

**Parallel surrogate-assisted optimization  
using Mesh Adaptive Direct Search**

B. Talgorn, S. Alarie,  
M. Kokkolaras

G-2020-38

July 2020

Revised: December 2020

---

La collection *Les Cahiers du GERAD* est constituée des travaux de recherche menés par nos membres. La plupart de ces documents de travail a été soumis à des revues avec comité de révision. Lorsqu'un document est accepté et publié, le pdf original est retiré si c'est nécessaire et un lien vers l'article publié est ajouté.

**Citation suggérée :** B. Talgorn, S. Alarie, M. Kokkolaras (Juillet 2020). Parallel surrogate-assisted optimization using Mesh Adaptive Direct Search, Rapport technique, Les Cahiers du GERAD G-2020-38, GERAD, HEC Montréal, Canada. Révision: décembre 2020.

**Avant de citer ce rapport technique**, veuillez visiter notre site Web (<https://www.gerad.ca/fr/papers/G-2020-38>) afin de mettre à jour vos données de référence, s'il a été publié dans une revue scientifique.

---

La publication de ces rapports de recherche est rendue possible grâce au soutien de HEC Montréal, Polytechnique Montréal, Université McGill, Université du Québec à Montréal, ainsi que du Fonds de recherche du Québec – Nature et technologies.

Dépôt légal – Bibliothèque et Archives nationales du Québec, 2020  
– Bibliothèque et Archives Canada, 2020

The series *Les Cahiers du GERAD* consists of working papers carried out by our members. Most of these pre-prints have been submitted to peer-reviewed journals. When accepted and published, if necessary, the original pdf is removed and a link to the published article is added.

**Suggested citation:** B. Talgorn, S. Alarie, M. Kokkolaras (July 2020). Parallel surrogate-assisted optimization using Mesh Adaptive Direct Search, Technical report, Les Cahiers du GERAD G-2020-38, GERAD, HEC Montréal, Canada. Revised version: December 2020.

**Before citing this technical report**, please visit our website (<https://www.gerad.ca/en/papers/G-2020-38>) to update your reference data, if it has been published in a scientific journal.

---

The publication of these research reports is made possible thanks to the support of HEC Montréal, Polytechnique Montréal, McGill University, Université du Québec à Montréal, as well as the Fonds de recherche du Québec – Nature et technologies.

Legal deposit – Bibliothèque et Archives nationales du Québec, 2020  
– Library and Archives Canada, 2020



# Parallel surrogate-assisted optimization using Mesh Adaptive Direct Search

Bastien Talgorn <sup>a,b</sup>

Stéphane Alarie <sup>c,b</sup>

Michael Kokkolaras <sup>d,b</sup>

<sup>a</sup> Intel Corporation, Programmable Solutions Group, Toronto (Ontario), Canada

<sup>b</sup> GERAD, Montréal (Québec), Canada

<sup>c</sup> Institut de recherche d'Hydro-Québec, Varennes (Québec) Canada

<sup>d</sup> Département de génie mécanique, Université McGill, Montréal (Québec), Canada

alarie.stephane@hydroquebec.com

michael.kokkolarase@mcgill.ca

July 2020

Revised: December 2020

Les Cahiers du GERAD

G–2020–38

Copyright © 2020 GERAD, Talgorn, Alarie, Kokkolaras

Les textes publiés dans la série des rapports de recherche *Les Cahiers du GERAD* n'engagent que la responsabilité de leurs auteurs. Les auteurs conservent leur droit d'auteur et leurs droits moraux sur leurs publications et les utilisateurs s'engagent à reconnaître et respecter les exigences légales associées à ces droits. Ainsi, les utilisateurs:

- Peuvent télécharger et imprimer une copie de toute publication du portail public aux fins d'étude ou de recherche privée;
- Ne peuvent pas distribuer le matériel ou l'utiliser pour une activité à but lucratif ou pour un gain commercial;
- Peuvent distribuer gratuitement l'URL identifiant la publication.

