

MOI-MBO: Multiobjective Infill for Parallel Model-Based Optimization

Bernd Bischl¹, Simon Wessing², Nadja Bauer¹, Klaus Friedrichs¹, and
Claus Weihs¹

¹ Department of Statistics, TU Dortmund, Germany

`{bischl,bauer,friedrichs,weihs}@statistik.tu-dortmund.de`

² Department of Computer Science, TU Dortmund, Germany
`simon.wessing@tu-dortmund.de`

Abstract. The aim of this work is to compare different approaches for parallelization in model-based optimization. As another alternative aside from the existing methods, we propose using a multi-objective infill criterion that rewards both the diversity and the expected improvement of the proposed points. This criterion can be applied more universally than the existing ones because it has less requirements. Internally, an evolutionary algorithm is used to optimize this criterion. We verify the usefulness of the approach on a large set of established benchmark problems for black-box optimization. The experiments indicate that the new method's performance is competitive with other batch techniques and single-step EGO.

1 Introduction

Efficient optimizers that work on a strictly reduced budget of function evaluations are crucial for parameter optimization of expensive black-box functions, such as industrial simulators or time-consuming algorithms. Classical optimization methods in design of experiments are based on the assumption of a simple (often linear or quadratic) relationship between input parameters and performance output. In this case, the optimal set of evaluation points to fit such a model can usually be specified in advance.

However, for computer experiments, these simple models often do not suffice, as their assumptions are often severely violated, leading to unsatisfying results if they are employed nevertheless. Therefore, general sequential model-based optimization (MBO) is a standard technique for cost expensive simulations nowadays. Here, evaluation points are proposed sequentially using an appropriate surrogate model that allows for nonlinear relationships. After defining an initial set of evaluation points, e. g., a space-filling latin hypercube design, the basic procedure of MBO is an iterating loop of the following steps. Firstly, a model is fitted on the evaluated points, secondly, a new evaluation point is proposed by an infill criterion and lastly, its performance is evaluated.

In the last decade, many MBO procedures were proposed and compared relying on kriging models. Particularly in the context of algorithm configuration, also

random forest surrogate models received attention recently, due to their capability to handle categorical parameters (see SMAC, [1]). Instead of just evaluating one point in each iteration, a batch-sequential extension, which enables parallel evaluating of several points, is natural due to modern multi-core architectures. While there are already some recent approaches for multi-point MBO, the idea proposed in this paper is using a new multiobjective perspective by considering multiple infill criteria simultaneously without aggregating them into a single criterion. Instead of searching for one optimal evaluation point, an approximate Pareto front of q points is generated, which contains a spectrum of trade-offs between the different criteria. Possible single-objective infill criteria are mean prediction or local uncertainty of the surrogate model, and we also consider distances of the points in the current batch to each other to ensure diverse new evaluations.

In Sect. 2, MBO is defined more formally and state-of-the-art approaches are described including several ones for parallel MBO. Our proposal using multiobjectivization is explained in Sect. 3. In Sect. 4, the conducted comparison experiments are described, and in Sect. 5, their results are presented and interpreted. Finally, in Sect. 6, the most important findings are summarized.

2 Model-based Optimization

2.1 The Basic Sequential Algorithm

Let us assume that we aim to minimize an expensive black-box function $f : \mathcal{X} \subset \mathbb{R}^d \rightarrow \mathbb{R}$, $f(\mathbf{x}) = y$, $\mathbf{x} = (x_1, \dots, x_d)^T$. Each x_i is a continuous parameter with box constraints $[\ell_i, u_i]$, $\mathcal{X} = [\ell_1, u_1] \times \dots \times [\ell_d, u_d]$ is the parameter space of \mathbf{x} , and y is the target value. An important distinction is whether we are observing y deterministically or corrupted due to stochastic fluctuation, i. e., whether our observed target values are actually realizations of a random variable. In this paper, we will only study the noiseless, deterministic case. With $\mathcal{D} = (\mathbf{x}_1, \dots, \mathbf{x}_n)^T$, we will later denote an indexed set (design) of n different points $\mathbf{x}_i \in \mathcal{X}$ and $\mathbf{y} = (f(\mathbf{x}_1), \dots, f(\mathbf{x}_n))^T$, the vector of associated target values.

The main idea in model-based optimization is to approximate the expensive function $f(\mathbf{x})$ in every iteration by a regression model, which is much cheaper to evaluate. This is also called a meta-model or surrogate. Such a regression model often not only provides a direct estimation $\hat{f}(\mathbf{x})$ of the true function value $f(\mathbf{x})$, but also an estimation of the prediction standard error $\hat{s}(\mathbf{x})$, also called a local uncertainty measure. This value allows to assess the “trustworthiness” of the prediction, or, in a more Bayesian terminology, the spread of the posterior distribution of $\hat{f}(\mathbf{x})$.

An outline of sequential model-based optimization (MBO) is given in Alg. 1. We start by exploring the parameter space with an initial design, often constructed in a space-filling fashion. The main sequential loop can be divided into two alternating stages: Fitting the response surface to the currently available design data, then optimizing the so-called infill criterion to propose a new promising point \mathbf{x}^* for the next expensive evaluation $f(\mathbf{x}^*)$.

