



HAL
open science

Walsh-based surrogate-assisted multi-objective combinatorial optimization: A fine-grained analysis for pseudo-boolean functions

Bilel Derbel, Geoffrey Pruvost, Arnaud Liefoghe, Sébastien Verel, Qingfu Zhang

► To cite this version:

Bilel Derbel, Geoffrey Pruvost, Arnaud Liefoghe, Sébastien Verel, Qingfu Zhang. Walsh-based surrogate-assisted multi-objective combinatorial optimization: A fine-grained analysis for pseudo-boolean functions. *Applied Soft Computing*, 2023, 136, pp.110061. 10.1016/j.asoc.2023.110061 . hal-04073811

HAL Id: hal-04073811

<https://hal.science/hal-04073811>

Submitted on 11 Sep 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Walsh-based Surrogate-assisted Multi-objective Combinatorial Optimization: A Fine-grained Analysis for Pseudo-boolean Functions*

Bilel Derbel^{a,*}, Geoffrey Pruvost^a, Arnaud Liefooghe^a, Sébastien Verel^b, Qingfu Zhang^c

^a*Univ. Lille, CNRS, Inria, Centrale Lille, UMR 9189 CRISTAL, F-59000 Lille, France*

^b*Univ. Littoral Côte d'Opale, UR 4491, LISIC, F-62100 Calais, France*

^c*City University of Hong Kong, Kowloon Tong, Hong Kong*

Abstract

The aim of this paper is to study surrogate-assisted algorithms for expensive multiobjective combinatorial optimization problems. Targeting pseudo-boolean domains, we provide a fine-grained analysis of an optimization framework using the Walsh basis as a core surrogate model. The considered framework uses decomposition in the objective space, and integrates three different components, namely, (i) an inner optimizer for searching promising solutions with respect to the so-constructed surrogate, (ii) a selection strategy to decide which solution is to be evaluated by the expensive objectives, and (iii) the strategy used to setup the Walsh order hyper-parameter. Based on extensive experiments using two benchmark problems, namely bi-objective NK-landscapes and unconstrained binary quadratic programming problems (UBQP), we conduct a comprehensive in-depth analysis of the combined effects of the considered components on search performance, and provide evidence on the effectiveness of the proposed search strategies. As a by-product, our work shed more light on the key challenges for designing a successful surrogate-assisted multi-objective combinatorial search process.

Keywords: Multi-objective optimization, discrete surrogates, decomposition

1. Introduction

Surrogate-assisted algorithms have been recognized as a mainstream methodology for expensive optimization where objective function evaluation is economically costly and/or time consuming [1–9]. These algorithms build and use surrogate models for estimating objective function values in order to reduce as many as possible the number of true expensive function evaluations. Although

*This work was supported by the French national research agency (ANR-16-CE23-0013-01) and the Research Grants Council of Hong Kong (RGC Project No. A-CityU101/16).

*Corresponding author

Email addresses: bilel.derbel@univ-lille.fr (Bilel Derbel), geoffrey.pruvost@univ-lille.fr (Geoffrey Pruvost), arnaud.liefooghe@univ-lille.fr (Arnaud Liefooghe), verel@univ-littoral.fr (Sébastien Verel), qingfu.zhang@cityu.edu.hk (Qingfu Zhang)



much effort has been made to develop surrogate-assisted algorithms, a number of difficult issues still need to be investigated [1]. In this paper, we are specifically concerned with the design of both *multi-objective* and *discrete* surrogate-assisted optimization algorithms. Interestingly, these two aspects, taken separately, have been identified in [1] as among of the five major challenges that require further in-depth investigations from the community.

Multi-objective optimization problems (MOP) [10, 11] are a natural outcome in different application fields where a number of conflicting objectives are to be optimized in a simultaneous manner. Our goal is to compute a representative set of solutions exposing different trade-offs in the objective space, i.e., a Pareto front approximation. Among the different techniques to tackle MOPs, evolutionary multi-objective algorithms (EMOA) [12, 13] have been proved to be extremely effective, especially due to their ability to evolve a whole set of solutions mapping the target approximation set. In an expensive optimization setting, EMOAs have been leveraged with surrogates in a number of different ways, and for different problems and application domains. Generally speaking, two major inter-dependent issues are considered: (i) the meta-model being used as a surrogate, and (ii) the class of the multi-objective technique used at the core of the underlying surrogate-assisted optimization process. On the one hand, standard statistical and machine-learning meta-models, ranging from kriging [14] and Gaussian process (GP) [15], to support vector regressions (SVR) [16], radial basis functions (RBF) [17], artificial neural networks (ANN) [18], decision trees and random forest (RF) [19, 20], etc, can be considered to learn a blackbox objective function. On the other hand, these models can be coupled with different evolutionary multi-objective search paradigms, such as decomposition [21, 22], dominance [23-25], and indicators [26, 27]; hence, ending-up with a variety of surrogate-assisted algorithms, such as Par-EGO [28], MOEA/D-EGO [29], MOEA/D-RBF [30], KRVEA [31], Multi-objective EGO [32], SMS-EGO [33], CSA-MOEA [34], CSEA [35], SGMOO [36], to cite a few. The existing literature is too vast to discuss in the scope of this paper. We hence refer the reader to the existing taxonomies [7, 8, 37] and surveys [2, 3, 6, 9] for a more thorough overview.

The multi-objective algorithms cited before were mainly intended to work with continuous decision variables. However, expensive optimization problems are not specific to continuous domains; in fact, many real-world applications are known to have discrete variables, e.g., strings, integers, graphs, permutations, etc. More importantly, continuous evolutionary approaches cannot be transferred as such to combinatorial domains. Firstly, given the difficulty of modeling the landscape of combinatorial problems, it is still not clear how to derive a discrete surrogate model that can accurately learn the structure of a combinatorial (single-objective) problem. This is to contrast with the fact that there exist well-established continuous surrogates with well-studied properties. However, one can find a number of investigations for designing discrete surrogates, ranging from naive approaches ignoring the discrete structure of the search space, to more domain-specific modeling techniques [3, 38-42]. In particular, some of the previously mentioned meta-models, such as GP, RBF, RF, ANN, can be adapted to support mixed (continuous and discrete) variables.

Secondly, assuming an accurate discrete surrogate can be used, it is not clear how to finely couple the surrogate model with the optimization process for an optimal performance. Such an issue is even more pressing for multi-objective combinatorial problems, given the variety of discrete spaces that can be considered, and the specific challenges underlying a multi-objective search. Finally, benchmarking surrogate-assisted combinatorial multi-objective algorithms is a challenging research question. Benchmarking is highlighted in [1] as *the* most important challenge, which is (quoting the authors in [1]) “an open issue throughout the whole field of surrogate-assisted optimization and a mandatory step towards identifying strengths and weaknesses of approaches”. In this respect, the systematic analysis of discrete surrogate-assisted approaches is even less developed compared to continuous domains, and a lot of effort has to be performed before ending up with a unified view of what makes an approach successful as a function of the target search space.

Our work aims at contributing to the development of new efficient combinatorial surrogate-assisted evolutionary multi-objective algorithms (CS-EMOA), while providing a rigorous analysis of the impact of the underlying design components through extensive benchmarking. Looking at the specialized literature, one can find a number of approaches targeting different discrete search spaces, and/or designed within a number of application-specific contexts, e.g., [43–49] to cite a few. For example, in [50], a data-driven multi-objective approach is designed specifically to constrained discrete problems. Random forests are employed as surrogates, together with logistic regression models to rectify the non-dominated ranking for the constraint functions. The derived approach is experimented on constrained multi-objective knapsack problems, and applied to trauma systems. In [51], a multi-objective task-oriented pattern mining problem is tackled using an ensemble of surrogates derived from a discrete hamming-distance based RBF network. In contrast to previous work, we focus on designing and analyzing CS-EMOAs for pseudo-boolean MOPs, viewed from a high-level perspective as a fundamental general-purpose class of optimization problems.

In a pseudo-boolean optimization problem, solutions are binary strings, and every solution is mapped to a real value for each objective. We do not make any further assumptions on the given problem, which is considered as blackbox. The blackbox problem is assumed to be expensive in the sense that the CPU time of evaluating the quality of one solution can range from at least some minutes, to hours. Pseudo-boolean optimization problems can arise in a number of real-world applications where decision variables typically model the presence/absence of design components. This is for instance the case in different engineering fields, e.g., bus stop design in public transportation systems [52, 53], bike sharing [54], drug discovery [55]. Independently of a particular real-world application, a number of relatively recent studies have investigated the design of discrete surrogate models, e.g., [39–42, 56–66]. Reviewing the theory behind existing models is out of the scope of this paper. We refer the reader to [3], and the reference therein, for a more detailed overview. In this paper, we rely on the so-called Walsh functions [67], which were shown to constitute a well-suited theoretical tool to derive custom surrogates for arbitrary pseudo-boolean functions [56, 57, 62–64]. In fact, discrete Walsh functions form an orthogonal set of functions allowing to uniquely represent

any blackbox pseudo-boolean function in an additive linear form. A first step towards integrating Walsh-based surrogates for expensive multi-objective optimization has been performed in our previous work presented in [68]. At this stage of the presentation, it is important to recall that besides a particular choice of the surrogate model, the successful design of a CS-EMOA depends on different other design choices and algorithmic components; especially, with respect to the smart configuration and training of the surrogate model, as well as, its combination with a multi-objective search process.

In this paper, we leverage our previous work from different perspectives. As in [68], we consider a decomposition-based paradigm [22] to handle the approximation set maintained by the (expensive) search process. However, we propose a number of novel design components and strategies to deal with three critical design aspects, namely, (i) the inner evolutionary optimizer used for searching promising offspring solutions with respect to the Walsh surrogate, (ii) the selection strategy allowing to decide which solution is to be evaluated by the expensive objectives, and (iii) the strategy used to setup the Walsh order hyper-parameter. Compared to the performance results presented in [68], we show that substantial improvements in terms of approximation quality can be obtained by the new proposed multi-objective search strategies. It is worth noticing that there exist no rule-of-thumb to configure the previously mentioned design components, which is a challenging research question. Therefore, besides being able to significantly improve over our previous work, we report our findings in light of a plug-and-play surrogate assisted framework which is benchmarked on the basis of a full factorial experimental design, and using a relatively broad range of benchmark functions. Hence, our results come with an in-depth comprehensive analysis on the combined effect of the different design choices, and sheds more light on the key ingredients for a successful surrogate-assisted multi-objective approach.

The rest of the paper is organized as follows. In Section 2, we recall the necessary background behind multi-objective decomposition and discrete Walsh surrogates. In Section 3, we describe the proposed approach in a step-by-step manner. In Section 4, we give our experimental setup. In Section 5, we report our empirical findings. In Section 6, we conclude the paper and we discuss future research directions.

2. Background

2.1. Multi-objective Combinatorial Optimization

We assume that we are given a black-box objective vector function $F = (f_1, f_2, \dots, f_m)$ to maximize, and a set X of solutions in the *variable space*. When X is a discrete set, we face a *multi-objective combinatorial optimization problem* (MCOP). In particular, we consider unconstrained pseudo-boolean multi-objective optimization problems, such that $F: \{0, 1\}^n \mapsto \mathbb{R}^m$, where n is the problem size. Let $Z = f(X) \subseteq \mathbb{R}^m$ be the set of feasible outcome vectors in the *objective space*. An objective vector $z \in Z$ is *dominated* by a vector $z' \in Z$ iff $\forall i \in \llbracket 1, m \rrbracket, z_i \leq z'_i$, and

$\exists j \in \llbracket 1, m \rrbracket$ s.t. $z_j < z'_j$. A solution $x \in X$ is dominated by a solution $x' \in X$ iff $F(x)$ is dominated by $F(x')$. A solution is *Pareto optimal*, or *non-dominated*, if there does not exist any other solution that dominates it. The set of all Pareto optimal solutions is the *Pareto set*. Its mapping in the objective space is the *Pareto front*. Identifying the Pareto set is known to be an NP-hard task for a wide range of MCOPs, and the Pareto set typically contains an exponential number of solutions [11]. As such, we often have to rely on a Pareto set *approximation*, for which a large number of multi-objective evolutionary algorithms have been proposed since the early nineties [12, 13].

2.2. Decomposition-based Multi-objective Optimization

In the broad range of multi-objective evolutionary algorithms, approaches based on *decomposition* are amongst the state-of-the-art [21]. In particular, MOEA/D [22] decomposes the multi-objective optimization problem into a set of single-objective sub-problems that target different regions of the Pareto front. The population of solutions $\mathcal{P}_\mu = \{x^1, \dots, x^\mu\}$ is evolved such that each solution x^i , $i \in \llbracket 1, \mu \rrbracket$, is assigned to a weight vector ω^i corresponding to a given sub-problem. A sub-problem then seeks for a high-quality solution with respect to an aggregation function $\mathbf{g}(x \mid \omega^i, F)$, parameterized by the weight vector ω^i . The population \mathcal{P}_μ is evolved following conventional evolutionary mechanisms, such as selection and variation, in order to optimize the different sub-problems.

Given a weight vector ω^i , the aggregation function transforms an objective vector into a scalar value. One recommended function that we consider in this paper is the Chebyshev function (to be minimized):

$$\mathbf{g}(x \mid \omega, F) = \max_{j \in \llbracket 1, m \rrbracket} \omega_j \cdot |z_j^* - f_j(x)|$$

where $x \in X$, $\omega = (\omega_1, \dots, \omega_m)$ is a positive weight vector, and z^* is a reference point s.t. $z_j^* > f_j(x)$, $\forall x \in X$.

In MOEA/D, the population is evolved in a cooperative manner. A solution that is currently assigned to a given sub-problem can become parent for an offspring generated by another sub-problem. A newly-generated offspring is compared to solutions assigned to other sub-problems by means of their corresponding aggregation function. At a given iteration, offspring can replace multiple solutions from the population, assuming that they improve multiple sub-problems. This cooperation may be limited by a neighborhood relation among sub-problems. Different variants [21, 69] of MOEA/D consider different ways of managing the neighborhood. Following [68], we consider a setting of MOEA/D where the whole population is used in the selection and replacements steps. This is clearly motivated by the fact that allowing an offspring to enter the population as soon as it improves any sub-problem privileges convergence, which is to be encouraged under a limited (expensive) budget setting.

2.3. Walsh Surrogates

The development of accurate surrogate models for black-box *combinatorial* problems is a difficult challenge [1]. With respect to a pseudo-boolean optimization problem, a surrogate model using the Walsh functions [67] was proved to be particularly accurate [56]. For completeness, we recall the background behind such a model.