Si vous pensez que ce document enfreint le droit d'auteur, contactez-nous en fournissant des détails. Nous supprimerons immédiatement l'accès au travail et enquêterons sur votre demande.

The authors are exclusively responsible for the content of their research papers published in the series *Les Cahiers du GERAD*. Copyright and moral rights for the publications are retained by the authors and the users must commit themselves to recognize and abide the legal requirements associated with these rights. Thus, users:

- June download and print one copy of any publication from the public portal for the purpose of private study or research;
- June not further distribute the material or use it for any profit-making activity or commercial gain;
- June freely distribute the URL identifying the publication.

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

**Abstract:** We consider computationally expensive blackbox optimization problems and present a method that employs surrogate models and concurrent computing at the search step of the mesh adaptive direct search (MADS) algorithm. Specifically, we solve a surrogate optimization problem using locally weighted scatterplot smoothing (LOWESS) models to find promising candidate points to be evaluated by the blackboxes. We consider several methods for selecting promising points from a large number of points. We conduct numerical experiments to assess the performance of the modified MADS algorithm with respect to available CPU resources by means of five engineering design problems.

# 1 Introduction

We consider the optimization problem

$$\begin{aligned} & \min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}) \\ & \text{subject to } c_j(\mathbf{x}) \leq 0, \quad j = 1, 2, \dots, m, \end{aligned} \tag{P}$$

where  $f(\mathbf{x})$  is the objective function,  $\mathbf{x} \in \mathbb{R}^n$  is the vector of decision variables,  $\mathcal{X}$  is a subset of  $\mathbb{R}^n$ , and  $c_j(\mathbf{x})$  are general nonlinear constraints. We assume that some (at least one) of the functions  $\{f, c_1, c_2, \dots, c_m\}$  are evaluated using simulations or other computational procedures that are blackboxes. In particular, we consider the case where these blackboxes are computationally expensive, possibly nonsmooth and/or nonconvex, and that the process used to evaluate them may crash or fail to return a value. Finally, we assume that function gradients either do not exist theoretically or, if they do, cannot be computed or approximated with reasonable computational effort.

Metaheuristics and derivative-free search algorithms are commonly used for solving (P). The former (e.g., genetic algorithms (GAs), particle swarm optimization (PSO), tabu search (TS), etc.) are commonly used for global exploration while the latter (e.g., generalized pattern search (GPS), mesh adaptive direct search (MADS), and trust-region methods (DFO, COBYLA, CONDOR)) are local methods with convergence properties [22]. In this work, we use the NOMAD implementation [8, 23] of the MADS algorithm [5] to solve (P).

Multiprocessor computers, supercomputers, cloud computing, or just a few connected PCs can provide parallel (or concurrent) computing opportunities to speed up so-called trajectory-based optimization algorithms. According to [1], three ways are commonly used to achieve this: (i) parallel evaluation of neighborhood solutions (distributed evaluations), (ii) parallel trajectories from the same (or different) initial guess(es) (independent optimization runs), (iii) the evaluation of a point  $\mathbf{x}$  is performed in parallel (i.e., the search is sequential). The implementation of (iii) depends only on the blackbox, while the other two are related to the optimization algorithm.

NOMAD offers implementations for (i) and (ii) through p-MADS for (i) and Coop-MADS for (ii) [24]. In both cases, the parallel evaluations are handled by NOMAD by means of MPI calls to the blackbox [23]. However, if one prefers to take charge of the distribution of evaluations, one can implement p-MADS with blocks by using NOMAD in batch mode [24]. In that case, instead of using MPI calls, NOMAD writes all points to be evaluated in an input file, waits for all evaluations to be completed, and reads the obtained values of  $f(\mathbf{x})$  and  $c_j(\mathbf{x})$  from an output file. Then, NOMAD either stops if a local optimum is reached or submits a new block of points to evaluate. We use here the block functionality of NOMAD, adopting option (i) for parallel evaluations, because of its flexibility and generality.