Algorithm 1: Sequential model-based optimization

```

1 Generate an initial design  $\mathcal{D} \subset \mathcal{X}$ ;
2 Compute  $\mathbf{y} = f(\mathcal{D})$ ;
3 while total evaluation budget is not exceeded do
4   Fit surrogate on  $\mathcal{D}$  and obtain  $\hat{f}$ ,  $\hat{s}$ ;
5   Get new design point  $\mathbf{x}^*$  by optimizing the infill criterion based on  $\hat{f}$ ,  $\hat{s}$ ;
6   Evaluate new point  $\mathbf{y}^* = f(\mathbf{x}^*)$ ;
7   Update:  $\mathcal{D} \leftarrow (\mathcal{D}, \mathbf{x}^*)$  and  $\mathbf{y} \leftarrow (\mathbf{y}, \mathbf{y}^*)$ ;
8 return  $y_{min} = \min(\mathbf{y})$  and the associated  $\mathbf{x}_{min}$ .

```

Quite a few infill criteria exist, both for the deterministic and noisy case. They are usually constructed pointwise by combining $\hat{f}(\mathbf{x})$ and $\hat{s}(\mathbf{x})$ in a certain way. For the former, lower values are more promising, as they indicate a low true function value $f(\mathbf{x})$. For the latter, higher values indicate less explored regions of the search space, as our model is less certain about the true landscape, usually because it lacks training points nearby. The task of many infill criteria is to balance these two conflicting criteria into one numerical formula.

Probably the simplest infill criterion, which can be used even if no local uncertainty estimator is available, is just considering the mean prediction $\hat{f}(\mathbf{x})$. This results in a greedy behavior, where promising regions are exploited at once and can quickly result in only local convergence.

In his seminal paper, Jones [2] recommended the following choices: Kriging [3], i. e., a Gaussian process should be used for regression. This model can fit multimodal landscapes with satisfying quality, even when only a low amount of data points are available. As a kernel method, it also offers flexibility, and roughness information regarding the target function can be encoded into the model via the covariance kernel. The posterior distribution of $\hat{f}(\mathbf{x})$ is now a univariate $N(\hat{f}(\mathbf{x}), \hat{s}(\mathbf{x})^2)$ one. Based on this, the now standard *expected improvement* (EI) criterion was proposed, which supposedly ensures global convergence [4–6]. It is defined as

$$\begin{aligned} \text{EI}(\mathbf{x}) &= \mathbb{E}[\max\{0, y_{min} - \hat{f}(\mathbf{x})\}] \\ &= (y_{min} - \hat{f}(\mathbf{x})) \Phi\left(\frac{y_{min} - \hat{f}(\mathbf{x})}{\hat{s}(\mathbf{x})}\right) + \hat{s}(\mathbf{x}) \phi\left(\frac{y_{min} - \hat{f}(\mathbf{x})}{\hat{s}(\mathbf{x})}\right), \end{aligned}$$

where ϕ and Φ are the density and cumulative distribution function of the standard normal distribution, respectively. Hence, the sought point is $\mathbf{x}^* = \arg \max_{\mathbf{x} \in \mathcal{X}} \text{EI}(\mathbf{x})$.

One further infill criterion which is relevant for this paper is the *lower confidence bound* (LCB) criterion. LCB combines the predicted mean response with the estimated standard error through a weighted sum:

$$\mathbf{x}^* = \arg \min_{\mathbf{x} \in \mathcal{X}} \text{LCB}(\mathbf{x}) = \arg \min_{\mathbf{x} \in \mathcal{X}} \hat{f}(\mathbf{x}) - \lambda \hat{s}(\mathbf{x}).$$

The weighting parameter λ has to be selected by the user and wrong choices will not guarantee global convergence. Obviously, the LCB criterion has to be minimized. If $\lambda = 0$, LCB coincides with predicted mean value. The larger λ is chosen, the more attractive unexplored regions of the search space become. Further infill criteria are discussed in [5].

2.2 Review of Parallel MBO Strategies

In the standard MBO procedure, only one point \mathbf{x}^* is proposed in each sequential step (see line 5 of Alg. 1). However, often dozens or hundreds of CPU cores are simultaneously available nowadays. In some cases, these processors can be used to parallelize a single function call of the simulator $f(\mathbf{x}^*)$, but in many cases this is not possible and such an evaluation has to be considered “atomic”. It is hence essential to exploit this computational potential in a reasonably efficient way, suggesting a batch-sequential approach. The aim is to propose q (instead of only 1) new points $\mathbf{x}_1^*, \dots, \mathbf{x}_q^*$ in each iteration so that q expensive function evaluations can be performed in parallel. Some possibilities have already been proposed to achieve this parallelization. However, they all introduce drawbacks that we want to avoid.

Arguably the mathematically most intuitive way to tackle the problem is to directly extend the EI-criterion for q points (q -EI) [7]. While for the 2-EI case an analytic solution is provided [7], for $q > 2$ an expensive Monte-Carlo simulation was implemented in [7] and [8]. Recently, Chevalier and Ginsbourger [9] proposed an analytic approach to more efficiently compute the q -EI for moderate values of $q \leq 10$.