2.3.1. Walsh Basis

Walsh functions [67] constitute an enumerable set of functions $\phi_\ell : [0, 1] \rightarrow \{-1, 1\}$ which composes a normal and orthogonal basis of the Hilbert space $L^2([0, 1])$. Like the trigonometric functions of the Fourier basis, they can be used to decompose any function from the Hilbert space [67], and have been used since the late seventies in the theory of evolutionary computation [70]. Given a pseudo-boolean function $f : \{0, 1\}^n \mapsto \mathbb{R}$, Walsh functions are defined as follows [56]. For any integer $\ell \in \llbracket 0, 2^n - 1 \rrbracket$, following the binary representation $\ell = \sum_i \ell_i 2^i$ with $\ell_i \in \{0, 1\}$, the Walsh function $\phi_\ell : \{0, 1\}^n \rightarrow \{-1, 1\}$ is defined for any binary string $x = (x_1, \dots, x_i, \dots, x_n) \in \{0, 1\}^n$ by:

$$\phi_\ell(x) = (-1)^{\sum_{i=1}^n \ell_i x_i} \quad (1)$$

The *order* of a Walsh function ϕ_ℓ , denoted by $o(\phi_\ell)$, is defined by the number of binary digits equal to 1 in the binary representation of ℓ . For example, the Walsh function of order 0 is ϕ_0 , the Walsh functions of order 1 are ϕ_{2^p} for all integers $p \geq 0$, the Walsh functions of order 2 are $\phi_{2^p + 2^{p'}}$ for all pairs of integers $p \neq p' \geq 0$, and so on.

2.3.2. Exact Walsh Transform

The so-defined (finite) set of discrete functions is a normal orthogonal basis for the space of pseudo-boolean functions, i.e., $\forall \ell, \ell' \in \llbracket 0, 2^n - 1 \rrbracket$, $\frac{1}{2^n} \sum_{x \in \{0, 1\}^n} \phi_\ell(x) \cdot \phi_{\ell'}(x) = \delta_{\ell \ell'}$. As such, *any* pseudo-boolean function f can be written as:

$$f(x) = \sum_{\ell=0}^{2^n-1} w_\ell \cdot \phi_\ell(x) \quad (2)$$

$$\text{s.t. } w_\ell = \frac{1}{2^n} \sum_{x \in \{0, 1\}^n} f(x) \cdot \phi_\ell(x) \quad (3)$$

2.3.3. Approximated Walsh Decomposition

On one hand, the Walsh functions ϕ_ℓ are uniquely defined and do not depend on the problem at hand; see Eq. (1). On the other hand, the values of the coefficients w_ℓ do depend on the considered function f , as given in Eq. (3). On top of that, there might exist an exponential number of non-zero coefficients in the exact Walsh transform as given in Eq. (2). Roughly speaking, the coefficients corresponding to some Walsh functions of a given order O capture the interaction among a set of O variables in function f ; i.e., they render how the function value is affected when

the variables change. This means that, unless a large number of variables interact with each other, many coefficients are expected to be zero. This is the case for multiple combinatorial problems with a quadratic or cubic number of variable interactions, and hence a reasonable assumption in practice for a wide range of problem classes. Therefore, the idea developed initially in [56, 57] is to approximate f using solely the Walsh functions up to a (small) constant order $\mathcal{O} \ll n$, and to use an estimate \hat{w}_ℓ of the (unknown) coefficient w_ℓ . More formally, given a *constant* order \mathcal{O} , a function f can be approximated by the following model:

$$\tilde{f}(x \mid d) = \sum_{\ell : o(\phi_\ell) \leq \mathcal{O}} \tilde{w}_\ell \cdot \phi_\ell(x) \quad (4)$$

For example, a second-order approximation can be rewritten:

$$\tilde{f}(x \mid 2) = \tilde{w}_0 + \sum_{i=1}^n \tilde{w}_i \cdot (-1)^{x_i} + \sum_{1 \leq i < j \leq n} \tilde{w}_{ij} \cdot (-1)^{x_i + x_j}$$

where \tilde{w}_0 is the zero-order estimated coefficient, and \tilde{w}_i and \tilde{w}_{ij} are first- and second-order estimated coefficients. Intuitively, the larger the order d , the more accurate the expansion to approximate the original function.

2.3.4. Walsh Surrogate Model

Constructing a discrete Walsh surrogate up to constant order \mathcal{O} consists of computing the approximate coefficients \tilde{w}_ℓ . This can be done by estimating the value of \tilde{w}_ℓ as in standard supervised machine-learning approaches. In fact, this turns out to be a standard linear regression problem, since Eq. (4) can be interpreted as a linear model whose predictors are the Walsh function values. In particular, sparse techniques can be used to minimize the number of non-zero coefficients when the number of predictors is large [71]. Following [56, 57, 62], we use the Lasso algorithm [72] to fit the approximate model, which is of particular interest given that the number of Walsh functions of up to a given constant order \mathcal{O} might be greater than the number of solutions used for training.

3. A Surrogate-assisted Evolutionary Multi-objective Framework based on Walsh Functions

In this section, we describe a surrogate-assisted multi-objective combinatorial optimization framework using the Walsh functions as surrogate (SMCO/w). The proposed framework integrates three main design components; (i) the inner optimizer of the substitute Walsh surrogate, (ii) the selection of the solution to be evaluated at each iteration, and (iii) the setting of the Walsh order at the training phase. These design components can be set interchangeably following different strategies which are described in a step-by-step detailed manner in the following.

3.1. General Description

The high-level pseudo-code of the proposed SMCO/w framework is depicted in the template of Algorithm 1. It follows the general computational flow of the (surrogate-less) MOEA/D algorithm, as described in Section 2.2, with a few exceptions. The original problem is decomposed into μ single-objective sub-problems using the Chebychev scalarizing function. A population \mathcal{P}_μ is initialized with μ randomly-generated solutions. Each solution is evaluated using the expensive objectives before being assigned to a unique sub-problem. This population is also used as the initial dataset \mathcal{D} for training the surrogates. A generation of the algorithm (line 6) consists of iterating once over all the sub-problems. The algorithm stops when a given termination condition is satisfied, which renders the computing budget one can afford. In an expensive setting, where the evaluation function is assumed to be very time consuming, the budget is expressed as a maximum number of calls to the objective vector function. In our work, we consider a range of values to better render the anytime behavior of the considered algorithms. Furthermore, as will be detailed later in our experimental study, we shall discuss the implication of using the Walsh surrogate on the computational time.

At the beginning of every iteration i dealing with a sub-problem $i \in \llbracket 1, \mu \rrbracket$, a Walsh order \mathcal{O} is chosen following the order selection component in line 7. The order selection component takes as input a *History* variable, which is an artifact indicating that some information about the search status may be used at this step. We then consider one surrogate for each of the m objectives. These surrogates $\tilde{\mathcal{F}} = (\tilde{f}_1, \dots, \tilde{f}_m)$ are trained (line 8) using the (training) dataset \mathcal{D} , which contains *all* solutions evaluated so far. The estimates of the Walsh coefficients (Eq. 4) for each surrogate \tilde{f}_i are computed following a sparse linear regression methodology, as discussed in Section 2.3.

In contrast to surrogate-less evolutionary algorithms, SMCO/w relies temporarily on the so-trained surrogates to intensively search for high-quality solutions. In other words, after the training step, a pool of offspring solutions \mathcal{S} is generated by a *surrogate optimizer* (line 10), without calling the true objective functions. This component takes the current population as input, together with the current surrogates $\tilde{\mathcal{F}}$ and a copy z^{**} of the reference point. It is important to use a copy of z^* in the surrogate optimizer because the reference point will likely be updated on the basis of the (unreliable) estimated objective values, whereas z^* is only updated according to the true objective values evaluated so far. The different design choices for this inner discrete surrogate optimization component are discussed in detail in Section 3.2.

The next step (line 11) selects one candidate solution from \mathcal{S} in order to be evaluated using the expensive objectives F . We call this component the *selection strategy*, and we describe it in Section 3.3. It takes the pool of solutions \mathcal{S} generated by the optimizer as input, together with the surrogate model $\tilde{\mathcal{F}}$, the reference point z^{**} as updated by the optimizer, and the current weight vector ω_i . As output, it returns the solution to be evaluated at the current iteration i . The weight vector ω_i given as input of the selection component is only an artifact indicating that this component may depend on iteration i .

Algorithm 1: Surrogate-assisted Multi-objective Combinatorial Optimization based on Walsh basis (SMCO/w)

Input: $\mathcal{W}_\mu := \{\omega^1, \dots, \omega^\mu\}$: weight vectors; $g(\cdot | \omega, \cdot)$: a scalarizing function to be minimized; O_{\max} : maximum order of Walsh functions;

- 1 $\mathcal{P}_\mu \leftarrow \{x^1, \dots, x^\mu\}$: initial population of size μ ;
- 2 $\mathcal{D} \leftarrow \{(x^1, F(x^1)), \dots, (x^\mu, F(x^\mu))\}$: training set;
- 3 $\mathcal{EP} \leftarrow$ initialize external archive (optional) ;
- 4 $z^* \leftarrow$ initialize reference point;
- 5 **while** *the termination condition is not satisfied* **do**
- 6 **for** $i \in \{1, \dots, \mu\}$ **do**
- 7 // Choose the Walsh order
 $O \leftarrow \text{WALSHORDER}(\text{History}, O_{\max})$;
- 8 // Train Walsh models
 $\tilde{\mathcal{F}} := (\tilde{f}_1, \dots, \tilde{f}_m) \leftarrow \text{TRAINWALSH}(\mathcal{D}, O)$;
- 9 // Copy the reference point
 $z^{**} \leftarrow z^*$;
- 10 // Optimize the Walsh surrogate
 $\mathcal{S} \leftarrow \text{RUNOPTIMIZER}(\mathcal{P}_\mu, \tilde{\mathcal{F}}, z^{**})$;
- 11 // Selection for true evaluation
 $x' \leftarrow \text{SELECTEVAL}(\mathcal{S}, \omega_i, \tilde{\mathcal{F}}, z^{**})$;
- 12 $F(x') \leftarrow$ evaluate x' ;
- 13 $\mathcal{EP} \leftarrow$ (optional) update external archive with x' ;
- 14 $z^* \leftarrow$ update reference point using $F(x')$;
- 15 // Replacement
for $j \in \{1, \dots, \mu\}$ **do**
- 16 **if** $g(x' | \omega^j, F) < g(x^j | \omega^j, F)$ **then**
- 17 $x^j \leftarrow x'$;
- 18 // Update training data
 $\mathcal{D} \leftarrow \mathcal{D} \cup \{(x', F(x'))\}$;

After the selection and the evaluation of the offspring solution x' , we follow the standard process of MOEA/D, with the update of the external archive \mathcal{EP} , the update of the reference point z^* , and the replacement of the population. At this step, we compare the scalarized fitness value of the candidate solution x' with the solutions from the current population \mathcal{P}_μ , according to their corresponding weight vectors. The candidate solution x' may enter the population if any improvement is found as discussed previously in Section 2.2. Finally, the training dataset \mathcal{D} is updated with the newly evaluated offspring x' (line 18).

To summarize, SMCO/w is based on three generic components that can be configured in different manners. In the next paragraphs, we propose different possible strategies.

Algorithm 2: Multiple Local Search (MLS)

Input: $\tilde{\mathcal{F}} := (\tilde{f}_1, \dots, \tilde{f}_m)$: surrogates; $\mathcal{W}_\mu := \{w^1, \dots, w^\mu\}$: weight vectors; $g(\cdot | \omega, \cdot)$: Chebychev scalarizing function; $\mathcal{N} : X \leftarrow 2^X$: a neighborhood relation;

```
1  $\mathcal{S} \leftarrow \emptyset$  ;
2 for  $\omega^i \in \mathcal{W}_\mu$  do
3   LocalOptimum  $\leftarrow$  False;
4    $x^* \leftarrow$  random (initial) solution;
5   while ! LocalOptimum do
6      $x'^* \leftarrow x^*$ ;
7     foreach  $x' \in \mathcal{N}(x^*)$  do
8       if  $g(x' | \omega^i, \tilde{\mathcal{F}}) < g(x'^* | \omega^i, \tilde{\mathcal{F}})$  then  $x'^* \leftarrow x'$  ;
9     if  $g(x'^* | \omega^i, \tilde{\mathcal{F}}) < g(x^* | \omega^i, \tilde{\mathcal{F}})$  then
10       $x^* \leftarrow x'^*$ ;
11     else
12      LocalOptimum  $\leftarrow$  True;
13    $\mathcal{S} \leftarrow \mathcal{S} \cup \{x^*\}$  ;
14 return  $\mathcal{S}$ 
```

3.2. Component #1: Surrogate Optimizer

The choice of an accurate surrogate optimizer is important since it allows us to search for high-quality solutions according to the model $\tilde{\mathcal{F}}$ without ever calling the true evaluation functions. We investigate three alternative optimizers using seemingly different search paradigms. The considered optimizers are summarized in the following:

Optim_{MOEA/D} performs the conventional MOEA/D in order to identify a Pareto set approximation with respect to the surrogates $\tilde{\mathcal{F}}$. It is initialized with the current population from the main SMCO/W algorithm, and returns the evolved population after a number of generations. Using MOEA/D as an optimizer was recommended in [68]; however, we shall show that it can be substantially out-performed by the other proposed techniques.

Optim_{MLS} (Algorithm 2) independently runs multiple local search, one for each defined sub-problem. Given the surrogates $\tilde{\mathcal{F}}$, a standard single-objective hill-climber is performed for each sub-problem defined by the weight vector ω^i , $i \in \llbracket 1, \mu \rrbracket$. Each hill-climber is initialized with a random solution, and iteratively selects the best improving solution, if any, using a standard 1-bit-flip neighborhood, until the search stops in a local optimum with respect to the aggregation function $g(\cdot | \omega^i, \tilde{\mathcal{F}})$. **Optim_{MLS}** then returns a pool of μ (local optimal) solutions, one per sub-problem.

Optim_{PLS} (Algorithm 3) is based on the so-called Pareto Local Search [73] algorithm. It maintains

Algorithm 3: Pareto Local Search (PLS)

Input: $\tilde{\mathcal{F}} := (\tilde{f}_1, \dots, \tilde{f}_m)$: surrogates; $\mathcal{N} : X \leftarrow 2^X$: a neighborhood relation;

```
1  $x \leftarrow$  random (initial) solution;
2  $\mathcal{S} \leftarrow \{x\}$ ; // Non dominated archive
3  $R \leftarrow \{x\}$ ; // Non visited solutions
4 while  $R \neq \emptyset$  do
5    $x' \leftarrow$  select a solution at random in  $R$ ;
6   foreach  $x'' \in \mathcal{N}(x')$  do
7     if  $x''$  is not dominated by any solution in  $\mathcal{S}$  then
8       for  $x \in \mathcal{S}$  do
9         if  $x$  is dominated by  $x''$  then  $\mathcal{S} \leftarrow \mathcal{S} \setminus \{x\}$ ;
10         $\mathcal{S} \leftarrow \mathcal{S} \cup \{x''\}$ ;
11        for  $x \in R$  do
12          if  $x$  is dominated by  $x''$  then  $R \leftarrow R \setminus \{x\}$ ;
13           $R \leftarrow R \cup \{x''\}$ ;
14    $R \leftarrow R \setminus \{x'\}$ ;
15 return  $\mathcal{S}$ 
```

an unbounded archive of mutually non-dominated solutions, initialized with a random solution. At each step, one solution is selected at random from the archive, all its neighbors are evaluated with respect to $\tilde{\mathcal{F}}$, and the archive is updated accordingly. The current solution is tagged as visited in order to avoid a useless re-exploration of its neighborhood. The search process stops once all solutions in the archive are visited. The content of the archive corresponds to the pool of solutions returned by $\text{Optim}_{\text{PLS}}$.

