_{y_l}^{y_u} \cdot \int_{x_l}^{x_u} x \phi_{i1}(x) dx \\ &+ c \cdot \Phi_{i1}(x)|_{x_l}^{x_u} \cdot \int_{y_l}^{y_u} y \phi_{i2}(y) dy + d \cdot \Phi_{i1}(x)|_{x_l}^{x_u} \cdot \Phi_{i2}(y)|_{y_l}^{y_u} \end{aligned} \quad (7)$$

where $\phi_{ih}(x) = \phi(x; \mu_{ih}, \kappa_{ih}, \nu_i)$, $\Phi_{ih}(x) = \Phi(x; \mu_{ih}, \kappa_{ih}, \nu_i)$, $\mu_{ih} = \bar{x}_{ih}$, $\kappa_{ih} = n_i(n_i + \tau_i) / \tau_i \hat{\sigma}_{ih}^2$ and $\nu_i = n_i - 1$. On the right hand side of Eq. (7), the most critical part is solving the integrals, and it can be done by calculating the corresponding indefinite integral, which is

$$\begin{aligned} \int x\phi(x; \mu, \kappa, \nu)dx &= \int (x - \mu)\phi(x; \mu, \kappa, \nu)dx + \mu\Phi(x; \mu, \kappa, \nu)dx \\ &= \psi(x; \mu, \kappa, \nu) + \mu\Phi(x; \mu, \kappa, \nu) \end{aligned} \tag{8}$$

with

$$\begin{aligned} \psi(x; \mu, \kappa, \nu) &:= \int (x - \mu)\phi(x; \mu, \kappa, \nu)dx \\ &= \sqrt{\frac{\nu}{\kappa\pi}} \cdot \frac{\Gamma(\frac{\nu+1}{2})}{(1-\nu)\Gamma(\frac{\nu}{2})} \left(1 + \frac{\kappa(x - \mu)^2}{\nu}\right)^{\frac{1-\nu}{2}} = \frac{\nu + \kappa(x - \mu)^2}{(1-\nu)\sqrt{\kappa}} \phi(x; \mu, \kappa, \nu). \end{aligned} \tag{9}$$

In the rest of this paper, for convenience, we will denote $\psi(x; \mu_{ih}, \kappa_{ih}, \nu_i)$ as $\psi_{ih}(x)$. Using the above results and gathering the terms with same integrals, Eq. (7) can be rewritten as

$$\begin{aligned} &\int_{x_l}^{x_u} \int_{y_l}^{y_u} (axy + bx + cy + d) \cdot \phi_{i1}(x) \cdot \phi_{i2}(y) dx dy \\ &= a\Psi_{i1}(x)|_{x_l}^{x_u} \Psi_{i2}(y)|_{y_l}^{y_u} + (b + a\mu_{i2})\Psi_{i1}(x)|_{x_l}^{x_u} \Phi_{i2}(y)|_{y_l}^{y_u} \\ &\quad + (c + a\mu_{i1})\Phi_{i1}(x)|_{x_l}^{x_u} \Psi_{i2}(y)|_{y_l}^{y_u} + (a\mu_{i1}\mu_{i2} + b\mu_{i1} + c\mu_{i2} + d)\Phi_{i1}(x)|_{x_l}^{x_u} \Phi_{i2}(y)|_{y_l}^{y_u} \end{aligned} \tag{10}$$

For example, considering the integral (4), we will have

$$a = 1; \quad b = -x_{p2}; \quad c = -x_{p+1,1} - \sum_{p < j < q} (x_{j+1,1} - x_{j1}); \quad d = x_{p+1,1}x_{p2} + \sum_{p < j < q} (x_{j+1,1} - x_{j1})x_{j2},$$

and then we can substitute them, in addition to

$$\mu_{mh} = \bar{x}_{mh}; \quad \kappa_{mh} = n_m(n_m + \tau_m) / \tau_m \hat{\sigma}_{mh}^2; \quad \nu_m = n_m - 1,$$

where $h = 1$ or 2 , into Eq. (10) to solve the integral (4).

4 M-MOBA-HV ALGORITHM

The overall Myopic Multi-Objective Budget Allocation procedure based on the HV change criterion is denoted as M-MOBA-HV and is summarized below.