A simple alternative for multi-point proposal using the EI-criterion is given by [7], where in the first step the kriging model is fitted based on the real data and \mathbf{x}_1^* is calculated according to the regular EI-criterion. Then, for $i = 1, \dots, q - 1$, a simple “guess” for $f(\mathbf{x}_i^*)$ is used to update the model in order to propose the subsequent point. Such estimation can be done due to $\hat{f}(\mathbf{x})$ or even any constant (like y_{min} or y_{max}) may be employed. In the first case, this approach is called *kriging believer* and in the second case *constant liar*. One should note that although the expensive re-estimation of the covariance kernel parameters is not performed during the batch generation but only after its batch evaluation (although without formal justification), the EI still needs to be optimized sequentially q times, which can in itself be very time consuming for higher values of d and q .

Hutter et al. [10] introduce another strategy. They use the LCB criterion and sample a new λ value from an exponential distribution with mean 1 for each point in the batch. This defines q different infill objective functions, one for each desired new point, each one – in principle – encoding a different trade-off between mean and standard error prediction of the model. The simple, independent optimization of the single-objective LCB criteria (which can also be performed in parallel) hence leads to a batch of q promising points. While this approach scales computationally very well with increasing q and is also extremely simple

to implement, there is no guaranteed diversity in the generated batch, so different values for λ can still lead to very similar global optima locations in the LCB landscape.

The aim of this article is to propose and compare other strategies for multi-point MBO, based on multicriteria evolutionary algorithms. The next section will detail our approach.

3 Proposal

Multiobjective optimization considers itself with the task of finding Pareto-optimal solutions to a set of objective functions. This methodology has become a standard tool under the belt of many researchers in black-box optimization, and a lot of excellent overview literature exists. As an introduction, we would like to refer the reader to [11].

Infill criteria for model-based optimization often consist of two conflicting parts: the mean response $\hat{f}(\mathbf{x})$ and the local uncertainty estimation $\hat{s}(\mathbf{x})$. In the case of LCB, it is very obvious that these two functions actually constitute a bicriteria problem. They have simply been scalarized into a weighted sum, where setting the actual weighting parameter is left up to the experimenter. This is usually a hard task if no further information is provided, and the best setting might not even be a constant value, but instead depend on the stage of the optimization process. For EI one might argue that a formal motivation exists which dictates the specific formula. But this derivation assumes that all model assumptions of the Gaussian process hold, which might not be true in practice. Furthermore, one might want to transfer the principle to other non-parametric regression models, e. g., because of discrete parameters in the input space or due to the disadvantageous runtime scaling behavior of the kriging model fits. This is already done in the case of SMAC, where a random forest instead of a kriging model is used in combination with an EI criterion, but there is no reason to assume that the posterior distribution of forest predictions is really normally distributed or that the derivation of the EI criterion is still valid.

Instead of dealing with the hassle of appropriately aggregating two or more objective functions into a single one (as, e. g., for EI or LCB), we simply accept the multiobjective nature of the problem and use a multiobjective optimizer. As we want to employ parallelization, we focus on a posteriori methods that return not only a single point, but a set of points \mathcal{P} that approximate the Pareto-set. We therefore choose $|\mathcal{P}| = q$, the number of points we want to process in parallel. In this case, it also suggests itself to use evolutionary algorithms (EA), as they are population-based approaches and can be used easily for multiobjective optimization.

So far, we have not discussed the fact that we need a set of q *distinct* points. Therefore, we add the distance to neighboring solutions as another objective, to respect the influence of the points on each other. By maximizing this distance, points are rewarded for being dissimilar to other solutions. The objective function is thus dynamic, i. e., it depends on the EA's population and may change in each

iteration of the optimization. If the distance is added as an artificial, secondary objective to a single-criterion problem, this approach is known under the name of multiobjectivization. The reason to proceed in such a fashion is that one might want to obtain a whole spectrum of promising, diverse solutions in the decision space, which is exactly what we want to achieve for multipoint proposal in MBO. This approach does not guarantee to result in a simpler problem, as shown by Brockhoff et al. [12], but gave decent results in practical applications similar to ours [13–15]. Summing up, the following objective functions are available for building our multiobjective infill criterion:

- **mean**: mean model prediction $\hat{f}(\mathbf{x})$,
- **se**: standard error / model uncertainty $\hat{s}(\mathbf{x})$,
- **ei**: expected improvement $\text{EI}(\mathbf{x})$,
- **dist.nn**: distance to the nearest neighbor of \mathbf{x} in the current population

$$\text{dist}_{\text{nn}}(\mathbf{x}, \mathcal{P}) = \min\{d(\mathbf{x}, \tilde{\mathbf{x}}) \mid \tilde{\mathbf{x}} \in \mathcal{P} \setminus \{\mathbf{x}\}\}, \mathcal{P} \subset \mathcal{X},$$