The MADS algorithm includes two steps at each iteration, the SEARCH and the POLL. The SEARCH step is flexible (defined by the user) and aims at determining one or more new points  $\mathbf{x} \in \mathcal{X}$  that improves the current best solution. The POLL step is defined according to the convergence analysis of the algorithm and generates trial points around the current best solution. The number of poll points at each iteration is either  $2n$  or  $n + 1$  depending on the utilized pattern with  $n$  being the size of  $\mathbf{x}$ .

The number of points evaluated in the SEARCH step depends on the methods chosen or defined by the user. Several techniques are already implemented and available in NOMAD, including the speculative search (SS), the quadratic model search (QUAD), the variable neighborhood search (VNS), the Latin hypercube search (LHS), and the Nelder Mead search (NMS). One can also implement their own technique if one so desires, which is called a *user search* (US).

With the exception of LHS, all provided techniques usually return only one trial point. When several techniques are used at once, they are called one after the other along the SEARCH step, each technique providing its own trial point, which is evaluated by the blackbox before proceeding to the next technique. Assuming that  $2n$  CPUs are available for solving (P), the POLL step can make good

use of these CPUs. However, since SEARCH step evaluations are sequential, progress is slowed down with almost all CPUs being idle. One may argue that we should then only use LHS since it can generate  $2n$  points. However, since LHS is random, its points will quickly become less promising after a few iterations.

Considering that the number of available CPUs are now, particularly with the emergence of cloud computing, relatively inexpensive and unlimited, we should rethink the SEARCH step to be as effective as the POLL step in terms of CPU use. In this work, we propose a SEARCH step technique that returns a large number of diverse points for evaluation.

The paper is structured as follows. The general idea behind the proposed technique is described in Section 2. In Section 3, six different methods are presented for selecting various candidates from a large set of points. In Section 4, the resulting model search for parallel computing is specified. In Section 5, we test our SEARCH step technique on five engineering design optimization problems using up to 64 processors. A discussion concludes the paper.

## 2 Proposed SEACH step technique

One of the practical challenges of the SEARCH step is that only one candidate is obtained at a significant computational investment [7, 27, 31, 32]. Specifically, regardless of the number of available CPUs, only one CPU is used in the SEARCH step for blackbox evaluations, with the exception of LHS. Before presenting our idea for mitigating this practical challenge, we will assume that computationally inexpensive surrogate models of the expensive blackboxes are available. We can then consider the surrogate problem of problem ( $P$ )

$$\begin{aligned} \min_{\mathbf{x} \in \mathcal{X}} \quad & \hat{f}(\mathbf{x}) \\ \text{subject to} \quad & \hat{c}_j(\mathbf{x}) \leq 0, \quad j = 1, 2, \dots, m, \end{aligned} \tag{\hat{P}}$$

where  $\{\hat{f}, \hat{c}_1, \hat{c}_2, \dots, \hat{c}_m\}$  are surrogate models of  $\{f, c_1, c_2, \dots, c_m\}$ , respectively. Note that we only need to ensure that the minimizers of ( $P$ ) and ( $\hat{P}$ ) are close enough, and not that the surrogate models are good approximations of the blackboxes globally. It then follows that a minimizer of ( $\hat{P}$ ) will be a good candidate for the solution of ( $P$ ).

If both problems have the same minimizers, they may share features in other areas of  $\mathcal{X}$  as well. Since the evaluations of  $\hat{f}(\mathbf{x})$  and  $\hat{c}_j(\mathbf{x})$  are rather inexpensive compared to  $f$  and  $c_j$ , one can allow a very large budget of model evaluations to solve ( $\hat{P}$ ), extending thus the number of design space areas that will be visited. This is acceptable as long as the solution of ( $\hat{P}$ ) is faster than any single evaluation of the blackboxes. Considering there are  $q$  CPUs available for blackbox evaluations, one may then select  $q$  points from the available budget by solving ( $\hat{P}$ ). The  $q$  points can be selected to consider areas of  $\mathcal{X}$  that have been neglected until now in the solution of ( $P$ ).