1. Specify a first-stage sample size $n_0 \geq 5$, and a number of samples $\tau > 0$ to allocate per subsequent stage. Specify stopping rule parameters.
2. Sample X_{ihm} , $i = 1, \dots, m; h = 1, \dots, H; n = 1, \dots, n_0$ independently, and initialize the number of samples $n_i \leftarrow n_0$
3. Determine the sample statistics \bar{x}_{ih} and $\hat{\sigma}_{ih}^2$, and the observed Pareto front.
4. WHILE stopping rule not satisfied DO another stage:
 - (a) For each alternative i , calculate the expected hypervolume change VC_i that would result from allocating τ additional sample to alternative i .
 - (b) Allocate τ samples to the alternative i that has the largest VC_i .
 - (c) Update sample statistics n_i , \bar{x}_{ik} and $\hat{\sigma}_{ik}^2$ and observe a new Pareto set.
5. Select alternatives on the observed Pareto set and calculate the difference between the selected and true Pareto set.

4.1 Empirical Simulation Results

In this section, we test M-MOBA-HV on three configurations, and compare it with two other methods, the original M-MOBA (Branke and Zhang 2015) which is based on PCS, and Equal allocation that allocates the same number of samples to each design. For each method, each design is sampled $n_0 = 5$ times during initialization, and then additional samples are allocated one at a time ($\tau = 1$) until a pre-set budget has been reached. Results are averaged over 1000 runs. In case we run into problems of numerical precision, τ is changed to 10 for the expected HV change calculation, but still only one sample is calculated. If the numerical precision problem persists, we will use Equal allocation until the problem stops and τ is then set back to 1.

4.1.1 Configuration 1: Standard benchmark

The first configuration has originally been proposed by Chen and Lee (2010) and is also used in (Branke and Zhang 2015) to test M-MOBA. There are 16 alternative designs, 2 objectives, and the standard deviation of each alternative in each objective is set to 2. The reference point has been set to (17,17) for these experiments. Expected values of each design are shown in Table 1.

Table 1: Configuration 1 with 16 alternatives and two objectives. Standard deviation for all designs is 2 in each objective.

Index	Obj. 1	Obj. 2	Index	Obj. 1	Obj. 2
0	0.5	5.5	8	4.8	5.5
1	1.9	4.2	9	5.2	5
2	2.8	3.3	10	5.9	4.1
3	3	3	11	6.3	3.8
4	3.9	2.1	12	6.7	7.2
5	4.3	1.8	13	7	7
6	4.6	1.5	14	7.9	6.1
7	3.8	6.3	15	9	9

Figure 11 reports the reduction of the HV difference as the number of samples allocated increases. It can be seen that the M-MOBA-HV method works much better than both the Equal and original M-MOBA methods in terms of HV difference between the selected and true Pareto set. Although M-MOBA has been shown to identify the Pareto optimal solutions much more quickly than Equal allocation on this problem (Branke and Zhang 2015), in terms of HV difference it is actually only slightly better than Equal allocation.




Figure 11: Comparison of hypervolume difference of Configuration 1.

Table 2: Configuration 2 with 10 alternatives and two objectives. Standard deviation for all designs is 2 in each objective.

Index	Obj. 1	Obj. 2
1	1	5
2	5	1
3	3	3
4	3.1	2
5	2	3.1
6	4	2.1
7	2.1	4
8	5.5	5
9	3.5	5
10	6	6




Figure 12: Configuration 2 with 10 alternatives.

Table 3: Configuration 3 with 8 alternatives and two objectives. Standard deviation for all designs is 2 in each objective.

Index	Obj. 1	Obj. 2
1	1	5
2	5	1
3	3.2	2.1
4	3	2
5	2	3.1
6	6	4
7	5	5
8	4	6




Figure 13: Configuration 3 with 8 alternatives.

4.1.2 Configuration 2: Borderline designs

The second configuration is designed to show the impact of solutions that are close to being dominated or non-dominated. These points have a small influence on the resulting HV, and whether they are actually identified as dominated or non-dominated may not matter so much to a decision maker. The configuration has 10 designs, 2 objectives, and the standard deviation of each alternative in each objective is also set to 2. The reference point is (10,10) in this case. Expected values of each design are shown in Table 2 and a visualization is provided in Figure 12. Designs 6 and 7 are dominated, but close to being non-dominated, and design 3 is non-dominated, but close to being dominated.