- **dist.nb**: distance to the nearest better neighbor of \mathbf{x} ,

$$\text{dist}_{\text{nb}}(\mathbf{x}, \mathcal{P}) = \min\{d(\mathbf{x}, \tilde{\mathbf{x}}) \mid f(\tilde{\mathbf{x}}) < f(\mathbf{x}) \wedge \tilde{\mathbf{x}} \in \mathcal{P}\}, \mathcal{P} \subset \mathcal{X}.$$

The latter distance definition is considered because it proved successful in experiments by Wessing et al. [15]. The distance measure $d(\mathbf{x}, \tilde{\mathbf{x}})$ used throughout this paper is simply the euclidean distance $\|\mathbf{x} - \tilde{\mathbf{x}}\|_2$. As there is in every set at least one solution \mathbf{x}^* that has no nearest better neighbor, we define $\text{dist}_{\text{nb}}(\mathbf{x}^*, \mathcal{P}) := \infty$.

Our multi-point MBO approach is geared to the evolutionary multiobjective optimization algorithm proposed by Beume et al. [16]. Alg. 2 introduces the resulting optimization procedure.

Function $\text{crossover}(\mathbf{x}, \tilde{\mathbf{x}}, \mathcal{X}, \eta_c, p_c)$ is the simulated binary crossover operator with a distribution index η_c and application probability p_c . The latter parameter specifies the probability to apply the operator at each variable. crossover produces one offspring from two parents. The function $\text{mutate}(\mathbf{x}, \mathcal{X}, \eta_m, p_m)$ is the polynomial mutation operator, which is applied to the offspring. Again, η_m denotes the distance parameter of the distribution and p_m the mutation probability for each variable. In this work, we fix these parameters to $\eta_c = \eta_m = 15$ and $p_c = p_m = 1$.

Lines 9–14 describe the survivor selection. After the distance values have been incorporated, a non-dominated sorting of all individuals is conducted. Then the worst individual of the worst non-dominated front is identified, according either to the hypervolume contribution [16] or to one of the single-objective functions. This individual is then removed from the population. We will call the former selection criterion **hv**. In the latter method, the first single infill criterion in the name – which is also the most important one – is always used for the selection, so that we call it **first** in the following.

Table 1 introduces our multiobjective infill (MOI-)MBO strategies. Note, that the two last approaches simply implement a multiobjectivization of the LCB criterion and hence do not use any distance measure. In this case, lines 9–10 of Alg. 2 should be ignored.

Algorithm 2: Evolutionary optimization of multiobjective infill criteria

```

1 Generate an initial population  $\mathcal{P} = \{\mathbf{x}_1, \dots, \mathbf{x}_\mu\} \subset \mathcal{X}$ ;
2 Evaluate  $\mathcal{P}$ ;
3 while number of iterations is not exceeded do
4   Sample two individuals (parents):  $\mathbf{x}_{p1} \in \mathcal{P}$  and  $\mathbf{x}_{p2} \in \mathcal{P}$ ;
5   Generate a new individual (child):  $\mathbf{x}_{ch} = \text{crossover}(\mathbf{x}_{p1}, \mathbf{x}_{p2}, \mathcal{X}, \eta_c, p_c)$ ;
6   Mutate the new individual:  $\mathbf{x}_{ch} := \text{mutate}(\mathbf{x}_{ch}, \mathcal{X}, \eta_m, p_m)$ ;
7   Evaluate selected infill criteria (except distance) for  $\mathbf{x}_{ch}$ ;
8   Update the current population:  $\mathcal{P} := \mathcal{P} \cup \{\mathbf{x}_{ch}\}$ ;
9   for  $\mathbf{x} \in \mathcal{P}$  do
10    | Calculate  $\text{dist}(\mathbf{x}, \mathcal{P})$  and append to objective values;
11   Compute non-dominated fronts  $\mathcal{F}_1, \dots, \mathcal{F}_k$  of  $\mathcal{P}$ ;
12   Sort  $\mathcal{F}_k$  by a selection criterion;
13    $\mathbf{x}_{worst}$  = last element of  $\mathcal{F}_k$ ;
14   Update the current population:  $\mathcal{P} := \mathcal{P} \setminus \{\mathbf{x}_{worst}\}$ ;
15 return  $\mathcal{P}$ ;

```

4 Comparison Experiments

In order to study the performances of all variants of our proposed MOI-MBO and the existing strategies for parallel multi-point infill, we conduct an extensive benchmark study. We also compare against the usual EGO 1-step algorithm – which of course works in a purely sequential, non-parallel fashion and is therefore able to exploit more information during the optimization. Finally, we also include a simple random search as a baseline comparison method.

Problem Instances and Budget. As problem instances for the benchmark, we selected all 24 test functions of the black-box optimization benchmark (BBOB) noise-free test suite [17]. It covers simple unimodal, ill-conditioned and multimodal functions. Some of the landscapes of these functions exhibit a strong global structure which a model-based optimizer can potentially exploit, while other functions do not have this characteristic. We study these functions in the dimensions $d \in \{5, 10\}$. For every function (and each dimension) we create 10 initial designs of size $5 \cdot d$. We then run each candidate optimizer 10 times, using the mentioned designs. Hence, this results in 10 statistical replications for each problem instance with fair and equal starting conditions for all competing optimizers. The optimizers are allowed an additional budget of $40 \cdot d$ function evaluations on top of the initial design. Parallel optimizers always propose batches of size $q = 5$ in our experiments, resulting in $8 \cdot d$ sequential iterations for them. As a quality measure, we choose the difference between the function values of the best obtained point during the optimization and the known global minimum.