The above proposition proposes the use of surrogate models  $\{\hat{f}, \hat{c}_1, \hat{c}_2, \dots, \hat{c}_m\}$  in a manner that is not reported in [18], which mentions two ways of exploiting surrogates in the context of parallelization. The simplest is to fit  $q$  different surrogate models at the same points already evaluated by the blackbox functions. This allows to get  $q$  different promising candidates and requires no uncertainty quantification for the surrogate models. One can also combine the surrogates; distance-based criteria can be added to ensure diversity between the candidates. The other way is to use a single surrogate model and consider  $q$  points where the blackboxes should be evaluated at to improve its accuracy.

We propose an intermediate approach. We use only one surrogate model for each blackbox. The  $q$  candidates are extracted from that single surrogate, but not with the aim of improving it. Instead, the  $q$  candidates are selected to be the most interesting to advance the optimization process.

### 3 Methods for selecting candidate points

Let  $\mathbf{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k\}$  be the set of all points evaluated by the blackbox. Note that  $\mathbf{X} \subset \mathcal{X} \subset \mathbb{R}^n$ . We denote  $\hat{\mathbf{X}}$  the *surrogate cache*, i.e., the set of all points for which  $\{\hat{f}, \hat{c}_1, \hat{c}_2, \dots, \hat{c}_m\}$  have been evaluated during the solution of  $(\hat{P})$ . Similarly,  $\hat{\mathbf{X}} \subset \mathcal{X} \subset \mathbb{R}^n$ . Let  $\mathbf{S}$  be the set of points that are selected by the SEARCH step to be evaluated with the blackbox. The set  $\mathbf{S}$  is initially empty and is built from the points of  $\hat{\mathbf{X}}$  (ensuring that  $\mathbf{S} \subset \hat{\mathbf{X}}$ ) with a greedy algorithm by means of up to six selection methods, each one having a different goal.

- Method 1 selects the best point of  $\hat{\mathbf{X}}$  not in  $\mathbf{X} \cup \mathbf{S}$ ;
- Method 2 selects the most distant point of  $\hat{\mathbf{X}}$  from  $\mathbf{X} \cup \mathbf{S}$ ;
- Method 3 selects the best point of  $\hat{\mathbf{X}}$  at a certain distance of  $\mathbf{X} \cup \mathbf{S}$ ;
- Method 4 selects the best point of  $\hat{\mathbf{X}}$  under additional constraints;
- Method 5 selects a point of  $\hat{\mathbf{X}}$  that is a possible local minimum of the surrogate problem;
- Method 6 selects a point of  $\hat{\mathbf{X}}$  in a non-explored area.

Note that some selection methods may fail to return a candidate, particularly methods 3 and 4. If this happens, the next method is used. We repeat and loop through all methods until we obtain enough points in  $\mathbf{S}$  matching the available CPUs. Some points  $\mathbf{x} \in \hat{\mathbf{X}}$  can also belong to the set  $\mathbf{X}$ ; this is not an issue since all methods only select point  $\mathbf{x} \in \hat{\mathbf{X}}$  to be added to  $\mathbf{S}$  if and only if  $\mathbf{x} \notin \mathbf{X}$ . The selection methods are detailed below after some definitions.

#### 3.1 Definitions

Let  $d(A, B)$  be the Euclidean distance between two subsets  $A$  and  $B$  of  $\mathbb{R}^n$

$$d(A, B) = \min_{a \in A} \min_{b \in B} \|a - b\|_2. \quad (1)$$

As a convention, the distance to an empty set is infinite:  $d(A, \emptyset) = d(\emptyset, \emptyset) = +\infty$ . By extension, we will denote the distance between an element  $a \notin B$  and the subset  $B$  simply by  $d(a, B)$ , which implies that  $a$  also refers to the particular subset containing only  $a$ , i.e.,  $\{a\}$ .