3.3. Component #2: Selection Strategy

The selection of the solution among the whole pool returned by the surrogate optimizer is a critically important component (line [11](#) in Algorithm [1](#)). Let us consider some iteration $i \in \llbracket 1, \mu \rrbracket$, which can be thought as corresponding to a particular sub-problem ω^i . An intuitive strategy is to attempt to improve sub-problems in a round-robin manner, and thus to select the solution whose estimated scalar value (computed on the basis of the surrogate) is the best for the current sub-problem. This was actually implemented in [\[68\]](#). However, we shall show that this local selection strategy can be sub-optimal, and is out-performed by other carefully designed global strategies. More specifically, for an iteration $i \in \llbracket i, \mu \rrbracket$, we study the following four selection strategies:

Select_{local} chooses the solution x' with the best aggregation value for the sub-problem considered at the current iteration and correspondingly to the weight vector ω^i .

$$x' := \operatorname{argmin}_{x \in \mathcal{S}} \mathbf{g}(x \mid \omega^i, \tilde{\mathcal{F}})$$

Select_{global} chooses the solution x' with the best aggregation value with respect to any sub-problem $\ell \in \llbracket 1, \mu \rrbracket$, independently of the current iteration:

$$x' := \operatorname{argmin}_{x \in \mathcal{S}} \min_{1 \leq \ell \leq \mu} \mathbf{g}(x \mid \omega^\ell, \tilde{\mathcal{F}})$$

Select_{BI} (best improvement) chooses the solution x' that improves the most the aggregation value of any solution in the population, independently of the current iteration:

$$x' := \operatorname{argmin}_{x \in \mathcal{S}} \min_{1 \leq \ell \leq \mu} \mathbf{g}(x_\ell \mid \omega^\ell, \tilde{\mathcal{F}}) - \mathbf{g}(x \mid \omega^\ell, \tilde{\mathcal{F}})$$

Select_{NBI} (best normalized improvement) is based on the previous strategy, but the improvement value is normalized by the actual aggregation value of the current solution associated with each sub-problem:

$$x' := \operatorname{argmin}_{x \in \mathcal{S}} \min_{1 \leq \ell \leq \mu} \frac{\mathbf{g}(x_\ell \mid \omega^\ell, \tilde{\mathcal{F}}) - \mathbf{g}(x \mid \omega^\ell, \tilde{\mathcal{F}})}{\mathbf{g}(x_\ell \mid \omega^\ell, \tilde{\mathcal{F}})}$$

3.4. Component #3: Walsh Order

As discussed in Section 2.3, a Walsh surrogate requires an order up to which the coefficients are expanded. The choice of the Walsh order is an important and difficult task for black-box problems, which was not deeply addressed in the past. This choice infers the number of coefficients in the surrogate model (Eq. 4). This does not only impact the model theoretical accuracy. It also impacts the fitting process itself, because the model complexity increases with the number of coefficients. Let us illustrate this for a problem of size $n = 50$. In such a case, there are 51 coefficients for an order $\mathcal{O} = 1$, 1 276 coefficients for an order $\mathcal{O} = 2$, and 20 876 coefficients for an order $\mathcal{O} = 3$. A low number of coefficients may be insufficient to accurately approximate the objectives. A high number of coefficients eventually improves the approximation quality, but a larger dataset should then be used in order to accurately train the model. Hence, this may seem contradictory with the fact that only a restricted number of evaluated solutions can be used for training.

In this paper, we assume that the order used to train the model is bounded by a constant value \mathcal{O}_{\max} . This is a reasonable assumption in practice, since otherwise the number of Walsh coefficients would grow exponentially. We also remark that such an assumption is in line with the fact that many challenging combinatorial optimization problems can be decomposed in an exact manner using Walsh functions of bounded order. However, it remains unclear which concrete order to use within the range $\llbracket 1, \mathcal{O}_{\max} \rrbracket$. Consequently, we investigate three simple strategies for setting the Walsh order, as described below.

Order_{static} corresponds to a baseline strategy for setting the Walsh order. The order is simply a static user-defined parameter, i.e., the Walsh surrogate is fitted considering a fixed order value

$O \in \llbracket 1, O_{\max} \rrbracket$, where O is assumed to be a parameter provided by the end-user.

Order_{random} is a basic dynamic strategy where the order is set randomly every time the model is trained. In other words, at each iteration, the value of the Walsh order is picked uniformly at random within the range $\llbracket 1, O_{\max} \rrbracket$. The advantage of this strategy is to use all admissible orders without being completely limited by an improper static setting. This simple strategy allows us to illustrate the benefits of changing dynamically the order value, in comparison to a static setting.

Order_{greedy} dynamically adjusts the Walsh order during the search process, depending on the data collected so far. At the end of each iteration $i \in \llbracket 1, \mu \rrbracket$, we compute the number of sub-problems that are improved (in line [16](#)). Let us denote by p_j the number of improved sub-problems, where j corresponds to the j^{th} true function evaluation. We thereby track the number of improved sub-problems over a window of size t (last) iterations, where t is a user-defined parameter. At the beginning of each iteration, we compute the average, over the window, of the number of improved sub-problems, denoted \bar{p}_t . If $\bar{p}_t < 1$, then we interpret this as a signal for changing the Walsh order. This is performed according to a particular schedule, as discussed in the following. Let O_{curr} be the order used at a current iteration $i \in \llbracket 1, \mu \rrbracket$. Let O_{prev} be the order used at the previous iteration (i.e., iteration $i - 1$ or μ). At the beginning of the next iteration, let us assume that $\bar{p}_t < 1$. Then, the new order O_{new} returned by the selection strategy is given by $O_{\text{new}} = O_{\text{curr}} + 1$ if either $O_{\text{curr}} = 1$ or $O_{\text{prev}} < O_{\text{curr}} < O_{\max}$; otherwise, we set $O_{\text{new}} = O_{\text{curr}} - 1$.

The idea behind this greedy dynamic strategy is to start the search with the smallest order 1. If we observe that it allows the current population to improve ($\bar{p}_t \geq 1$), then we keep using the same successful order. Otherwise, we increase the order, and hence the model complexity, in the hope of obtaining a more accurate surrogate. However, when attaining the maximum allowed model complexity O_{\max} , we continue scanning the possible order range by decreasing the order value, and increasing it again if an order of 1 is reached. The current order value stops changing whenever the model is accurate enough to improve the population, and the whole dynamic schedule is repeated as soon as no more improvements are observed.

4. Experimental Setup

In the remainder of the paper, we report a comprehensive analysis of the combined effect of the previously described components. Obviously, extensive experiments with real-world expensive black-box MCOPs would be computationally infeasible. Besides, our primary goal is not to tackle an application-specific problem, but to conduct a general-purpose informative study providing generic guidelines for designing a discrete surrogate-assisted approach. For this purpose, we consider two classes of well-established MCOP benchmarks. This is described below together with the experimented algorithm variants.

4.1. Benchmark Problems

We consider bi-objective NK-landscapes [74, 75] and bi-objective unconstrained binary quadratic programming problems (UBQP) [76], as a set of challenging and representative pseudo-boolean MCOPs. In fact, NK-landscapes allows one to control the degree of non-linearity of the problem. We hence consider instances with a variable degree of difficulty. UBQP is a computationally challenging problem which is known to embrace a remarkable range of applications in combinatorial optimization [77].

4.1.1. Multi-objective NK-Landscapes (MNK-Landscapes)

Solutions are binary strings of size n and the objective vector, to be maximized, is $F = (f_1, \dots, f_m): \{0, 1\}^n \mapsto [0, 1]^m$. Each objective function is defined by [78]:

$$f_i(x) = \frac{1}{n} \sum_{j=1}^n c_j^i(x_j, x_{j_1}, \dots, x_{j_k})$$

where the $c_j^i: \{0, 1\}^{k+1} \rightarrow [0, 1)$ are component functions, and k is a parameter specifying the number of epistatic interactions. The component function c_j^i assigns a real-valued contribution for every combination of x_j and its k epistatic interactions $\{x_{j_1}, \dots, x_{j_k}\}$. The parameter k defines the degree of non-linearity of the problem, and hence the *ruggedness* of the landscape.

We consider bi-objective MNK-landscapes [75] with the following setting: $n \in \{25, 50\}$ and $k \in \{0, 1, 2\}$. In line with previous studies [57, 68], this allows us to span instances ranging from a small to a relatively large number of variables, and following a linear ($k = 0$), quadratic ($k = 1$) and cubic ($k = 2$) variable interaction scheme.

4.1.2. Multi-objective UBQP (MUBQP)

Given a collection of n items such that each pair of items is associated with multidimensional profit values, the multi-objective unconstrained binary quadratic programming (MUBQP) problem seeks a subset of items that maximizes the sum of their paired values in each dimension [76]. The value of a pair is summed up only if the two corresponding items are selected. A solution can be represented as a binary string of size n . Each position from the binary string maps to a particular variable that indicates whether the corresponding item is included in the subset of selected items or not. Given a solution $x = (x_1, \dots, x_n)$, each objective f_k , $k \in \llbracket 1, m \rrbracket$, is defined as follows:

$$f_k(x) = \sum_{i=1}^n \sum_{j=1}^n q_{ij}^k x_i x_j$$

where $Q_k = q_k^{ij}$ is an $n \times n$ matrix of constant values, either positive, negative or zero, $k \in \llbracket 1, m \rrbracket$. We consider bi-objective UBQP instances with $n \in \{25, 50\}$ variables and a density of 0.9 for non-zero entries in the matrix Q_k to infer difficult and highly non-linear search landscapes.

Table 1: Summary of the considered SMCO/w components.

Component	Proposed configuration values	analyzed in Sec.
Optimizer	Optim _{MOEA/D} [68], Optim _{MLS} , Optim _{PLS}	5.1 & 5.2
Selection	Select _{local} [68], Select _{global} , Select _{BI} , Select _{NBI}	5.1
Order	Order _{static} [68] ($O \in \{1, 2, O_{\max} = 3\}$), Order _{random} , Order _{greedy}	5.3

4.2. Experimental set-up

In order to fully analyze the impact of the proposed strategies, we systematically evaluate the performance of *all* possible combinations. Furthermore, we consider to evaluate the gain in approximation quality in comparison to an evolutionary solving process without any surrogate. For this purpose, we consider the three multi-objective algorithms used to optimize the learned problem model as baselines. In other words, we shall analyze the search performance obtained with MOEA/D, MLS and PLS as a main optimization procedure for the original expensive problem.

For all experimented algorithms, we use 50 weight vectors such that $\omega^i = \left(\frac{i-1}{\mu-1}, \frac{1-(i-1)}{\mu-1}\right)$, $i \in \llbracket 1, 50 \rrbracket$. For Walsh surrogates, a Lasso regression is used to fit the model [72]. The maximum value allowed for the Walsh order is set to $O_{\max} = 3$ which corresponds to a cubic model. When using the Order_{static} strategy, the actual order d used for training is a user-defined parameter, which we set in the range $O \in \{1, 2, 3 = O_{\max}\}$, i.e., three O -values are experimented for each variant using Order_{static}. The greedy strategy to adapt the Walsh order Order_{greedy} uses a window of size $t = 5$. Besides, each surrogate optimizer comes with different parameters. Optim_{MOEA/D} is executed for 10 generations, as in [68]. It uses a one-point crossover followed by a uniform bit-flip mutation with a rate of $1/n$. The solution neighborhood considered in Optim_{MLS} and Optim_{PLS} is based on the bit-flip operator; i.e., two solutions are neighbors if their Hamming distance is 1. Finally, Optim_{MLS} uses the same 50 weight vectors as defined previously.

Overall, we experiment 60 possible configurations (See Table 1 for a summary): 3 surrogate optimizers \times 4 selection strategies \times (3 + 2) Walsh order settings, together with 3 additional surrogate-less algorithms. Each algorithm is independently executed 10 times on the 8 considered MCOP instances (2×3 MNK-landscapes + 2 MUBQP), for a maximum budget of 1 500 (expensive) evaluations and not exceeding 64 CPU hours per run. The experiments were conducted in Python 3.7 on an Intel Core Xeon E5-2630 (2.20 GHz, 256 GB RAM) under Debian 10.

For performance assessment, we use the additive epsilon indicator [79]. Notice, however, that our results were found to be consistent when using the hypervolume indicator. This is not reported explicitly to avoid flooding the reader with too much data. The epsilon indicator gives the minimum factor by which an approximation set has to be translated in the objective space in order to (weakly) dominate a reference set. The lower the indicator, the better the approximation. The reference set is constructed by merging the solutions found over all runs and configurations for a given instance, and removing dominated ones. For a given run, the archive of all non-dominated solutions found during the search process is used to measure approximation quality. Finally, we shall report the

approximation quality for different intermediate budgets, so as to render the relative anytime performance.

5. Experimental Analysis

In this section, we report our analysis and empirical findings. Since we consider an extensive number of configurations, a main challenge is to fairly address the impact of the combined effect of the designed components on search performance. In fact, as our analysis will reveal, we found that there is a non-trivial interaction between the different components, which we report in a step-by-step manner. In Section 5.1, we start by studying the interaction between the surrogate optimizer (Component #1) and the selection strategy (Component #2). In Section 5.2, we analyze the relative performance obtained when the surrogate optimizer is combined with the corresponding best selection strategy. In Section 5.3, we study the impact of the Walsh order setting (Component #3). In Section 5.4, we report the performance gain against surrogate-less approaches.

5.1. Impact of Selection on Surrogate Optimizers

In the following, we address the relative impact of the selection strategy (Component #2) on each *individual* SMCO/W variant obtained with a particular surrogate optimizer (Component #1). For MNK-landscapes, this is summarized in Tables 2, 3, and 4, respectively when using $\text{Optim}_{\text{MOEA/D}}$, $\text{Optim}_{\text{MLS}}$, and $\text{Optim}_{\text{PLS}}$. For MUBQP, this is summarized in Table 5 for all optimizers. The tables show the relative rank of each selection strategy and each *static* setting of the Walsh order. A rank c indicates that the corresponding algorithm is significantly outperformed by c other strategies using a Wilcoxon test with a Bonferroni correction at a significance level of 0.05. Ranks in bold correspond to approaches that are not significantly outperformed by any other. Different budgets, ranging from a small to a medium and a relatively high number of true evaluations, are considered. This is to better render the *anytime* behavior. We also show the average epsilon indicator value obtained for each variant (in parentheses in the tables), which is to provide the reader with a preliminary idea of the overall relative performance (underlined values in the tables show the best averages). Let us comment that the ranks shown in Tables 2, 3, 4, and 5, still do not allow to elicit in a comprehensive manner what is the best possible configuration. A more specific analysis is to follow in the subsequent sections, after analyzing the relative impact of the selection strategy when combined with a particular optimizer. In fact, from these tables, we can for now draw two main general observations.