EGO and Constant Liar Variants. To avoid interferences in our performance measurements due to how the infill criterion of the surrogate is optimized

Table 1. Proposed MOI-MBO approaches

Single Infill Criterion				Sel. Criterion		Abbreviation	
mean	se	ei	dist		first		hv
			nn	nb			
×	×		×		×		<code>moi_mean.se.dist_nn_first</code>
×	×			×	×		<code>moi_mean.se.dist_nb_first</code>
×	×		×			×	<code>moi_mean.se.dist_nn_hv</code>
×	×			×		×	<code>moi_mean.se.dist_nb_hv</code>
		×	×		×		<code>moi_ei.dist_nn_first</code>
		×		×	×		<code>moi_ei.dist_nb_first</code>
		×	×			×	<code>moi_ei.dist_nn_hv</code>
		×		×		×	<code>moi_ei.dist_nb_hv</code>
×	×				×		<code>moi_mean.se_first</code>
×	×					×	<code>moi_mean.se_hv</code>

(in EGO this is usually a sophisticated combination of a gradient-based and an evolutionary/restart mechanism), we reimplemented EGO with a simple $(\mu + 1)$ evolutionary algorithm with simulated binary crossover and polynomial mutation as variation operators for optimization of the infill criterion. Here, we set $\mu = 20$. We also include an EGO variant where we use the mean value $\hat{f}(\mathbf{x})$ as a simple infill criterion.

Software. All of our experiments are conducted in the statistical programming language R. The BBOB test functions are made available in the R package `soobench` [18]. We have implemented all our code regarding model-based optimization (including the below mentioned evolutionary variants of 1-step EGO and constant liar, all multicriteria methods and parallel LCB) in the experimental R package `mlrMBO`³. The toolbox allows a generic combination of regression models and optimization strategies and builds upon the `mlr` R package for machine learning from Bischl [19]. The kriging models are fitted via the `DiceKriging` package and we compare against the EGO and constant liar implementations of the popular `DiceOptim` package implementation. Both Dice packages are published by Ginsbourger et al. [7].

Summary. In order to provide a succinct overview, we again list all compared approaches in Table 2.

5 Results

As the test functions differ w. r. t. their respective characteristics, we divide the BBOB function in the following three sets, using the same numbering as in [17]:

³ See <https://github.com/berndbischl/mlrMBO>.

Table 2. Overview of compared model-based optimization strategies.

Method	Abbreviation	R Package	Infill Optimizer
EGO	<code>ego</code>	<code>DiceOptim</code>	gradient-based <code>rgeoud</code>
EGO	<code>ego_ea_ei</code>	<code>mlrMBO</code>	simple EA
EGO Infill $\hat{f}(\mathbf{x})$	<code>ego_ea_mean</code>	<code>mlrMBO</code>	simple EA
Constant Liar	<code>par_cl</code>	<code>DiceOptim</code>	gradient-based <code>rgeoud</code>
Constant Liar	<code>par_cl_ea</code>	<code>mlrMBO</code>	simple EA
Parallel LCB	<code>par_lcb</code>	<code>mlrMBO</code>	simple EA per LCB function
Multiobj. Infill (see Table 1)	<code>moi</code>	<code>mlrMBO</code>	multicrit EA
Random Search	<code>random_search</code>	–	–

set1 Unimodal functions: 1, 2, 5–14,

set2 Multimodal functions with adequate global structure: 3, 4, 15–19,

set3 Multimodal functions with weak global structure: 20–24.

To interpret the results, we first rank the considered optimization approaches among each other for each dimension / test function / replication by their performance and subsequently calculate the mean rank for each approach across test functions and replications (per dimension). Thus, the lower the mean rank, the better the approach. We also conduct statistical sign tests to analyse significant differences in performance between candidate algorithms. Here, we compute for each pair of optimization approaches the performance differences over all replications for a set of test functions (per dimension). The sign distribution of this difference vector is then used to decide whether one approach significantly outperforms the other, which is in essence a binomial test, see [20]. Each test is conducted at the 5% significance level without further adjustment for multiple testing, as our aim is to use the test as an exploratory tool to provide a descriptive visualization of the stochastic results.

We illustrate the results using preference relation graphs, containing two kinds of information: the mean ranks of the optimization approaches on the one hand and the decisions of pairwise sign tests on the other hand. The results are presented in Fig. 1–4. Each colored node represents an optimization approach, where its mean rank is given in braces. The main goal is to compare the different MOI-MBO strategies (yellow nodes) with other parallel approaches (red nodes). Of these, our most relevant competitor is LCB, as it scales a lot better with increasing sizes of q than constant liar. For comparison, also non-parallel approaches are considered (green nodes). But these purely sequential optimizers are expected to perform better than the parallel ones, as they benefit from a higher number of surrogate model fits and can incorporate more information into the models. Two nodes are connected with an edge if one approach (the upper) is significantly better than the other (the lower). Note, it is possible that one approach is significantly better than another one, although it has a (slightly) worse mean rank.