Regarding feasibility, we consider the aggregate constraint violation function used in [15], i.e.,  $h(\mathbf{x}) = \sum_{j=1}^m \max\{0, c_j(\mathbf{x})\}^2$ . The same function is used in the *progressive barrier* mechanism in NOMAD [6].

We also define the order operators between two points  $\mathbf{x}$  and  $\mathbf{x}' \in \mathcal{X}$ :

$$\mathbf{x} \prec \mathbf{x}' \Leftrightarrow \begin{cases} h(\mathbf{x}) < h(\mathbf{x}') \\ \text{or} \\ h(\mathbf{x}) = h(\mathbf{x}') \text{ and } f(\mathbf{x}) < f(\mathbf{x}') \end{cases} \quad (2)$$

$$\mathbf{x} \preceq \mathbf{x}' \Leftrightarrow \text{not}(\mathbf{x}' \prec \mathbf{x}), \quad (3)$$

which are transitive. By those definitions, an incumbent solution  $\mathbf{x}'$  of the original problem  $(P)$  is such that  $\mathbf{x}' \preceq \mathbf{x}$ ,  $\forall \mathbf{x} \in \mathbf{X}$ . Similarly, a global minimizer  $\mathbf{x}^*$  is such that  $\mathbf{x}^* \preceq \mathbf{x}$ ,  $\forall \mathbf{x} \in \mathcal{X}$ . In the same manner as we define  $\prec$  and  $\preceq$  for  $f$  and  $h$ , we define  $\succsim$  and  $\succeq$  for  $\hat{f}$  and  $\hat{h}$ .

Finally, to simplify the description of the proposed selection methods, we define  $\mathbf{s}^\infty$  as a *virtual point* (in the sense that it does not have coordinates in  $\mathbb{R}^n$ ), which represents the worst possible candidate in  $\mathbb{R}^n$ :

$$\hat{h}(\mathbf{s}^\infty) = \hat{f}(\mathbf{s}^\infty) = +\infty \quad \text{and} \quad d(\mathbf{s}^\infty, \mathbf{X}) = 0. \quad (4)$$

#### 3.2 Detailed description of selection methods

**Method 1.** The first selection method selects the best point  $\mathbf{s}$  of  $\hat{\mathbf{X}}$  under the constraint that  $d(\mathbf{s}, \mathbf{X} \cup \mathbf{S}) > 0$ , which means that  $\mathbf{s}$  is not in the set  $\mathbf{X}$  of evaluated points nor already selected (i.e.,  $\notin \mathbf{S}$ ). This method reflects how surrogate models are typically used for finding new candidate points.

**Algorithm 1** Selection of the best point (Method 1)

---

```

s* ← s∞
for all s ∈  $\hat{\mathbf{X}}$ , do:
  | if s  $\succ$  s* and d(s,  $\mathbf{X} \cup \mathbf{S}$ ) > 0, then:
  | | s* ← s
  | end
end
if s* ≠ s∞, then:
  |  $\mathbf{S} \leftarrow \mathbf{S} \cup \{\mathbf{s}^*\}$ 
end

```

---

**Method 2.** The second method aims to maximize the diversity of the candidates to be evaluated. It selects the point  $\mathbf{s}$  of  $\hat{\mathbf{X}}$  that maximizes the distance  $d(\mathbf{s}, \mathbf{X} \cup \mathbf{S})$ , i.e., as far as possible from points already evaluated.