The result is shown in Figure 14. Again, M-MOBA-HV works very well. The PCS-based version M-MOBA now is even worse than Equal allocation. To investigate this further, Figure 15 shows the percentage of samples allocated to a particular design. M-MOBA allocates quite a few samples to the borderline designs 3, 6 and 7, because it aims to improve the probability of correct selection, and for these designs the classification is most difficult. For a decision maker, however, these designs are probably less relevant. M-MOBA-HV instead focuses on the designs 1,2,4 and 5, which are the clearly Pareto optimal

solutions probably most relevant to a decision maker. Thus, it creates reliable performance estimates where it is most relevant.




Figure 14: Comparison of hypervolume difference for Configuration 2.




Figure 15: Allocation of samples to the different alternatives for Configuration 2.

4.1.3 Configuration 3: Similar designs

The third configuration is designed to show the impact of very similar designs. Again, for PCS-based MORS algorithms, it is difficult to distinguish between them. On the other hand, the distinction is probably not very relevant for a decision maker. There are 8 designs, 2 objectives, and the standard deviation of each alternative in each objective is set to 2. Expected values of each design are shown in Table 3, with a visualization in Figure 13.

The results are similar to Configuration 2 in the sense that M-MOBA-HV works best, and the PCS-based M-MOBA is worse than Equal allocation. Again, Figure 17 provides further detail on the distribution of samples onto the different alternatives.




Figure 16: Hypervolume difference for Configuration 3.




Figure 17: Allocation of samples to the different alternatives for Configuration 3.

5 CONCLUSION

In this paper, a new myopic budget allocation algorithm for solving multi-objective problems with the goal of minimizing the difference in hypervolume has been proposed. To the best of our knowledge, this constitutes the first approach based on hypervolume. We argue that the hypervolume difference criterion leads to information collection that is more relevant to a decision maker, taking more samples from the more "interesting" areas of the Pareto front. Empirical comparisons to Equal allocation and the PCS-based M-MOBA show that the new method indeed works very well with respect to the hypervolume difference performance measure.

REFERENCES

- Andradottir, S., and S.-H. Kim. 2010. "Fully sequential procedures for comparing constrained systems via simulation". *Naval Research Logistics* 57 (5): 403–421.
- Beume, N., B. Naujoks, and M. Emmerich. 2007. "SMS-EMOA: Multiobjective selection based on dominated hypervolume". *European Journal of Operational Research* 161 (3): 1663–1669.
- Branke, J., and J. Gamer. 2007. "Efficient sampling in interactive multi-criteria selection". In *INFORMS Simulation Society Research Workshop*, 42–46: INFORMS Simulation Society.
- Branke, J., and W. Zhang. 2015. "A new myopic sequential sampling algorithm for multi-objective problems". In *Proceedings of the 2015 Winter Simulation Conference*, edited by L. Yilmaz, W. K. V. Chan, I. Moon, T. M. K. Roeder, C. Macal, and M. D. Rossetti, 3589–3598. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Chen, C. H., and L. H. Lee. 2010. *Stochastic Simulation Optimization: An Optimal Computing Budget Allocation*. World Scientific.
- Chen, C.-H., J. Lin, E. Yücesan, and S. Chick. 2000. "Simulation budget allocation for further enhancing the efficiency of ordinal optimization". *Discrete Event Dynamic Systems: Theory and Applications* 10:251–270.
- DeGroot, M. H. 2005. *Optimal Statistical Decisions*, Volume 82. John Wiley & Sons.
- Emmerich, M., and J. Klinkenberg. 2008. "The computation of the expected improvement in dominated hypervolume of Pareto front approximations". Technical report, Leiden University.
- Feldman, G., S. R. Hunter, and R. Pasupathy. 2015. "Multi-objective simulation optimization on finite sets: Optimal allocation via scalarization". In *Proceedings of the 2015 Winter Simulation Conference*, edited by L. Yilmaz, W. K. V. Chan, I. Moon, T. M. K. Roeder, C. Macal, and M. D. Rossetti, 3106–3107. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Frazier, P., and K. Kazachkov. 2011. "Guessing preferences: A new approach to multi-attribute ranking and selection". In *Proceedings of the 2011 Winter Simulation Conference*, edited by S. Jain, R. R. Creasey, J. Himmelspach, K. P. White, and M. Fu, 4319 – 4331. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- He, D., S. E. Chick, and C.-H. Chen. 2007. "Opportunity cost and OCBA selection procedures in ordinal optimization for a fixed number of alternative systems". *IEEE Transactions on Systems, Man, and Cybernetics, Part C* 37 (5): 951 – 961.
- Hu, S., and S. Andradottir. 2014. "A penalty function approach for simulation optimization with stochastic constraints". In *Proceedings of the 2014 Winter Simulation Conference*, edited by A. Tolk, S. Y. Diallo, I. O. Ryzhov, L. Yilmaz, S. Buckley, and J. A. Miller, 3730–3736. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Hunter, S. R., and G. Feldman. 2015. "Optimal sampling laws for bi-objective simulation optimization on finite sets". In *Proceedings of the 2015 Winter Simulation Conference*, edited by L. Yilmaz, W. K. V. Chan, I. Moon, T. M. K. Roeder, C. Macal, and M. D. Rossetti, 3749–3757. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Kim, S., and J.-H. Ryu. 2011. "The sample average approximation method for multi-objective stochastic optimization". In *Proceedings of the 2011 Winter Simulation Conference*, edited by S. Jain, R. R. Creasey, J. Himmelspach, K. P. White, and M. Fu, 4021 – 4032. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Lee, L. H., E. P. Chew, and S. Teng. 2007. "Finding the Pareto set for multi-objective simulation models by minimization of expected opportunity cost". In *Proceedings of the 2007 Winter Simulation Conference*, edited by S. G. Henderson, B. Biller, M.-H. Hsieh, J. Shortle, J. D. Tew, and R. R. Barton, 513 – 521. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers.
- Lee, L. H., E. P. Chew, and S. Teng. 2010. "Computing budget allocation rules for multi-objective simulation models based on different measures of selection quality". *Automatica* 46 (12): 1935–1950.