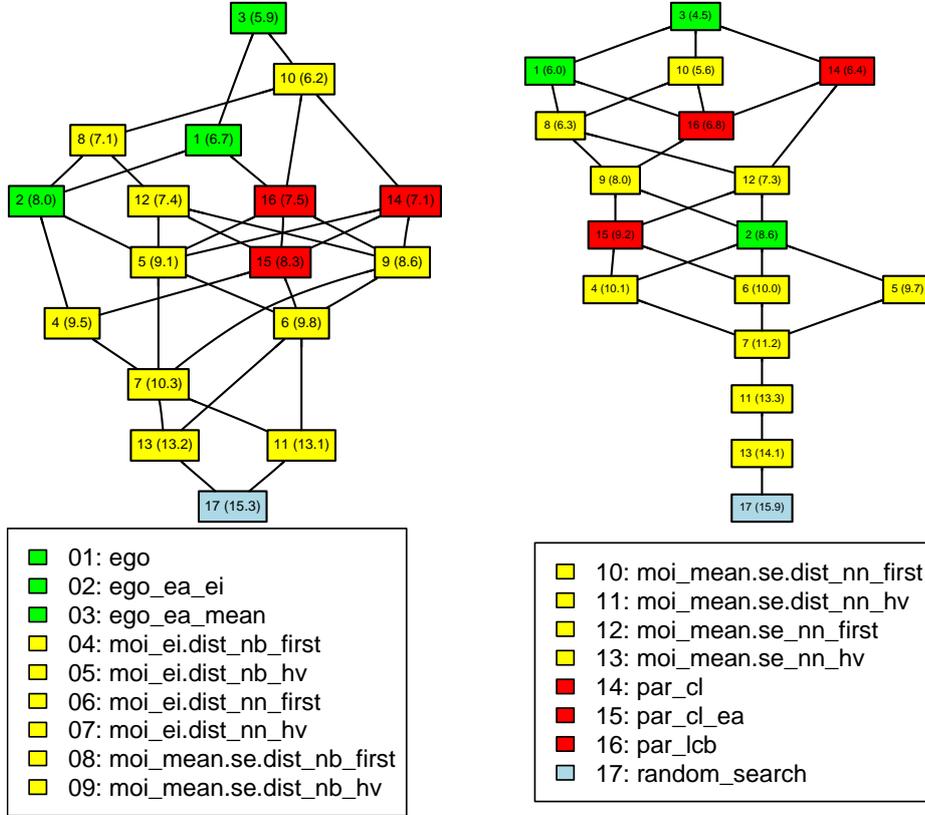
Fig. 1. Comparison over all test functions ($d = 5$ and $d = 10$).

Figure 1 illustrates the results regarding all test functions. Overall, the best strategy is `ego_ea_mean`, even outperforming `ego` and `ego_ea_ei`. This is somewhat surprising, particularly as the same holds true for the set of multimodal functions (Fig. 3 and 4), where exploration of the parameter space is expected to be more beneficial than just exploitation of the surrogate function.

Over all test functions, the best parallel approach seems to be `moi_mean.se.dist_nn_first`, which has the best mean rank and – with one exception – outperforms all other parallel strategies according to the sign test. Only slightly worse perform the strategies `moi_mean.se.dist_nb_first` and `moi_mean.se.first`, which outperform all other considered MOI-MBOs. Since all three methods just differ in the applied distance criterion, this seems to have only a relatively weak influence. In contrast, the graphs indicate significant improvements by using the selection criterion `first` – which focuses more on the first criterion `mean` or `ei` – instead of `hv`. Regarding our experiments, the best state-of-the-art parallel MBO is `par_cl`, although this holds in a strong sense (significance of the sign test when we compare it to `par_lcb`) only for $d = 10$. While `par_cl` performs

Fig. 2. Comparison over test functions of set1 ($d = 5$ and $d = 10$).