**Algorithm 2** Selection of the most distant point to  $\mathbf{X} \cup \mathbf{S}$  (Method 2)

---

```

s* ← s∞
for all s ∈  $\hat{\mathbf{X}}$ , do:
  | if d(s,  $\mathbf{X} \cup \mathbf{S}$ ) > d(s*,  $\mathbf{X} \cup \mathbf{S}$ ), then:
  | | s* ← s
  | end
end
if s* ≠ s∞, then:
  |  $\mathbf{S} \leftarrow \mathbf{S} \cup \{\mathbf{s}^*\}$ 
end

```

---

**Method 3.** This method selects the best point  $\mathbf{s}$  of  $\hat{\mathbf{X}}$  under the constraint that  $d(\mathbf{s}, \mathbf{X} \cup \mathbf{S}) \geq d_{\min}$ , where  $d_{\min}$  is initialized at 0 at the beginning of the selection process and increased progressively as the method is applied. Method 3 may fail to select a candidate  $\mathbf{s}$  when  $d_{\min}$  becomes too large. Since the selected points  $\mathbf{S}$  must be projected on the current mesh  $\mathcal{M} = \{\mathbf{x} + \Delta^m \mathbf{D}\mathbf{z}, \mathbf{z} \in \mathbb{N}^{n_D}, \mathbf{x} \in \mathbf{X}\}$  as required by MADS, incrementing  $d_{\min}$  by the current mesh size  $\Delta^{\mathcal{M}}$  allows to avoid that several candidates become identical after the projection.

**Algorithm 3** Selection of the best point with a constraint on the distance to  $\mathbf{X} \cup \mathbf{S}$  (Method 3)

---

```

s* ← s∞
if first use of Method 3, then:
  | dmin ← 0
end
for all s ∈  $\hat{\mathbf{X}}$ , do:
  | if s  $\succ$  s* and d(s,  $\mathbf{X} \cup \mathbf{S}$ ) ≥ dmin, then:
  | | s* ← s
  | end
end
if s* ≠ s∞, then:
  |  $\mathbf{S} \leftarrow \mathbf{S} \cup \{\mathbf{s}^*\}$ 
  | dmin ← dmin +  $\Delta^{\mathcal{M}}$ 
end

```

---

**Method 4.** Considering that the surrogate models  $\hat{c}_j$  may fail to predict correctly if  $\mathbf{s}$  is feasible, the present method tries to select points that will be likely to be feasible when evaluated by the

blackboxes  $c_j$ . This is done by selecting the best feasible point  $\mathbf{s}$  of  $\hat{\mathbf{X}}$  under the constraint  $\hat{c}_{\max}(\mathbf{s}) \leq \hat{c}_{\text{margin}}$ , where  $\hat{c}_{\max}(\mathbf{s})$  is defined as being the most violated constraint of  $\mathbf{s}$ , i.e.,

$$\hat{c}_{\max}(\mathbf{s}) = \max_{j=1,2,\dots,m} \hat{c}_j(\mathbf{s}), \quad (5)$$

where  $\hat{c}_{\text{margin}}$  is set as

$$\hat{c}_{\text{margin}} \leftarrow \max_{\substack{\mathbf{s} \in \hat{\mathbf{X}} \\ \hat{c}_{\max}(\mathbf{s}) < 0}} \hat{c}_{\max}(\mathbf{s}) \quad (6)$$

and quantifies, among all feasible points of  $\hat{\mathbf{X}}$ , the smallest amount by which these are satisfied.

By definition,  $\hat{c}_{\max}(\mathbf{s}) \leq 0$  if  $\mathbf{s}$  is predicted to be feasible by the surrogate models. The more negative  $\hat{c}_{\max}(\mathbf{s})$  is, the more likely is  $\mathbf{s}$  to be feasible when evaluated by the blackboxes. Decreasing progressively the value of  $\hat{c}_{\text{margin}}$  after each call of this selection method will favor candidates that are increasingly likely to be feasible (but possibly with worse objective function values).