- Lee, L. H., E. P. Chew, S. Teng, and D. Goldsman. 2010. "Finding the non-dominated Pareto set for multi-objective simulation models". *IEEE Transactions* 42 (9): 656–674.
- Marceau-Caron, G., and M. Schoenauer. 2014. "Racing Multi-objective selection probabilities". In *Parallel Problem Solving from Nature*, edited by T. Bartz-Beielstein et al., Volume 8672 of *LNCS*, 631–640: Springer.
- Pasupathy, R., S. R. Hunter, N. A. Pujowidianto, L. H. Lee, and C. Chen. 2014. "Stochastically constrained ranking and selection via SCORE". *ACM Transactions on Modeling and Computer Simulation* 25 (1): 1–26.
- Ponweiser, W., T. Wagner, D. Biermann, and M. Vincze. 2008. "Multiobjective optimization on a limited budget of evaluations using model-assisted S-Metric selection". In *Parallel Problem Solving from Nature*, edited by G. Rudolph et al., Volume 5199 of *LNCS*, 784–794: Springer.
- Teng, S., L.-H. Lee, and E. P. Chew. 2010. "Integration of indifference-zone with multi-objective OCBA". *European Journal of Operational Research* 203:419–429.
- Zhang, T., M. Georgiopoulos, and G. C. Anagnostopoulos. 2013. "S-Race: A multi-objective racing algorithm". In *Genetic and Evolutionary Computation Conference*, edited by C. Blum, 1565–1572. ACM.

AUTHOR BIOGRAPHIES

JUERGEN BRANKE is Professor of Operational Research and Systems of Warwick Business School, University of Warwick, UK. He is Associate Editor for *IEEE Transaction on Evolutionary Computation*, Associate Editor for the *Evolutionary Computation Journal*, and Area Editor for the *Journal of Heuristics*. His research interests include metaheuristics, multiobjective optimisation and decision making, optimisation in the presence of uncertainty, and simulation-based optimisation, and he has published over 150 peer-reviewed papers in international journals and conferences. His e-mail address is Juergen.Branke@wbs.ac.uk.

WEN ZHANG is a PhD student of the Operational Research and Management Science Group at Warwick Business School, University of Warwick. Her research interests mainly lie in multi-objective optimisation and decision making. Her email address is phd14wz@mail.wbs.ac.uk.

YANG TAO is a PhD student of School of Engineering and Computing Sciences, University of Durham. His research interests mainly lie in computational fluid mechanics and asymptotic analysis. His email address is yang.tao@dur.ac.uk.