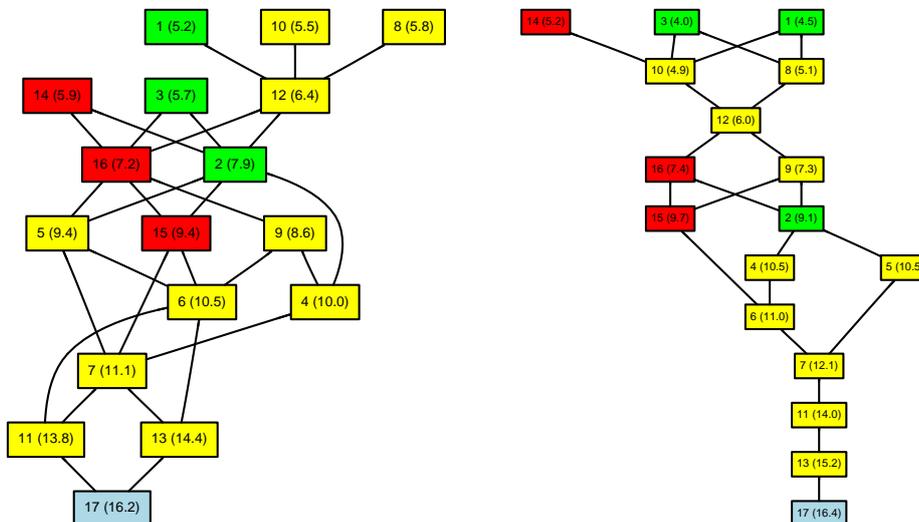
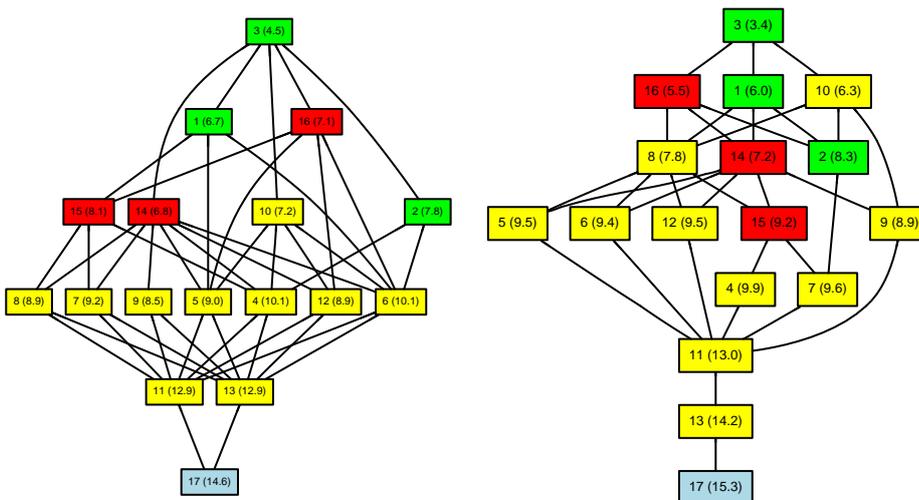
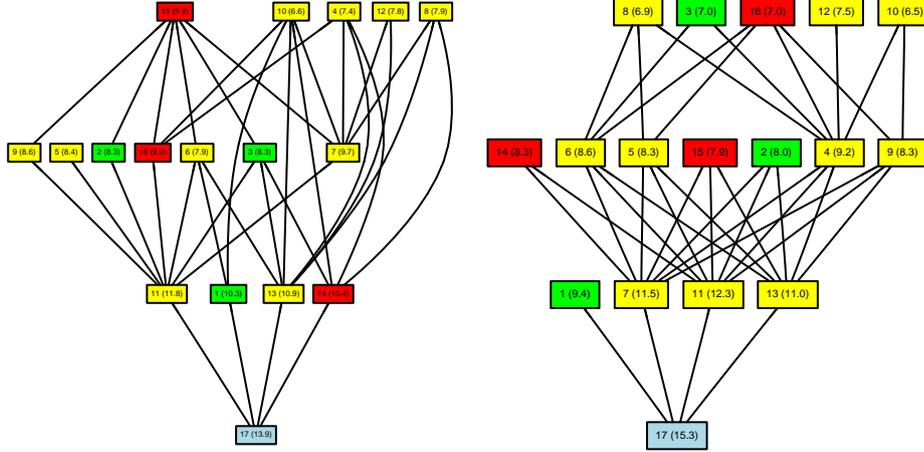


Fig. 3. Comparison over test functions of set2 ($d = 5$ and $d = 10$).



approximately as good as the best MOI-MBO `moi_mean.se.dist_nn_first` for the unimodal and the multimodal functions with adequate global structure (Fig. 2 and 3), our new approach performs better on multimodal functions with weak global structure (Fig. 4). Contrary, `par_cl_ea` and `par_lcb` – the other two considered state-of-the-art parallel MBOs – perform quite weakly on unimodal functions (Fig. 2).

Fig. 4. Comparison over test functions of set3 ($d = 5$ and $d = 10$).

6 Conclusion

In this paper, a multiobjective approach for parallel model-based optimization was introduced. Therefore, ten different strategies – each relying on a reasonable subset of five infill criteria – were compared with several state-of-the-art approaches including EGO, parallel LCB and constant liar. While three of the infill criteria have been applied before, the concept of multiobjectivization and the consideration of the distance to neighboring points is a new approach in the context of MBO. Regarding the 24 considered test functions in five and ten dimensions, a MOI-MBO strategy, which relies on the mean model prediction $\hat{f}(\mathbf{x})$, the model uncertainty $\hat{s}(\mathbf{x})$ and the distance to the nearest neighbor as infill criteria, performs best on average. As shown in the previous section it even outperforms existing parallel methods in many situations. Additionally, its runtime behavior is not significantly inhibited if the number of cores q is increased, as this only results in a larger population size of the EA. Furthermore, the experiment shows a bias in favor of more exploitative methods versus more explorative ones, although this might be an artefact of the considered benchmark set.