---

**Algorithm 4** Selection of the best point with a constraint on the feasibility (Method 4)

---

```

 $\mathbf{s}^* \leftarrow \mathbf{s}^\infty$ 
if first use of Method 4, then:
     $\hat{c}_{\text{margin}} \leftarrow \min\{0, \max_{\substack{\mathbf{s} \in \hat{\mathbf{X}} \\ \hat{c}_{\max}(\mathbf{s}) < 0}} \hat{c}_{\max}(\mathbf{s})\}$ 
end
for all  $\mathbf{s} \in \hat{\mathbf{X}}$ , do:
    if  $\hat{c}_{\max}(\mathbf{s}) \leq \hat{c}_{\text{margin}}$  and  $\hat{f}(\mathbf{s}) < \hat{f}(\mathbf{s}^*)$  and  $d(\mathbf{s}, \mathbf{X} \cup \mathbf{S}) > \Delta^{\mathcal{M}}$ , then:
         $\mathbf{s}^* \leftarrow \mathbf{s}$ 
    end
end
if  $\mathbf{s}^* \neq \mathbf{s}^\infty$ , then:
     $\mathbf{S} \leftarrow \mathbf{S} \cup \{\mathbf{s}^*\}$ 
     $\hat{c}_{\text{margin}} \leftarrow 2\hat{c}_{\max}(\mathbf{s}^*)$ 
end

```

---

Note that this method requires that  $\hat{c}_{\text{margin}} \leq 0$  is always satisfied. Moreover, we also assume that there is at least one  $\mathbf{s} \in \hat{\mathbf{X}}$  that is feasible. If it is not the case, which may happen in the first iteration, we will end up with  $\hat{c}_{\text{margin}} > 0$  and an inappropriate candidate will be selected. To avoid this, we initialize  $\hat{c}_{\text{margin}}$  to 0 so that the method may fail to return a candidate if that is the case.

**Method 5.** The isolation distance is used here to detect local minima of the surrogate problem. This concept is inspired from the *topographic isolation* of a mountain summit, which measures the local significance of a summit. It is defined as the distance to the closest higher summit.<sup>1</sup>

Transferred to optimization, the concept of topographic isolation is used to quantify the importance of a local minimum. Its strict application is however impossible since it will require to prove that no other point within a certain distance of  $\mathbf{x}$  is better than  $\mathbf{x}$ . We can only compute isolation distance of the already evaluated points. Consequently, we define the isolation distance as being the distance from  $\mathbf{s}$  to the closest point of  $\hat{\mathbf{X}}$  that is better than  $\mathbf{s}$

$$d_{\text{iso}}(\mathbf{s}) = \min_{\substack{\mathbf{s}' \in \hat{\mathbf{X}} \\ \mathbf{s}' \succ \mathbf{s}}} d(\mathbf{s}, \mathbf{s}'). \quad (7)$$

---

<sup>1</sup> In mountaineering, the topographic isolation of Mount Everest is infinite and the summit with the second highest isolation is the Aconcagua in Argentina. The Aconcagua is not the second highest summit but there is no higher mountain in a 16,518 km range, making it the most important summit in the Americas and in the southern hemisphere.

Constraints are taken into account by using the order relationship defined in Equation (2). As a convention, if no point of  $\hat{\mathbf{X}}$  is better than  $\mathbf{s}$ , then  $d_{\text{iso}}(\mathbf{s}) = +\infty$ . With this definition, the point of  $\hat{\mathbf{X}}$  with the highest isolation distance is also the best candidate in  $\hat{\mathbf{X}}$ . However, we have observed that the other points with a high isolation distance are often poor points far from any other point of  $\hat{\mathbf{X}}$ . To address this problem, we define the *isolation number* of  $\mathbf{s} \in \hat{\mathbf{X}}$  as the number of points of  $\hat{\mathbf{X}}$  within the ball of centre  $\mathbf{s}$  and radius  $d_{\text{iso}}(\mathbf{s})$

$$n_{\text{iso}}(\mathbf{s}) = \text{card}\{\mathbf{s}' : \mathbf{s}' \in \hat{\mathbf{X}}, d(\mathbf{s}, \mathbf{s}') < d_{\text{iso}}(\mathbf{s})\}. \quad (8)$$