Firstly, the relative rank obtained by a given selection strategy is overall consistent over the different considered instances, although for MNK-landscapes, the relative performance is a function of the value of k , which is the degree of non-linearity/ruggedness of the landscape. This is also related to the Walsh order setting. We can observe that the Walsh order O providing the best rank depends consistently on the value of k , such that $O = k + 1$. For MUBQP instances, and since

Table 2: Rank and average epsilon indicator value ($\times 10^2$, between brackets) after 200, 700, and 1500 evaluations for MNK-landscapes. The results are w.r.t. $\text{Optim}_{\text{MOEA/D}}$ and $\text{Order}_{\text{static}}$.

			#eval	Select _{local}			Select _{global}			Select _{BI}			Select _{NBI}		
				O=1	O=2	O=3	O=1	O=2	O=3	O=1	O=2	O=3	O=1	O=2	O=3
Optim _{MOEA/D}	k = 0	n = 25	200	0 _(1.4)	0 _(1.8)	3 _(2.4)	9 ₍₅₎	9 _(5.4)	9 _(5.7)	0 ₍₁₎	0 _(1.2)	0 _(1.8)	0 _(1.3)	0 _(1.4)	3 _(2.4)
			700	0 ₍₁₎	0 _(1.1)	2 _(1.3)	9 _(4.4)	9 _(4.1)	9 _(4.1)	0 _(0.7)	0 _(0.7)	0 _(0.9)	0 _(1.1)	0 ₍₁₎	0 _(1.1)
			1500	0 _(0.8)	0 _(0.8)	0 _(0.9)	9 _(3.9)	9 _(3.6)	9 _(3.2)	0 _(0.6)	0 _(0.6)	0 _(0.7)	0 _(0.8)	0 _(0.9)	0 _(0.9)
		avg. rank		0.56			7.17			0.67			3.5		
	k = 1	n = 25	200	3 _(4.2)	0 _(1.5)	0 _(2.6)	8 _(8.7)	6 _(6.8)	8 _(7.7)	3 _(4.5)	0 _(1.1)	0 _(2.5)	3 _(4.2)	8 _(7.2)	0 _(1.7)
			700	5 _(3.4)	0 _(0.5)	0 _(0.5)	7 _(6.6)	6 _(5.7)	7 _(6.2)	5 _(3.8)	0 _(0.5)	0 _(0.6)	5 _(2.8)	7 ₍₆₎	0 _(0.6)
			1500	5 _(2.9)	0 _(0.3)	0 _(0.4)	7 _(5.7)	5 _(4.4)	6 _(4.8)	5 _(3.5)	0 _(0.4)	0 _(0.4)	5 _(2.6)	7 _(5.7)	0 _(0.4)
		avg. rank		0.89			5.67			0.78			5		
	k = 2	n = 25	200	4 ₍₈₎	0 _(5.6)	0 _(5.9)	6 ₍₉₎	6 _(8.3)	4 _(8.4)	2 _(7.5)	0 _(4.3)	0 _(4.7)	5 _(8.2)	0 _(4.7)	0 _(5.2)
			700	6 _(6.5)	0 _(3.9)	0 ₍₃₎	6 _(7.7)	1 _(5.4)	0 _(6.2)	6 _(7.2)	0 _(2.9)	0 _(2.7)	5 _(6.8)	0 _(3.6)	0 _(3.2)
			1500	6 _(6.2)	0 _(3.4)	0 ₍₂₎	6 _(7.2)	0 ₍₅₎	0 _(5.3)	5 _(5.9)	0 _(2.4)	0 _(2.1)	5 _(6.1)	0 _(2.8)	0 _(2.7)
		avg. rank		0.94			4.17			0.72			3.67		
n = 50	n = 25	200	0 _(5.2)	0 _(4.8)	1 _(6.3)	6 _(9.7)	6 _(10.6)	6 _(10.9)	0 _(5.8)	0 _(4.8)	0 _(5.6)	6 _(9.8)	6 _(10.7)	6 _(10.9)	
		700	1 _(4.1)	0 _(3.2)	0 _(3.7)	6 _(7.7)	5 _(7.7)	6 _(9.1)	1 _(4.4)	0 _(2.8)	0 _(3.4)	6 _(8.7)	6 _(9.6)	7 _(9.9)	
		1500	1 _(3.8)	0 _(2.9)	0 _(3.1)	4 _(6.5)	0 _(5.6)	6 ₍₆₎	0 _(3.9)	0 _(2.8)	0 _(2.8)	6 _(8.3)	6 _(8.6)	6 _(8.7)	
	avg. rank		0.89			5.67			0.78			5			

Table 3: Rank and average epsilon indicator value ($\times 10^2$, between brackets) after 200, 700, and 1500 evaluations for MNK-landscapes. The results are w.r.t. $\text{Optim}_{\text{MLS}}$ and $\text{Order}_{\text{static}}$.

			#eval	Select _{local}			Select _{global}			Select _{BI}			Select _{NBI}		
				O=1	O=2	O=3	O=1	O=2	O=3	O=1	O=2	O=3	O=1	O=2	O=3
Optim _{MLS}	k = 0	n = 25	200	5 _(4.7)	3 _(3.3)	2 _(2.4)	3 _(2.2)	4 _(4.7)	7 _(5.2)	5 _(3.7)	3 _(3.2)	3 _(2.8)	0 _(1.1)	0 _(0.4)	1 ₍₁₎
			700	4 _(3.1)	3 _(2.4)	3 _(1.5)	3 _(2.1)	8 _(4.1)	8 _(4.7)	3 _(2.9)	3 _(2.3)	2 _(1.8)	0 _(0.3)	0 _(0.3)	0 _(0.3)
			1500	3 _(2.3)	3 _(1.7)	0 _(1.1)	3 _(2.1)	10 _(3.9)	10 ₍₄₎	3 _(2.2)	3 _(1.7)	2 _(1.3)	0 _(0.3)	0 _(0.3)	0 _(0.3)
		avg. rank		3			6.89			2.44			0.17		
	k = 1	n = 25	200	2 _(3.9)	1 _(2.7)	1 _(2.8)	6 _(6.3)	4 _(5.4)	6 _(6.3)	3 _(4.8)	1 _(1.5)	1 _(3.2)	4 _(4.7)	0 _(0.6)	1 _(1.7)
			700	6 _(2.9)	3 _(1.2)	1 _(0.6)	9 _(5.5)	7 _(4.8)	8 _(5.2)	6 _(3.9)	2 ₍₁₎	1 _(0.9)	6 _(3.8)	0 _(0.4)	0 _(0.3)
			1500	6 _(2.6)	3 _(0.9)	1 _(0.5)	9 _(5.1)	7 _(4.4)	9 _(4.9)	6 _(3.7)	2 _(0.8)	1 _(0.8)	6 _(3.5)	0 _(0.4)	0 _(0.3)
		avg. rank		2.5			7.33			2.22			1.61		
	k = 2	n = 25	200	0 _(7.4)	0 _(6.4)	0 _(5.9)	4 _(8.7)	0 ₍₅₎	1 _(7.1)	1 _(8.5)	0 _(5.7)	0 _(6.1)	1 _(8.3)	0 _(5.3)	0 _(5.5)
			700	4 _(5.6)	3 _(2.8)	0 _(1.1)	8 _(7.4)	3 _(3.6)	4 _(4.8)	7 _(7.1)	2 _(2.9)	0 _(0.9)	7 _(7.3)	2 _(2.5)	0 _(0.5)
			1500	6 _(5.1)	3 ₍₂₎	0 _(0.7)	8 _(7.1)	6 _(3.5)	6 _(4.6)	7 _(6.6)	3 _(2.1)	0 _(0.5)	8 _(6.9)	3 ₍₂₎	0 _(0.5)
		avg. rank		1.33			5.17			1.83			1.94		
n = 50	n = 25	200	0 ₍₇₎	0 _(7.7)	0 _(8.4)	8 _(10.7)	0 _(8.8)	1 _(9.5)	0 _(8.3)	0 _(7.3)	0 _(7.6)	1 _(9.1)	0 ₍₇₎	0 _(6.3)	
		700	2 _(4.7)	0 _(3.6)	0 _(3.7)	9 ₍₁₀₎	5 _(7.2)	7 _(8.3)	6 _(6.3)	0 _(3.2)	0 ₍₄₎	6 _(6.6)	0 _(2.8)	0 _(3.3)	
		1500	6 _(4.3)	0 _(1.9)	0 _(1.9)	11 _(9.8)	6 _(6.7)	6 _(7.3)	7 _(6.1)	0 _(1.8)	0 ₍₂₎	7 ₍₆₎	0 _(1.4)	0 _(1.5)	
	avg. rank		1.33			5.17			1.83			1.94			

Table 4: Rank and average epsilon indicator value ($\times 10^2$, between brackets) after 200, 700, and 1500 evaluations for MNK-landscapes. The results are w.r.t. **OptimPLS** and **Order_{static}**.

			#eval	Select _{local}			Select _{global}			Select _{BI}			Select _{NBI}		
				O=1	O=2	O=3	O=1	O=2	O=3	O=1	O=2	O=3	O=1	O=2	O=3
OptimPLS	k = 0	n = 25	200	0 (0.4)	0 (0.4)	6(1.7)	10(6.4)	9(5.4)	9(4.9)	0 (0.5)	3(0.8)	3(1.2)	0 (0.4)	0 (0.5)	2(1)
			700	0 (0.4)	0 (0.4)	0 (0.3)	9(5.4)	9(4.6)	9(4.4)	1(0.5)	4(0.8)	4(1.1)	0 (0.4)	0 (0.5)	0 (0.5)
			1500	0 (0.4)	0 (0.4)	0 (0.3)	10(4.7)	9(4.3)	9(3.9)	1(0.5)	4(0.8)	4(1.1)	0 (0.4)	0 (0.5)	0 (0.5)
		avg. rank		0.94		9.06		2.56		0.39					
	k = 1	n = 25	200	5(3.7)	0 (0.8)	1(1.5)	8(6.8)	8(6.5)	7(6.1)	6(5)	0 (0.9)	1(2.1)	5(4.4)	0 (0.5)	1(1.3)
			700	6(2.8)	0 (0.3)	0 (0.3)	8(5.6)	8(5.6)	7(5.5)	7(4.6)	0 (0.6)	2(0.6)	6(4.1)	0 (0.4)	0 (0.4)
			1500	6(2.8)	0 (0.3)	0 (0.3)	8(5.2)	7(4.8)	7(5)	7(4.3)	0 (0.6)	3(0.6)	6(4)	0 (0.4)	0 (0.4)
		avg. rank		2		8		2.44		1.83					
	k = 2	n = 25	200	3(6.4)	0 (4.8)	0 (5)	9(8.9)	0 (5.8)	0 (6.4)	3(7.7)	0 (3.9)	0 (5.3)	1(6.8)	0 (3.8)	0 (4)
			700	7(6)	3(2.3)	0 (1)	8(8)	4(3.6)	4(4.7)	5(5.9)	3(2.5)	0 (0.7)	6(5.9)	3(2.1)	0 (0.7)
			1500	6(5.7)	3(1.8)	0 (0.2)	8(7.1)	6(3.6)	6(4.5)	6(5.4)	3(2.3)	0 (0.6)	7(5.7)	3(1.9)	0 (0.5)
		avg. rank		1.89		5.61		1.94		2					
n = 50	n = 25	200	1(4.6)	0 (3.5)	2(5.4)	7(7.1)	8(7.5)	7(7.7)	1(4.9)	1(4)	1(5.2)	1(4.6)	0 (2.4)	1(4.3)	
		700	6(3.5)	0 (0.6)	3(1.7)	9(6.7)	9(7.3)	9(6.7)	6(4.1)	0 (0.6)	3(1.9)	6(4)	0 (0.4)	1(1)	
		1500	6(3.4)	0 (0.4)	0 (0.4)	9(6.6)	9(7.1)	9(6.4)	6(4)	0 (0.6)	0 (0.5)	6(4)	0 (0.4)	0 (0.4)	
	avg. rank		2		8		2.44		1.83						
n = 50	n = 25	200	2(8)	0 (5.6)	0 (6.6)	7(10.7)	0 (8.1)	5(9.5)	5(9)	0 (6.4)	0 (6.7)	4(8.5)	0 (6.2)	0 (5.9)	
		700	4(5.9)	0 (3)	0 (3.8)	10(9.9)	4(6.6)	6(7.8)	4(6.2)	0 (2.7)	0 (4)	6(6.3)	0 (2.4)	0 (3.3)	
		1500	6(5.2)	0 (2.2)	0 (2.3)	11(9.3)	6(5.6)	7(7.2)	6(5.7)	0 (1.9)	0 (2.6)	6(6)	0 (1.8)	0 (1.8)	
	avg. rank		1.89		5.61		1.94		2						

Table 5: Rank and average epsilon indicator value ($\times 10^2$, between brackets) after 200, 700, and 1500 evaluations for MUBQP. The results are w.r.t. **Order_{static}**. A dash (—) indicates that the algorithm was not able to reach the corresponding budget within the maximum CPU time allowed in our experiments.