In future comparison studies, also the recent approach of Chevalier and Ginsbourger [9] should be considered once their code is released. While in the experiments above all approaches are applied on only five CPU cores ensuring acceptable run times even for the most complex ones, in a next step the influence of the number of cores regarding the performance should be analyzed. Furthermore, we also would like to investigate the applicability of MOI-MBO in the noisy case.

Acknowledgements. This paper is based on investigations of the projects B3 and C2 of the Collaborative Research Center SFB 823, which are kindly

supported by Deutsche Forschungsgemeinschaft (DFG). The authors also thank Tobias Wagner for fruitful discussions of multiobjective infill criteria.

References

1. Hutter, F., Hoos, H.H., Leyton-Brown, K.: Sequential model-based optimization for general algorithm configuration. In Coello, C.A.C., ed.: *Learning and Intelligent Optimization*. Volume 6683 of *Lecture Notes in Computer Science*. Springer (2011) 507–523
2. Jones, D.R., Schonlau, M., Welch, W.J.: Efficient global optimization of expensive black-box functions. *Journal of Global Optimization* **13**(4) (1998) 455–492
3. Krige, D.: A statistical approach to some basic mine valuation problems on the witwatersrand. *Journal of the Chemical, Metallurgical and Mining Society of South Africa* **52**(6) (1951) 119–139
4. Locatelli, M.: Bayesian algorithms for one-dimensional global optimization. *Journal of Global Optimization* **10**(1) (1997) 57–76
5. Jones, D.R.: A taxonomy of global optimization methods based on response surfaces. *Journal of Global Optimization* **21**(4) (2001) 345–383
6. Vazquez, E., Bect, J.: Convergence properties of the expected improvement algorithm with fixed mean and covariance functions. *Journal of Statistical Planning and Inference* **140**(11) (2010) 3088–3095
7. Ginsbourger, D., Le Riche, R., Carraro, L.: Kriging is well-suited to parallelize optimization. In Tenne, Y., Goh, C.K., eds.: *Computational Intelligence in Expensive Optimization Problems*. Volume 2 of *Adaptation, Learning, and Optimization*. Springer (2010) 131–162
8. Janusevskis, J., Le Riche, R., Ginsbourger, D., Girdziusas, R.: Expected improvements for the asynchronous parallel global optimization of expensive functions: Potentials and challenges. In Hamadi, Y., Schoenauer, M., eds.: *Learning and Intelligent Optimization*. *Lecture Notes in Computer Science*. Springer (2012) 413–418
9. Chevalier, C., Ginsbourger, D.: Fast computation of the multi-points expected improvement with applications in batch selection. In Nicosia, G., Pardalos, P., eds.: *Learning and Intelligent Optimization*. *Lecture Notes in Computer Science*. Springer (2013) 59–69
10. Hutter, F., Hoos, H.H., Leyton-Brown, K.: Parallel algorithm configuration. In Hamadi, Y., Schoenauer, M., eds.: *Learning and Intelligent Optimization*. *Lecture Notes in Computer Science*. Springer (2012) 55–70
11. Coello Coello, C.A., Lamont, G.B., Van Veldhuizen, D.A.: *Evolutionary Algorithms for Solving Multi-Objective Problems*. Second edn. Springer (2007)
12. Brockhoff, D., Friedrich, T., Hebbinghaus, N., Klein, C., Neumann, F., Zitzler, E.: Do additional objectives make a problem harder? In: *Proceedings of the 9th Annual Conference on Genetic and Evolutionary Computation*. GECCO '07, ACM (2007) 765–772
13. Tran, T.D., Brockhoff, D., Derbel, B.: Multiobjectivization with NSGA-II on the Noiseless BBOB Testbed. In: *GECCO (Companion), workshop on Black-Box Optimization Benchmarking (BBOB'2013)*, Amsterdam, Pays-Bas (July 2013)
14. Ulrich, T., Thiele, L.: Maximizing population diversity in single-objective optimization. In: *Proceedings of the 13th annual conference on Genetic and evolutionary computation*. GECCO '11, New York, NY, USA, ACM (2011) 641–648

15. Wessing, S., Preuss, M., Rudolph, G.: Niching by multiobjectivization with neighbor information: Trade-offs and benefits. In: IEEE Congress on Evolutionary Computation (CEC). (2013) 103–110
16. Beume, N., Naujoks, B., Emmerich, M.: SMS-EMOA: Multiobjective selection based on dominated hypervolume. *European Journal of Operational Research* **181**(3) (2007) 1653–1669
17. Hansen, N., Finck, S., Ros, R., Auger, A.: Real-Parameter Black-Box Optimization Benchmarking 2009: Noiseless Functions Definitions. Technical Report RR-6829, INRIA (2009) <http://hal.inria.fr/inria-00362633/en>.
18. Mersmann, O., Bischl, B., Bossek, J., Judt, L.: soobench: Single Objective Optimization Benchmark Functions. R package version 1.1-164.
19. Bischl, B.: mlr: Machine Learning in R. R package version 1.2.
20. Conover, W.: Practical Nonparametric Statistics. Second edn. John Wiley, New York, NY (1980)