			#eval	Select _{local}			Select _{global}			Select _{BI}			Select _{NBI}		
				O=1	O=2	O=3	O=1	O=2	O=3	O=1	O=2	O=3	O=1	O=2	O=3
OptimMOEA/D	n = 25	200	1(371.5)	0 (194.6)	0 (252.2)	5(756.2)	4(639.8)	3(658.9)	3(515.5)	0 (179.3)	0 (234.9)	5(776.2)	4(666.2)	5(709.6)	
		700	4(280.7)	0 (33.7)	1(112)	6(638.3)	4(499.9)	4(388.3)	4(389.5)	0 (55.8)	1(98)	5(663.9)	4(565.1)	4(476)	
		1500	4(260)	0 (24.6)	0 (49.7)	8(576.6)	4(344.4)	4(285.6)	4(370.6)	1(47.6)	0 (46.8)	8(577.3)	4(528.6)	4(278.7)	
	avg. rank		0.67		5.06		0.83		5.28						
OptimPLS	n = 25	200	0 (290.1)	0 (399.1)	2(407.8)	10(707.4)	2(424)	2(469.3)	0 (360.2)	0 (333.2)	0 (335.2)	0 (431.7)	0 (233.6)	0 (315.6)	
		700	5(241.7)	1(71.8)	5(239.1)	11(648.2)	7(412.5)	6(350.6)	6(339.6)	0 (48.2)	1(108.1)	5(404.7)	0 (16.3)	1(59.7)	
		1500	5(239.4)	1(38.3)	2(90.4)	11(584.1)	7(408.6)	6(327.2)	6(333.1)	0 (36.3)	1(57.3)	6(383.3)	0 (14.9)	1(31.5)	
	avg. rank		1.89		5.22		1.5		1.44						
OptimPLS	n = 50	200	0 (1132.8)	3(2800.1)	3(2670.3)	3(2352.1)	3(2661.8)	3(2775.1)	0 (1260.6)	3(3077.3)	3(2741.5)	0 (1239.1)	3(2602.7)	3(2823.9)	
		700	0 (602.4)	3(1562)	—	7(2309.4)	3(1585)	—	1(857.1)	3(1373)	—	1(830.1)	3(1598.5)	—	
		1500	2(557.9)	2(793.6)	—	7(2273.3)	6(1401.1)	—	3(856.7)	0 (149.2)	—	3(830.1)	0 (115.1)	—	
	avg. rank		1.22		5.94		1.5		1.39						
OptimPLS	n = 25	200	0 (343.9)	0 (223)	0 (248.4)	11(765.4)	4(516.3)	1(459.7)	1(409.6)	0 (201.5)	0 (269.4)	2(440.3)	0 (341.2)	0 (363.6)	
		700	6(272.3)	0 (19.4)	1(116.7)	10(644.9)	9(486.2)	6(391.7)	6(310.7)	1(69.8)	1(74.9)	6(320.2)	0 (55.4)	1(107.1)	
		1500	6(264.3)	0 (15.3)	0 (50.4)	11(606)	9(470)	6(369.6)	6(287.5)	1(69.8)	1(58.5)	6(307.9)	0 (55.4)	1(68.4)	
	avg. rank		1.22		5.94		1.5		1.39						
OptimPLS	n = 50	200	0 (1187.5)	3(2139.8)	3(2106.1)	8(2939)	3(2471.9)	3(2410.4)	0 (1308.5)	3(2352.7)	3(2403.8)	0 (1309.1)	3(2471.9)	3(2410.4)	
		700	0 (781.3)	0 (1049.5)	—	7(2726)	6(1588.1)	—	1(1118.7)	0 (899.5)	—	0 (1015.3)	0 (945.9)	—	
		1500	3(734.3)	0 (142.1)	—	7(2600.6)	6(1588)	—	3(1054.8)	0 (177.8)	—	3(960.5)	0 (115.3)	—	
	avg. rank		1.22		5.94		1.5		1.39						

the underlying problem is quadratic, a Walsh order of 2 is always performing better. This is with no surprise since performance is tightly related to the ability of the Walsh surrogate to faithfully approximate the underlying landscape, which is directly related to the choice of the Walsh order.

Table 6: Sorted selection strategies by average rank ($L = \text{Select}_{\text{local}}$, $G = \text{Select}_{\text{global}}$, $B_C = \text{Select}_{\text{BI}}$, $B_N = \text{Select}_{\text{NBI}}$).

		Optim _{MOEA/D}	Optim _{MLS}	Optim _{PLS}
MNK-landscapes	$k = 0$	$L \prec B_C \prec B_N \prec G$	$B_N \prec B_C \prec L \prec G$	$B_N \prec L \prec B_C \prec G$
MNK-landscapes	$k = 1$	$B_C \prec L \prec B_N \prec G$	$B_N \prec B_C \prec L \prec G$	$B_N \prec L \prec B_C \prec G$
MNK-landscapes	$k = 2$	$B_C \prec L \prec B_N \prec G$	$L \prec B_C \prec B_N \prec G$	$L \prec B_C \prec B_N \prec G$
MUBQP		$L \prec B_C \prec G \prec B_N$	$B_N \prec B_C \prec L \prec G$	$L \prec B_N \prec B_C \prec G$

Overall, we can hence conclude that the rank obtained when using a surrogate optimizer with a particular selection strategy changes consistently with the choice of the order, suggesting that there is an optimal setting that depends on the underlying landscape. Additionally, we observe that the ranks may vary depending on the considered budget. This indicates that the different algorithm configurations may expose different anytime behaviors. This is to be analyzed in more detail afterwards, in light of the other proposed *dynamic* strategies for setting the Walsh order.

Secondly, we can clearly see that the selection strategy providing the best rank is *different* depending on which surrogate optimizer is considered. This is an important finding showing that there is a strong dependency between these two components, i.e., for optimal performance, the selection strategy has to be configured differently depending on the adopted surrogate optimizer. More specifically, the two local search optimizers (Optim_{MLS} in Table 3 and Optim_{PLS} in Table 4) perform at their best when using the Select_{NBI} strategy for MNK-landscapes with $k = 0$ and $k = 1$, whereas using the Select_{local} strategy is a slightly better choice when $k = 2$. In contrast, for MUBQP (Table 5), these two optimizers perform at their best when using different selection strategies: Optim_{PLS} performs better with the Select_{local} strategy, and Optim_{MLS} performs better with the Select_{NBI} and Select_{BI} strategies. Looking at the third Optim_{MOEA/D} optimizer, the situation is seemingly different: Select_{BI} is the best performing strategy for MNK-landscapes (in Table 2), whereas Select_{local} is better for MUBQP (in Table 5). We also found that these two selection strategies are definitely a better choice when using Optim_{MOEA/D} compared to the Select_{NBI} and Select_{global} strategies. In contrast to the work in [68], where Select_{local} is recommended, these observations show that other selection strategies can be more accurate, depending on both the considered problem and the surrogate optimizer.

To summarize these complex dependencies, we provide in Table 6 an overview of the relative performance of the selection strategies for the different surrogate optimizers, obtained by sorting the selection strategies according to the average ranks from Tables 2, 3, 4 and 5 over the considered instances and budgets (as given by the ‘avg. rank’ rows). Looking at the impact of the selection strategy over the different surrogate optimizers, it is interesting to remark that for Optim_{PLS} and Optim_{MLS}, the Select_{NBI} strategy is to be preferred, with an average rank of 1.40 and 1.29 respectively, whereas its average rank is of 4.36 when used in combination with Optim_{MOEA/D}. In other words, the Select_{NBI} strategy can be viewed as a relatively good selection strategy, except for Optim_{MOEA/D}. Besides, the Select_{global} strategy shows the worst results overall and independently

of the considered surrogate optimizer.

In the rest of the paper, and unless explicitly stated, we shall always consider the best performing selection strategy when dealing with a particular surrogate optimizer, as depicted in Table 6.

5.2. Impact of the Surrogate Optimizer

In this section, we compare the different surrogate optimizers when different static Walsh orders are used. This is depicted in Fig. 1 and 2 showing the convergence (anytime) profile of the different configurations respectively for MNK-landscapes and MUBQP.

As a first observation, we clearly see that the $\text{Optim}_{\text{MLS}}$ and $\text{Optim}_{\text{PLS}}$ optimizers have significantly better convergence profiles compared to $\text{Optim}_{\text{MOEA/D}}$. This is consistent over all considered instances, with very few exceptions. At this stage of the analysis, we recall that the proposed Walsh-based surrogate approach uses the decomposition paradigm to both structure the population, and to select the next solution to evaluate. Hence, this observation shows that, (i) such an approach does *not* necessarily need to be configured with a decomposition-based *inner* optimizer to handle the underlying surrogate; and more importantly, (ii) other more accurate optimizers should be used specifically to the underlying Walsh surrogate. This is again to be contrasted with the work in [68], where the decomposition-based MOEA/D algorithm was employed at the different design stages.

Let us now analyze in more detail the two $\text{Optim}_{\text{MLS}}$ and $\text{Optim}_{\text{PLS}}$ optimizers, which were found to lead to a better approximation quality. Although overall they expose a similar behavior, their relative anytime profile depends both on the difficulty of the tackled problem, and on the complexity of the Walsh surrogate, as implied by the choice of the Walsh order. This is discussed in the next paragraphs by splitting our experimented instances into three classes: (i) linear, i.e., MNK-landscapes with $k = 0$, (ii) quadratic, i.e., MNK-landscapes with $k = 1$ and MUBQP, and (iii) cubic, i.e., MNK-landscapes with $k = 2$. It is important to remark that the order of the exact Walsh transform of any function in these three classes is bounded respectively by 1, 2, and 3.

For the two linear instances (with $k = 0$), $\text{Optim}_{\text{PLS}}$ used in combination with the (exact) Walsh order of 1 converges very quickly to a high-quality approximation set, and $\text{Optim}_{\text{MLS}}$ is only able to obtain a better quality when a higher number of evaluations is affordable. When using an over-estimated Walsh order of 2 or 3, $\text{Optim}_{\text{MLS}}$ performs significantly better than $\text{Optim}_{\text{PLS}}$ independently of the available budget.

For quadratic instances, when setting the Walsh order to the exact transform order of 2, we found that $\text{Optim}_{\text{MLS}}$ and $\text{Optim}_{\text{PLS}}$ have seemingly the same convergence profile for $n = 25$. By contrast, for $n = 50$, $\text{Optim}_{\text{PLS}}$ is slightly better under a restricted budget, while $\text{Optim}_{\text{MLS}}$ is slightly better under higher budgets, which is to recall their behavior for linear instances. When the Walsh order is set to an underestimated value of 1, the approximation quality obtained with all optimizers drops down in comparison to an (exact) order of 2. Interestingly, using an underestimated order value does not prevent to find improvements during the very first iterations, that

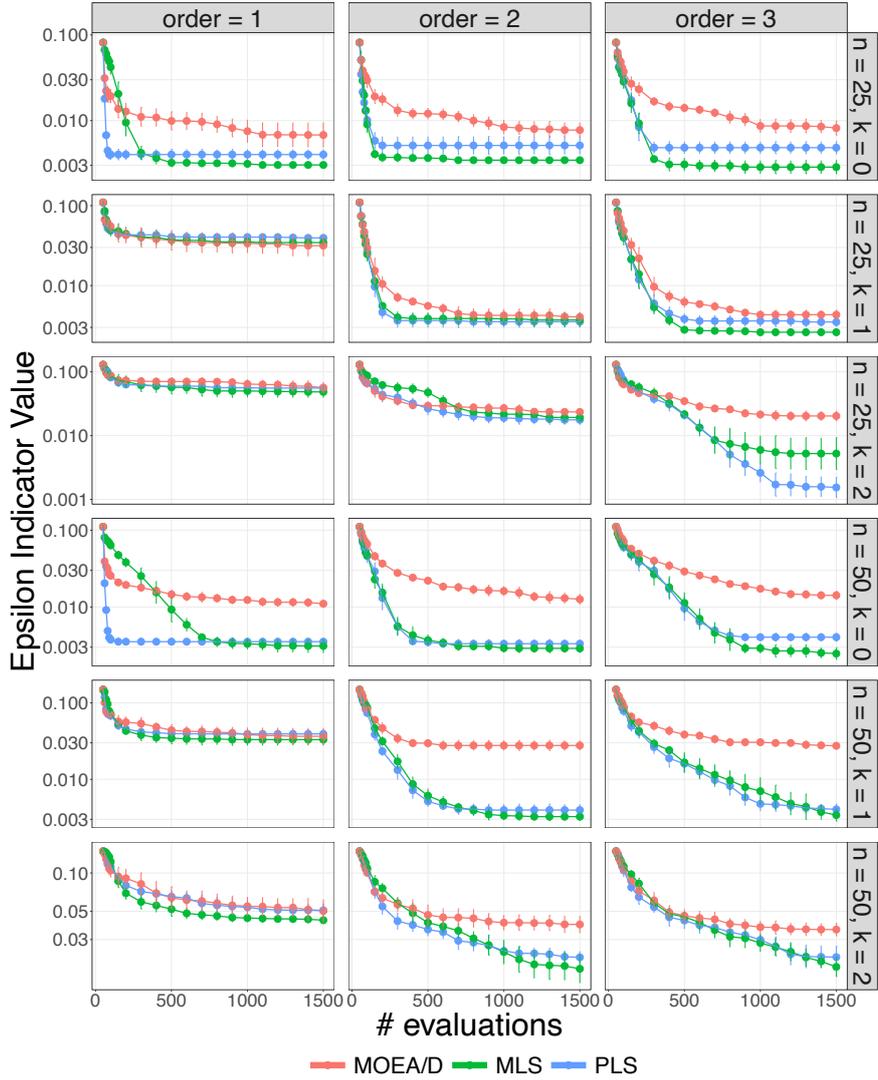


Figure 1: Convergence profiles when using the surrogate optimizers with their corresponding best selection strategy for MNK-landscapes.

is, with a very restricted budget. When the Walsh order is set to an overestimated value of 3, the overall approximation quality of both optimizers seems to be comparable to the exact setting of the order.

Finally, for the most complex cubic instances ($k = 2$) and for $n = 25$, an underestimated order of 1 or 2 is not competitive with the exact transform order 3, independently of the surrogate optimizer. For dimension $n = 50$, all algorithms are clearly performing very poorly when an underestimated order of 1 is used. However, the situation improves when using an order of 2, which indicates that for such difficult problems, a sufficiently large but non-exact Walsh order still allows for a reasonable approximation.

From this set of observations, we can draw two general conclusions. First, the two local search

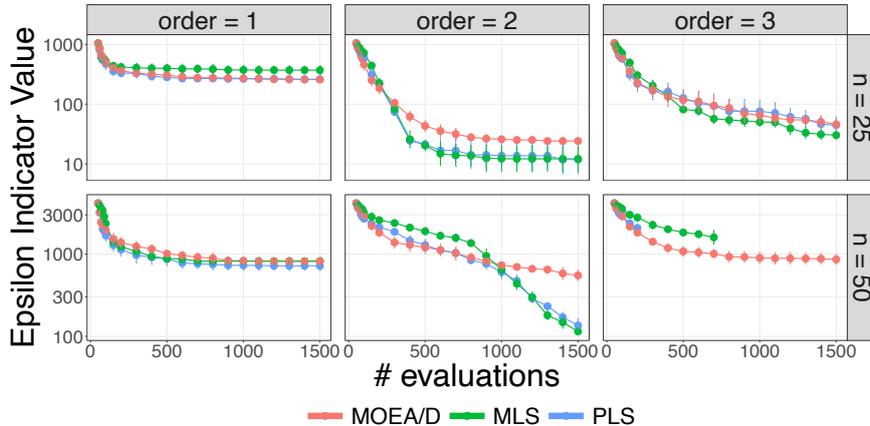


Figure 2: Convergence profiles when using the surrogate optimizers with their corresponding best selection strategy for MUBQP.

optimizers $\text{Optim}_{\text{MLS}}$ and $\text{Optim}_{\text{PLS}}$ are efficient in dealing with the discrete Walsh surrogate, which is to contrast with the $\text{Optim}_{\text{MOEA/D}}$ optimizer. Second, an effective surrogate optimizer alone is not able to provide a good performance independently of the available budget and the Walsh order setting. The impact of the order choice strategy is studied next, hence addressing the last component of the considered framework.

5.3. Impact of the Walsh Order

To study the impact of the strategy used for the Walsh order, we focus on $\text{Optim}_{\text{MLS}}$ and $\text{Optim}_{\text{PLS}}$, since they were found to substantially outperform $\text{Optim}_{\text{MOEA/D}}$. This is shown in Fig. 3 and 4, respectively for MNK-landscapes and MUBQP.

From these two figures, it becomes very clear that underestimating the order value (relatively to the exact Walsh transform) has a dramatic impact on the overall performance. Besides, overestimating the order value does not always allow for competitive results. This might seem surprising, since overestimating the order makes the Walsh model more complex, and should not decrease its accuracy. While such a claim sounds true in theory, increasing the Walsh order makes the model much more challenging to fit accurately in practice. To illustrate this issue, we show in Fig. 5 the mean absolute error of the Walsh model as a function of the size of the training dataset. We clearly see that overestimating the order can lead to a worst fitting quality depending on the characteristic on the degree of non-linearity k , and the size of the training set. This is to be attributed to the fact that the more complex the model is, the largest the training data should be in order for the Lasso regression to be successful in finding a good fit.

Furthermore, in comparison to a static setting of the Walsh order ($\text{Order}_{\text{static}}$), the considered dynamic strategies ($\text{Order}_{\text{random}}$ and $\text{Order}_{\text{greedy}}$) are clearly a better option. Interestingly, there is no clear separation between the greedy and the random strategy independently of the underlying surrogate optimizer, problem instance and budget. For example, we can see that $\text{Order}_{\text{greedy}}$ out-

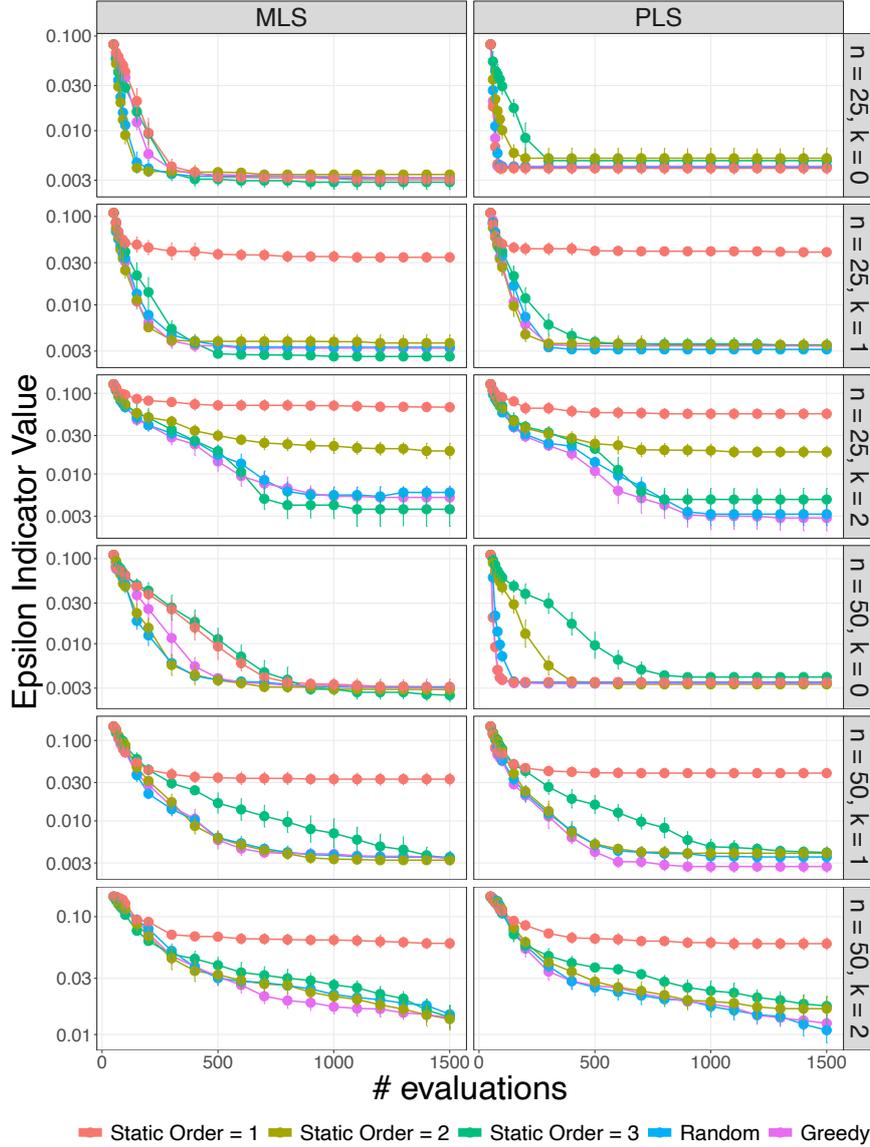


Figure 3: Convergence profile of when using different Walsh order settings with the $\text{Select}_{\text{NBI}}$ strategy for MNK-landscapes.

performs $\text{Order}_{\text{random}}$ for MUBQP instances of size $n = 50$. However, this is no more true for $n = 25$ and the $\text{Optim}_{\text{PLS}}$ surrogate optimizer. Similarly, $\text{Order}_{\text{greedy}}$ obtains a better convergence profile than $\text{Order}_{\text{random}}$ when using $\text{Optim}_{\text{PLS}}$ for a MNK-landscapes with $n = 50$ and $k = 0$, whereas this is no more true when combining $\text{Order}_{\text{greedy}}$ with $\text{Optim}_{\text{MLS}}$. Overall, we conclude that both dynamic strategies for setting the Walsh order obtain very competitive results. A dynamic order strategy is always better than an underestimated static order, and is better than an overestimated static order in most cases. Surprisingly, it is even better than a static strategy using the exact Walsh transform order for some instances; e.g. MNK-landscapes with $n = 50$ and $k = 1$. These

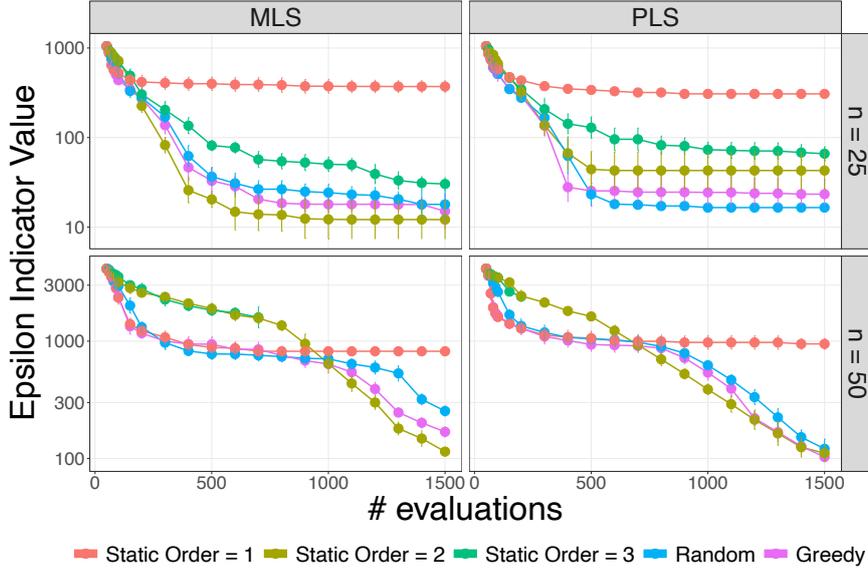


Figure 4: Convergence profile when using different Walsh order settings with the $\text{Select}_{\text{NBI}}$ strategy for MUBQP.

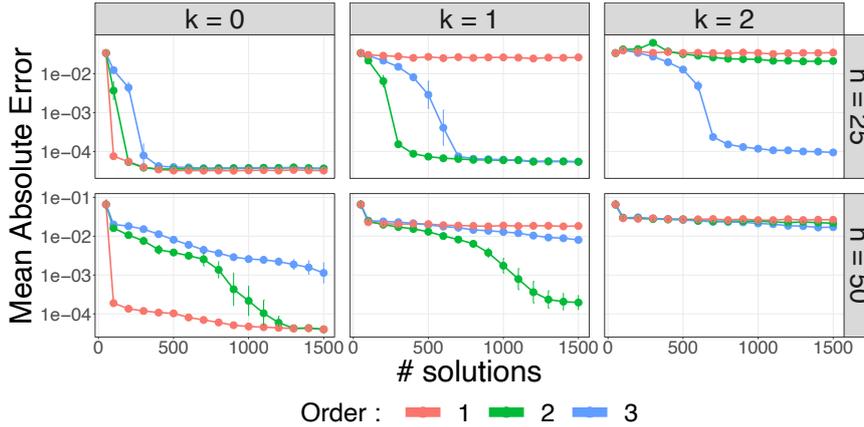


Figure 5: Mean absolute error (MAE) obtained by the Walsh surrogate according to the number of solutions in the training dataset with the $\text{Optim}_{\text{MLS}}$ optimizer combined to the $\text{Select}_{\text{NBI}}$ strategy for order $O \in \{1, 2, 3\}$ on MNK-landscapes.

results suggest that more sophisticated dynamic strategies for setting the Walsh order are worth to be investigated in the future.

5.4. Approximation quality of SMCO/w vs. Surrogate-less algorithms

In this section, we show the benefits behind using a surrogate model to assist the search process. For this purpose, we compare the approximation quality obtained when running the three algorithms MOEA/D, PLS, and MLS on the original problems (without using Walsh surrogate). Our results are reported in Fig. 6. We only show the results obtained with the $\text{Select}_{\text{NBI}}$ strategy and the $\text{Order}_{\text{greedy}}$ strategy, since these two strategies were found to provide fairly good results over

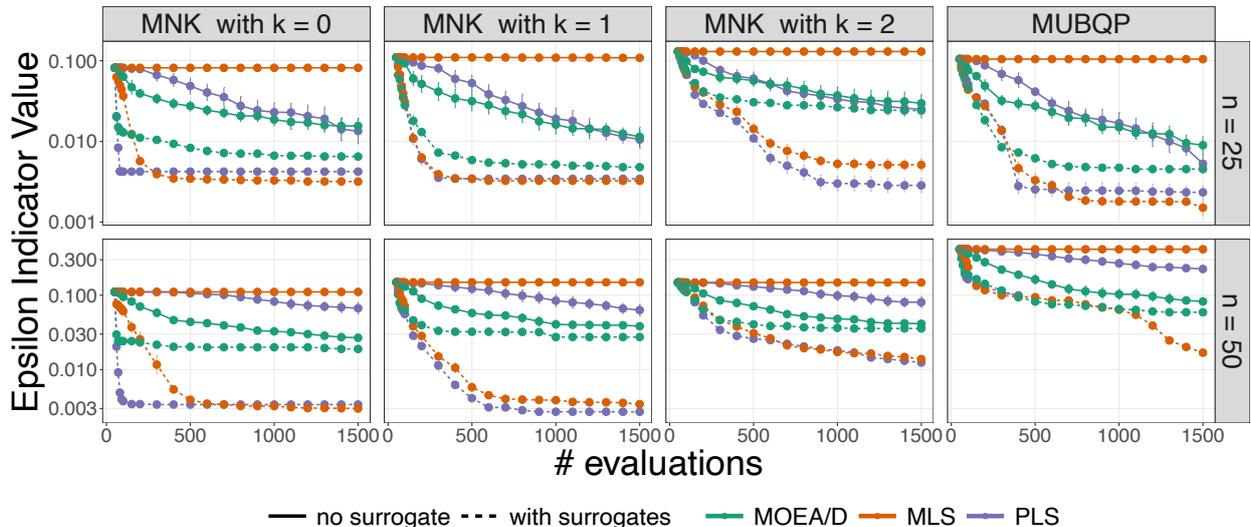


Figure 6: Comparison of surrogate-less approaches with SMCO/w using a greedy Walsh order ($\text{Order}_{\text{greedy}}$) and the best normalized improvement selection strategy ($\text{Select}_{\text{NBI}}$). For MUBQP, the indicator has been scaled by a factor 10^{-2} for a better readability of the figures.

all the experimented instances. We can see that the so-obtained surrogate-assisted algorithms lead to substantially better approximation sets, independently of the considered instance and budget. This provides evidence on the accuracy of the designed approach and its effectiveness in reducing the number of function evaluations required to compute a high quality approximation set.

The results presented in Fig. 6 hold under the standard assumption that the evaluation function is very costly, in the sense that its CPU time cost dominates all the other parts of the considered algorithms. However, it is well-known that using a surrogate model, like in any machine learning technique, can incur some computational time. Mitigating such a computational cost with respect to the real cost of the expensive evaluation can be important depending on the problem considered in practice. This aspect is only discussed to a small extent in the literature. For completeness, and although our work is not concerned with a particular target application problem, we conclude our analysis by conducting additional experiments allowing us to fairly discuss the computational CPU times incurred by using a Walsh surrogate and its implications, while adopting the same benchmark-oriented methodology.

5.4.1. Considerations on the Walsh surrogate CPU time

To analyze the CPU time cost incurred by the Walsh surrogate, we adopt the following experimental procedure. For each problem instance, we generate successively $s \in \{250, 500, 1000, 1500\}$ random binary string solutions. The Walsh surrogate is trained as discussed previously for each order $O \in \{1, 2, 3\}$. The average CPU (*training*) time required for this phase is then reported (over 100 independent repetitions). Using each trained model, we predict the objective value for a randomly generated binary string solution, and the average CPU (*prediction*) time is reported as well.

Table 7: CPU time of Walsh surrogate training phase (in *seconds*) and prediction phase (in *milliseconds*) using a sample size $s \in \{250, 500, 1000, 1500\}$. Results are shown for every static order O and each problem instance.

		s	$O = 1$				$O = 2$				$O = 3$			
			MNK-landscapes			MUBQP	MNK-landscapes			MUBQP	MNK-landscapes			MUBQP
			$k = 1$	$k = 2$	$k = 3$		$k = 1$	$k = 2$	$k = 3$		$k = 1$	$k = 2$	$k = 3$	
Training (in s)	$n = 25$	250	0.04	0.05	0.06	0.06	0.17	0.19	0.23	0.19	1.35	1.49	1.54	1.25
		500	0.04	0.05	0.05	0.06	0.27	0.28	0.28	0.28	2.43	2.51	3.03	2.36
		1000	0.05	0.05	0.06	0.06	0.53	0.52	0.55	0.51	4.50	4.60	4.85	4.48
		1500	0.05	0.06	0.07	0.07	0.79	0.82	0.81	0.78	6.60	6.57	6.74	6.60
	$n = 50$	250	0.07	0.08	0.09	0.16	0.67	0.78	0.77	0.72	13.82	13.66	14.63	10.49
		500	0.07	0.09	0.10	0.17	1.15	1.23	1.39	1.24	27.68	24.19	28.31	19.05
		1000	0.09	0.10	0.12	0.18	2.23	2.28	2.94	2.38	43.38	35.01	48.07	32.86
		1500	0.10	0.12	0.13	0.19	3.31	3.30	3.55	3.37	56.00	52.96	58.71	46.64
Prediction (in ms)	$n = 25$	250	0.06	0.06	0.06	0.06	0.06	0.13	0.41	0.51	0.07	0.51	0.60	4.27
		500	0.06	0.06	0.06	0.06	0.06	0.11	0.55	0.55	0.07	0.16	1.12	4.58
		1000	0.06	0.06	0.06	0.06	0.07	0.11	0.60	0.59	0.07	0.11	0.32	4.47
		1500	0.06	0.06	0.06	0.06	0.07	0.18	0.63	0.61	0.08	0.12	0.30	4.34
	$n = 50$	250	0.10	0.10	0.10	0.10	0.24	0.46	0.50	2.00	1.11	1.32	1.77	14.53
		500	0.10	0.10	0.10	0.10	0.12	0.26	0.81	2.13	0.50	1.80	2.40	41.41
		1000	0.10	0.10	0.10	0.10	0.12	0.21	1.70	2.02	0.25	0.58	3.57	34.21
		1500	0.10	0.10	0.10	0.10	0.13	0.22	2.40	2.23	0.20	0.50	3.41	34.02

These additional experiments were conducted in Python 3.7 on a personal Intel Core i7, 6-core processor (2.6 GHZ, 16 GB RAM) under macOS, without any further low-level code optimization. The obtained average CPU times for training and prediction are reported in Table 7.

We can first see that all CPU times (for both training and prediction) are strongly correlated with the considered Walsh order O and the problem size n , i.e., CPU times are higher for larger orders and larger dimensions. This is without surprise since higher orders/dimensions imply a more complex model, with more Walsh coefficients to estimate.

The CPU time required for the prediction phase appears to be overall relatively low, i.e., from few to at most some dozens of milliseconds. The CPU time for prediction can be viewed as the cost of the inferred surrogate function, that will be optimized internally by the embedded surrogate optimizer (e.g., MOEA/D, PLS, MLS). Hence, we can conclude that the obtained Walsh surrogate function is relatively cheap. Having such a cheap surrogate function is very desirable in order to be later extensively tackled by the embedded evolutionary optimizer. However, on the other hand, the CPU training time appears to be more substantial. In the worst case scenarios, using the largest training sample of size 1500 and the largest dimension $n = 50$, the training cost ranges from 0.13 seconds in average (over all instances) for a Walsh order $O = 1$, to 3.38 seconds for $O = 2$, and 53.58 seconds for $O = 3$. Nonetheless, such CPU times can still be considered as reasonable in a practical expensive context, where the objective function cost can typically be at least on the order of minutes or hours. Notice in fact that the Walsh model needs to be trained once, before serving as a cheap substitute that can in turn be searched more extensively, hopefully in less time than the cost of one real expensive evaluation. For example, using the highest Walsh order $O = 3$, it costs less than 2 minutes for an internal surrogate optimizer embedded in SMCO/w to evaluate the expected quality of more than 10^4 (respectively 10^3) candidate solutions for the

hardest MNK-landscape with $n = 50$ and $k = 2$ (respectively MUBQP with $n = 50$).

5.4.2. Considerations on the acceleration ratio

Having these CPU times in mind, we push our analysis further by eliciting how many function evaluations a surrogate-less optimizer needs in order to reach the same approximation quality than the proposed Walsh-assisted SMCO/w approach. This shall fairly complement the observations from Fig. 6, where all algorithms including the surrogate-less ones have been run with exactly the same number of real (expensive) function evaluations. More precisely, we execute the two surrogate-less algorithms, MOEA/D and PLS, using 10^5 real function evaluations, which is a reasonably high budget for the considered instances. We then compute the number of function evaluations required for each surrogate-less algorithm to reach an approximation set having a given target epsilon indicator value. We choose different target values so as to render to relative difference with the SMCO/w approach when executed with a range of different budgets of $B_q = q \times 100$ evaluations, with $q \in [2, 15]$. As such, we choose a target value t_q to be the average epsilon indicator value attained by SMCO/w after B_q evaluations, and using the greedy Walsh order ($\text{Order}_{\text{greedy}}$) and the best normalized improvement selection strategy ($\text{Select}_{\text{NBI}}$), as previously considered in Fig. 6. This experiment is repeated for 10 different seeds. Let B be the number of true function evaluations needed by a surrogate-less algorithm to hit target t_q as defined previously. We then define the *acceleration ratio* to be: B/B_q . Since SMCO/w was found to be always better than the considered surrogate-less approaches (Fig. 6), the acceleration ratio allows us to *quantify* in a more fine-grained manner *by how much* SMCO/w is superior. In fact, an acceleration ratio value of $x > 1$, for some target t_q , means that the corresponding surrogate-less algorithm needs x times more function evaluations to reach the same average approximation quality attained by SMCO/w when using B_q evaluations. The average acceleration ratios are depicted in Fig. 7 under the different experimented conditions, and as a function of the target affordable number of real function evaluations B_q . Notice that it might happen that one execution of the surrogate-less algorithms (MOEA/D and PLS) terminates *without* hitting a given target value t_q . In that case, following the standard definition of the so-called empirical running time (ERT) [80], we adopt an empirical-based assessment technique, where the surrogate-less algorithm is assumed to be restarted with a seed sampled uniformly from the 10 considered ones. In the case none of the 10 seeds were successful in hitting a target value t_q , Fig. 7 does not show any ratio; hence, indicating that all the 10 executions of the surrogate-less algorithm failed in hitting the target t_q after exhausting the 10^5 evaluations.

We can observe different trends depending on the characteristics of considered problem instance. For the easiest MNK-landscape with $n = 25$ and $k = 0$, PLS performs better than MOEA/D, and the acceleration ratio due to use of the Walsh surrogate decreases as the affordable budget B_q increases. For example for a very restricted budget $B_q = 200$ evaluations, SMCO/w requires 11.5 (respectively, 47.5) times less evaluations than PLS (respectively, MOEA/D), whereas it requires around 2 (respectively, 27) times less evaluations than PLS (respectively, MOEA/D), for the highest

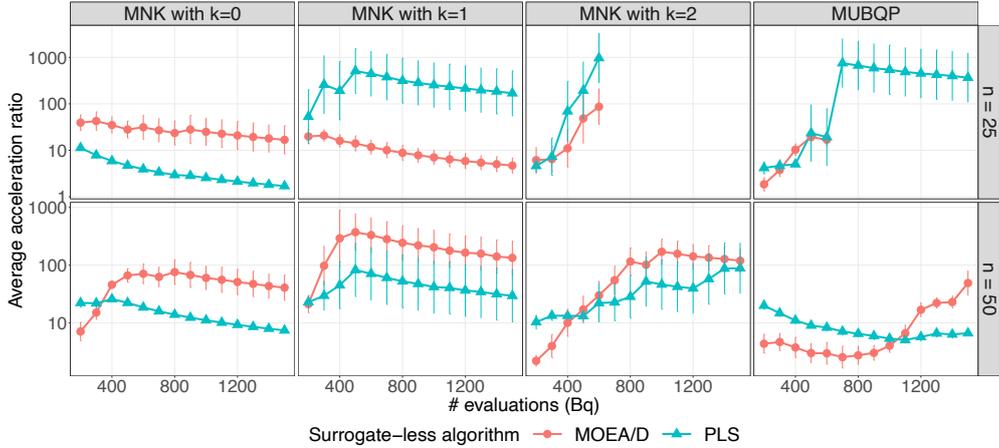


Figure 7: Average acceleration ratio between surrogate-less algorithms and SMCO/w (using $\text{Order}_{\text{greedy}}$ and $\text{Select}_{\text{NBI}}$) for every target approximation quality t_q and as a function of the number of corresponding function evaluations B_q .

budget $B_q = 1500$. For the hardest MNK-landscape with $n = 50$ and $k = 2$, MOEA/D performs better than PLS for a very restricted budget $B_q = 200$, with an acceleration ratio of about 2.1 in favor of SMCO/w. With a high budget $B_q = 1500$, SMCO/w requires around 150 times less evaluations than both MOEA/D and PLS. For MUBQP, only PLS is able to hit all the targets. The acceleration ratio is again found to be substantial depending on the different targets and dimensions.

Generally speaking, the acceleration ratios presented in Fig. 7 indicate that the SMCO/w approach requires overall many order of magnitudes less expensive evaluations than its surrogate-less counterparts to attain a high-quality approximation set. Combined with the CPU computational times incurred by the Walsh surrogate from Table 7, these results provide a relatively clear picture on the benefits of using the Walsh-assisted SMCO/w approach when tackling a concrete optimization scenario having a computationally expensive objective vector.

6. Conclusion

In this paper, we have investigated the design and analysis of a modular surrogate-assisted framework for expensive multi-objective combinatorial optimization using Walsh basis. Based on extensive experiments, we provided a systematic study of the combined effects of different design components. In particular, we found that there is a non-trivial interaction between the strategy allowing to optimize the surrogate, and the strategy used for selecting the next solution to evaluate. We showed that a selection strategy based on the predicted improvements of candidate solutions with respect to a set of decomposed single-objective sub-problems is highly effective. Besides, we found that local search is to be preferred for the inner optimization of the discrete Walsh surrogate. Finally, we highlighted the importance of the Walsh order, and we proposed simple

dynamic strategies to accommodate the search to a range of black-box optimization functions having different degrees of difficulty.

We leave open a number of research issues are worthwhile exploring. For instance, we rely on the decomposition paradigm to structure the population of evaluated solutions, and to estimate which new offspring is the most promising for a true evaluation. It would be interesting to support different selection strategies based on other multi-objective paradigms. Another challenging issue is to deal with problems having a high number of objectives, as well as to address other discrete optimization domains, such as permutation problems, for which other single-objective surrogates exist in the specialized literature. A main question is then to study at what extent our findings hold for different objective and decision spaces.

Finally, the work conducted in this paper is of fundamental nature in the sense that we adopt a benchmark-oriented validation methodology, without targeting a specific real-world optimization setting. Considering a real-world application implies to deal with a number of additional and difficult questions. For instance, it could happen that for some complex applications, a restricted, yet not empty, subset of Walsh functions of possibly different high orders, are to be used so that the considered real-world problem can be fit accurately. This suggests to derive improved learning techniques to discover those functions without experiencing scalability issues at the training phase. Besides, the amount of computational resources, in terms of parallel and distributed CPU cores, that are available in practice when tackling a real-world expensive application, can be substantial. This suggests that an ensemble of Walsh surrogates can be coordinated and trained in parallel, hence allowing one to sample an eventually large number of solutions that cover the Pareto front in a more effective way. It is our hope that the comprehensive study conducted in this paper will accelerate the establishment of a unified methodology for designing and analyzing discrete surrogate-assisted multi-objective optimization algorithms.

Acknowledgement

This work was partially supported by the French National Research Agency (ANR-16-CE23-0013-01) and the Research Grants Council of Hong Kong (A-CityU101/16). We gratefully acknowledge the anonymous reviewers for their comments and suggestions that helped us improving our work.

References

- [1] J. Stork, M. Friese, M. Zaefferer, T. Bartz-Beielstein, A. Fischbach, B. Breiderhoff, B. Naujoks, T. Tusar, Open issues in surrogate-assisted optimization, in: T. Bartz-Beielstein, B. Filipic, P. Korosec, E. Talbi (Eds.), *High-Performance Simulation-Based Optimization*, Vol. 833 of *Studies in Computational Intelligence*, Springer, 2020, pp. 225–244.
- [2] T. Chugh, K. Sindhya, J. Hakanen, K. Miettinen, A survey on handling computationally expensive multiobjective optimization problems with evolutionary algorithms, *Soft Computing* 23 (9) (2019) 3137–3166.

- [3] T. Bartz-Beielstein, M. Zaefferer, Model-based methods for continuous and discrete global optimization, *Applied Soft Computing* 55 (2017) 154–167.
- [4] T. Bartz-Beielstein, A survey of model-based methods for global optimization, in: *International Conference on Bioinspired Optimization Methods and Their Applications (BIOMA)*, 2016, pp. 1–18.
- [5] Y. Jin, Surrogate-assisted evolutionary computation: Recent advances and future challenges, *Swarm and Evolutionary Computation* 1 (2) (2011) 61–70.
- [6] Y. Jin, H. Wang, T. Chugh, D. Guo, K. Miettinen, Data-driven evolutionary optimization: An overview and case studies, *IEEE Transactions on Evolutionary Computation* 23 (3) (2019) 442–458.
- [7] K. Deb, R. Hussein, P. C. Roy, G. T. Pulido, A taxonomy for metamodeling frameworks for evolutionary multiobjective optimization, *IEEE Transactions on Evolutionary Computation* 23 (1) (2019) 104–116.
- [8] D. Horn, T. Wagner, D. Biermann, C. Weihs, B. Bischl, Model-based multi-objective optimization: Taxonomy, multi-point proposal, toolbox and benchmark, in: *Evolutionary Multi-Criterion Optimization*, Springer International Publishing, 2015, pp. 64–78.
- [9] R. Allmendinger, M. T. M. Emmerich, J. Hakanen, Y. Jin, E. Rigoni, Surrogate-assisted multicriteria optimization: Complexities, prospective solutions, and business case, *Journal of Multi-Criteria Decision Analysis* 24 (1-2) (2017) 5–24.
- [10] K. Miettinen, *Nonlinear multiobjective optimization*, Vol. 12 of *International series in operations research and management science*, Kluwer, 1998.
- [11] M. Ehrgott, *Multicriteria optimization*, 2nd Edition, Springer, 2005.
- [12] K. Deb, *Multi-Objective Optimization using Evolutionary Algorithms*, John Wiley & Sons, 2001.
- [13] C. A. Coello Coello, G. B. Lamont, D. A. Van Veldhuizen, *Evolutionary Algorithms for Solving Multi-Objective Problems*, 2nd Edition, Springer, 2007.
- [14] D. Krige, A statistical approach to some basic mine valuation problems on the Witwatersrand, *Journal of the Southern African Institute of Mining and Metallurgy* 52 (6) (1951) 119–139.
- [15] C. E. Rasmussen, Gaussian processes in machine learning, in: O. Bousquet, U. von Luxburg, G. Rätsch (Eds.), *Advanced Lectures on Machine Learning, ML Summer Schools*, Canberra, Australia, February 2-14, Tübingen, Germany, August 4-16, Revised Lectures, Vol. 3176 of *Lecture Notes in Computer Science*, Springer, 2003, pp. 63–71.
- [16] H. Drucker, C. J. C. Burges, L. Kaufman, A. J. Smola, V. Vapnik, Support vector regression machines, in: M. Mozer, M. I. Jordan, T. Petsche (Eds.), *Advances in Neural Information Processing Systems*, NIPS, MIT Press, 1996, pp. 155–161.
- [17] M. D. Buhmann, *Radial Basis Functions - Theory and Implementations*, Vol. 12 of *Cambridge monographs on applied and computational mathematics*, Cambridge University Press, 2009.
- [18] I. J. Goodfellow, Y. Bengio, A. C. Courville, *Deep Learning*, Adaptive computation and machine learning, MIT Press, 2016.
- [19] L. Breiman, J. H. Friedman, R. A. Olshen, C. J. Stone, *Classification and Regression Trees*, Chapman and Hall/CRC, Monterey, CA, 1984.
- [20] L. Breiman, Random forests, *Machine Learning* 45 (1) (2001) 5–32.
- [21] A. Trivedi, D. Srinivasan, K. Sanyal, A. Ghosh, A survey of multiobjective evolutionary algorithms based on decomposition, *IEEE Transactions on Evolutionary Computation* 21 (3) (2017) 440–462.
- [22] Q. Zhang, H. Li, MOEA/D: A multiobjective evolutionary algorithm based on decomposition, *IEEE Transactions on Evolutionary Computation* 11 (6) (2007) 712–731.
- [23] K. Deb, A. Pratap, S. Agarwal, T. Meyarivan, A fast and elitist multiobjective genetic algorithm: NSGA-II, *IEEE Transactions on Evolutionary Computation* 6 (2) (2002) 182–197.
- [24] K. Deb, H. Jain, An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, Part I: Solving problems with box constraints, *IEEE Transactions on Evolutionary*

- Computation 18 (4) (2014) 577–601.
- [25] H. Jain, K. Deb, An evolutionary many-objective optimization algorithm using reference-point based nondominated sorting approach, Part II: Handling constraints and extending to an adaptive approach, *IEEE Transactions on Evolutionary Computation* 18 (4) (2014) 602–622.
 - [26] N. Beume, B. Naujoks, M. Emmerich, SMS-EMOA: Multiobjective selection based on dominated hypervolume, *European Journal of Operational Research* 181 (3) (2007) 1653–1669.
 - [27] E. Zitzler, S. Künzli, Indicator-based selection in multiobjective search, in: *Parallel Problem Solving from Nature - PPSN VIII, Lecture Notes in Computer Science*, Springer, 2004, pp. 832–842.
 - [28] J. Knowles, ParEGO: A hybrid algorithm with on-line landscape approximation for expensive multiobjective optimization problems, *IEEE Transactions on Evolutionary Computation* 10 (1) (2006) 50–66.
 - [29] Q. Zhang, W. Liu, E. Tsang, B. Virginas, Expensive multiobjective optimization by MOEA/D with gaussian process model, *IEEE Transactions on Evolutionary Computation* 14 (3) (2010) 456–474.
 - [30] S. Zapotecas Martinez, C. A. Coello Coello, MOEA/D assisted by RBF networks for expensive multi-objective optimization problems, in: *The Genetic and Evolutionary Computation Conference, GECCO, ACM*, 2013, pp. 1405–1412.
 - [31] T. Chugh, Y. Jin, K. Miettinen, J. Hakanen, K. Sindhya, A surrogate-assisted reference vector guided evolutionary algorithm for computationally expensive many-objective optimization, *IEEE Transactions on Evolutionary Computation* 22 (1) (2018) 129–142.
 - [32] R. Hussein, K. Deb, A generative kriging surrogate model for constrained and unconstrained multi-objective optimization, in: *The Genetic and Evolutionary Computation Conference, GECCO, ACM*, 2016, pp. 573–580.
 - [33] W. Ponweiser, T. Wagner, D. Biermann, M. Vincze, Multiobjective optimization on a limited budget of evaluations using model-assisted s-metric selection, in: *Parallel Problem Solving from Nature – PPSN X, Lecture Notes in Computer Science*, Springer, 2008, pp. 784–794.
 - [34] J. Li, P. Wang, H. Dong, J. Shen, C. Chen, A classification surrogate-assisted multi-objective evolutionary algorithm for expensive optimization, *Knowledge-Based Systems* 242 (2022) 108416.
 - [35] L. Pan, C. He, Y. Tian, H. Wang, X. Zhang, Y. Jin, A classification-based surrogate-assisted evolutionary algorithm for expensive many-objective optimization, *IEEE Transactions on Evolutionary Computation* 23 (1) (2019) 74–88.
 - [36] H. Dong, J. Li, P. Wang, B. Song, X. Yu, Surrogate-guided multi-objective optimization (SGMOO) using an efficient online sampling strategy, *Knowledge-Based Systems* 220 (2021) 106919.
 - [37] N. Bervagliari, B. Derbel, A. Liefoghe, H. Aguirre, K. Tanaka, Surrogate-assisted multiobjective optimization based on decomposition: A comprehensive comparative analysis, in: *The Genetic and Evolutionary Computation Conference, GECCO, ACM*, 2019, pp. 507–515.
 - [38] M. Zaefferer, Surrogate models for discrete optimization problems, Ph.D. thesis, University of Dortmund (2018).
 - [39] A. Moraglio, A. Kattan, Geometric generalisation of surrogate model based optimisation to combinatorial spaces, in: *Evolutionary Computation in Combinatorial Optimization*, Springer, 2011, pp. 142–154.
 - [40] M. Zaefferer, J. Stork, M. Friese, A. Fischbach, B. Naujoks, T. Bartz-Beielstein, Efficient global optimization for combinatorial problems, in: *The Genetic and Evolutionary Computation Conference, GECCO, ACM*, 2014, pp. 871–878.
 - [41] R. Baptista, M. Poloczek, Bayesian optimization of combinatorial structures (2018). [arXiv:1806.08838](https://arxiv.org/abs/1806.08838).
 - [42] H. Dong, P. Wang, B. Song, Y. Zhang, X. An, Kriging-assisted discrete global optimization (KDGO) for black-box problems with costly objective and constraints, *Applied Soft Computing* 94 (2020) 106429.
 - [43] Q. Gu, Q. Wang, X. Li, X. Li, A surrogate-assisted multi-objective particle swarm optimization of expensive constrained combinatorial optimization problems, *Knowledge-based Systems* 223 (2021) 107049.
 - [44] Q. Gu, Q. Wang, N. N. Xiong, S. Jiang, L. Chen, Surrogate-assisted evolutionary algorithm for expensive constrained multi-objective discrete optimization problems, *Complex & Intelligent Systems* 8 (2021) 2699–2718.

- [45] R. Sun, Q. Duan, X. Mao, A multi-objective adaptive surrogate modelling-based optimization algorithm for constrained hybrid problems, *Environmental Modelling & Software* 148 (2022) 105272.
- [46] R. G. Regis, A two-phase surrogate approach for high-dimensional constrained discrete multi-objective optimization, in: *The Genetic and Evolutionary Computation Conference, GECCO Companion Volume*, ACM, 2021, pp. 1870–1878.
- [47] V. Drouet, S. Verel, J.-M. Do, Surrogate-assisted asynchronous multiobjective algorithm for nuclear power plant operations, in: *The Genetic and Evolutionary Computation Conference, GECCO*, ACM, 2020, pp. 1073–1081.
- [48] J. Hakanen, J. Malmberg, V. Ojalehto, K. Eyvindson, Data-driven interactive multiobjective optimization using a cluster-based surrogate in a discrete decision space, in: *International Conference on Machine Learning, Optimization, and Data Science, Lecture Notes in Computer Science*, Springer, 2018, pp. 104–115.
- [49] R. G. Regis, High-dimensional constrained discrete multi-objective optimization using surrogates, in: *International Conference on Machine Learning, Optimization, and Data Science*, Springer, 2020, pp. 203–214.
- [50] H. Wang, Y. Jin, A random forest-assisted evolutionary algorithm for data-driven constrained multiobjective combinatorial optimization of trauma systems, *IEEE Transactions on Cybernetics* 50 (2) (2020) 536–549.
- [51] Y. Tian, S. Yang, L. Zhang, F. Duan, X. Zhang, A surrogate-assisted multiobjective evolutionary algorithm for large-scale task-oriented pattern mining, *IEEE Transactions on Emerging Topics in Computational Intelligence* 3 (2) (2019) 106–116.
- [52] F. Leprêtre, C. Fonlupt, S. Verel, V. Marion, Combinatorial surrogate-assisted optimization for bus stops spacing problem, in: *International Conference on Artificial Evolution (Evolution Artificielle)*, Springer, 2020, pp. 42–52.
- [53] M. Rieser, Adding transit to an agent-based transportation simulation: Concepts and implementation, Ph.D. thesis, der Technischen Universität Berlin (06 2010). [doi:10.14279/depositonce-2581](https://doi.org/10.14279/depositonce-2581).
- [54] J. P. Romero, A. Ibeas, J. L. Moura, J. Benavente, B. Alonso, A simulation-optimization approach to design efficient systems of bike-sharing, *Procedia - Social and Behavioral Sciences* 54 (2012) 646–655, *Proceedings of EWGT2012 - 15th Meeting of the EURO Working Group on Transportation*, Paris.
- [55] D. M. Negoescu, P. I. Frazier, W. B. Powell, The knowledge-gradient algorithm for sequencing experiments in drug discovery, *INFORMS Journal on Computing* 23 (3) (2011) 346–363.
- [56] S. Vérel, B. Derbel, A. Liefoghe, H. E. Aguirre, K. Tanaka, A surrogate model based on walsh decomposition for pseudo-boolean functions, in: *Parallel Problem Solving from Nature - PPSN XV*, Vol. 11102 of *Lecture Notes in Computer Science*, Springer, 2018, pp. 181–193.
- [57] F. Leprêtre, S. Verel, C. Fonlupt, V. Marion, Walsh functions as surrogate model for pseudo-boolean optimization problems, in: *The Genetic and Evolutionary Computation Conference, GECCO*, ACM, 2019, pp. 303–311.
- [58] A. Dushatskiy, A. M. Mendrik, T. Alderliesten, P. A. Bosman, Convolutional neural network surrogate-assisted GOMEA, in: *The Genetic and Evolutionary Computation Conference, GECCO*, ACM, 2019, pp. 753–761.
- [59] A. Deshwal, S. Belakaria, J. R. Doppa, Bayesian optimization over hybrid spaces, in: *International Conference on Machine Learning*, PMLR, 2021, pp. 2632–2643.
- [60] K. Swingler, Learning and searching pseudo-boolean surrogate functions from small samples, *Evolutionary Computation* 28 (2) (2020) 317–338.
- [61] L. Han, H. Wang, A random forest assisted evolutionary algorithm using competitive neighborhood search for expensive constrained combinatorial optimization, *Memetic Computing* 13 (2021) 19–30.
- [62] F. Leprêtre, C. Fonlupt, S. Verel, V. Marion, Combinatorial surrogate-assisted optimization for bus stops spacing problem, in: *International Conference on Artificial Evolution (Evolution Artificielle)*, Springer, 2019, pp. 42–52.
- [63] A. Dushatskiy, T. Alderliesten, P. A. Bosman, A novel approach to designing surrogate-assisted genetic algorithms by combining efficient learning of Walsh coefficients and dependencies, *ACM Transactions on Evolutionary Learning and Optimization* 1 (2) (2021) 1–23.
- [64] I. Unanue, M. Merino, J. A. Lozano, A general framework based on Walsh decomposition for combinatorial

- optimization problems, in: Congress on Evolutionary Computation (CEC), IEEE, 2021, pp. 391–398.
- [65] A. Deshwal, S. Belakaria, J. R. Doppa, A. Fern, Optimizing discrete spaces via expensive evaluations: A learning to search framework, in: The AAAI Conference on Artificial Intelligence, AAAI Press, 2020, pp. 3773–3780.
- [66] F. Hutter, H. H. Hoos, K. Leyton-Brown, Sequential model-based optimization for general algorithm configuration, in: Learning and Intelligent Optimization, Springer, 2011, pp. 507–523.
- [67] J. L. Walsh, A closed set of normal orthogonal functions, *American Journal of Mathematics* 45 (1) (1923) 5.
- [68] G. Pruvost, B. Derbel, A. Liefoghe, S. Verel, Q. Zhang, Surrogate-assisted multi-objective combinatorial optimization based on decomposition and Walsh basis, in: The Genetic and Evolutionary Computation Conference, GECCO, ACM, 2020, pp. 542–550.
- [69] H. Li, Q. Zhang, Multiobjective optimization problems with complicated Pareto sets, MOEA/D and NSGA-II, *IEEE Transactions on Evolutionary Computation* 13 (2) (2009) 284–302.
- [70] A. D. Bethke, Genetic algorithms as function optimizers, Ph.D. thesis, University of Michigan (1980).
- [71] T. Hastie, R. Tibshirani, M. Wainwright, *Statistical Learning with Sparsity: The Lasso and Generalizations*, 1st Edition, Chapman and Hall/CRC, 2015.
- [72] R. Tibshirani, Regression shrinkage and selection via the lasso, *Journal of the Royal Statistical Society: Series B (Methodological)* 58 (1) (1996) 267–288.
- [73] L. Paquete, T. Stützle, A study of stochastic local search algorithms for the biobjective QAP with correlated flow matrices, *European Journal of Operational Research* 169 (3) (2006) 943–959.
- [74] H. Aguirre, K. Tanaka, Working principles, behavior, and performance of MOEAs on MNK-landscapes, *European Journal of Operational Research* 181 (3) (2007) 1670–1690.
- [75] S. Verel, A. Liefoghe, L. Jourdan, C. Dhaenens, On the structure of multiobjective combinatorial search space: MNK-landscapes with correlated objectives, *European Journal of Operational Research* 227 (2) (2013) 331–342.
- [76] A. Liefoghe, S. Verel, J.-K. Hao, A hybrid metaheuristic for multiobjective unconstrained binary quadratic programming, *Applied Soft Computing* 16 (2014) 10–19.
- [77] G. Kochenberger, J.-K. Hao, F. Glover, M. Lewis, Z. Lü, H. Wang, Y. Wang, The unconstrained binary quadratic programming problem: A survey, *Journal of Combinatorial Optimization* 28 (1) (2014) 58–81.
- [78] S. A. Kauffman, *The Origins of Order*, Oxford University Press, 1993.
- [79] E. Zitzler, L. Thiele, M. Laumanns, C. M. Fonseca, V. Grunert da Fonseca, Performance assessment of multi-objective optimizers: An analysis and review, *IEEE Transactions on Evolutionary Computation* 7 (2) (2003) 117–132.
- [80] N. Hansen, A. Auger, S. Finck, R. Ros, Real-Parameter Black-Box Optimization Benchmarking 2009: Experimental Setup, Research Report RR-6828, INRIA (2009).