To have a high isolation number, a point must be better than many of its neighbors, which means that this criterion allows to detect local minima. Note that Equation (4) implies that  $d_{\text{iso}}(\mathbf{s}^\infty) = n_{\text{iso}}(\mathbf{s}^\infty) = 0$ . Method 5 selects the point of  $\hat{\mathbf{X}}$  that has the highest isolation number not yet in  $\mathbf{X} \cup \mathbf{S}$ .

---

**Algorithm 5** Selection of the most isolated point (Method 5)

---

```

 $\mathbf{s}^* \leftarrow \mathbf{s}^\infty$ 
for all  $\mathbf{s} \in \hat{\mathbf{X}}$ , do:
  | if  $n_{\text{iso}}(\mathbf{s}) > n_{\text{iso}}(\mathbf{s}^*)$  and  $d(\mathbf{s}, \mathbf{X} \cup \mathbf{S}) > 0$ , then:
  |   |  $\mathbf{s}^* \leftarrow \mathbf{s}$ 
  |   end
end
if  $\mathbf{s}^* \neq \mathbf{s}^\infty$ , then:
  |  $\mathbf{S} \leftarrow \mathbf{S} \cup \{\mathbf{s}^*\}$ 
end

```

---

**Method 6.** The purpose of this method is to select points in neglected areas of the design space. To do so, it selects points in areas heavily explored while solving ( $\hat{P}$ ) but overlooked when solving ( $P$ ). The *density number* of  $\mathbf{s} \in \hat{\mathbf{X}}$  is defined as

$$n_{\text{density}}(\mathbf{s}) = \text{card}\{\mathbf{s}' : \mathbf{s}' \in \hat{\mathbf{X}}, d(\mathbf{s}, \mathbf{s}') < d(\mathbf{s}, \mathbf{X} \cup \mathbf{S})\}. \quad (9)$$

Method 6 selects the point of  $\hat{\mathbf{X}}$  with the highest density number. Note that, as for  $n_{\text{iso}}$ , Equation (4) implies that  $n_{\text{density}}(\mathbf{s}^\infty) = 0$ .

---

**Algorithm 6** Selection of a point in a populated area (Method 6)

---

```

 $\mathbf{s}^* \leftarrow \mathbf{s}^\infty$ 
for all  $\mathbf{s} \in \hat{\mathbf{X}}$ , do:
  | if  $n_{\text{density}}(\mathbf{s}) > n_{\text{density}}(\mathbf{s}^*)$ , then:
  |   |  $\mathbf{s}^* \leftarrow \mathbf{s}$ 
  |   end
end
if  $\mathbf{s}^* \neq \mathbf{s}^\infty$ , then:
  |  $\mathbf{S} \leftarrow \mathbf{S} \cup \{\mathbf{s}^*\}$ 
end

```

---

## 4 Parallel computing implementation

We now describe how we implement the proposed SEARCH step. We start with the surrogate models and follow with the algorithms used for solving ( $\hat{P}$ ). We conclude with how all of this is integrated with the MADS algorithm to solve ( $P$ ).

## 4.1 Surrogate models

Several surrogate models are mentioned in [18] regarding their use in parallel computing approaches, including Kriging, radial basis functions (RBF), support vector regression (SVR), and polynomial response surfaces (PRs). For a more exhaustive description and comparison of surrogate models, see [2]. Based on our previous work reported in [31], we choose to use the locally weighted scatterplot smoothing (LOWESS) surrogate modeling approach [10, 11, 12, 13].

LOWESS models generalize PRs and kernel smoothing (KS) models. PRs are good for small problems, but their efficacy decreases for larger, highly nonlinear, or discrete problems. KS models tend to overestimate low function values and underestimate high ones, but usually predict correctly which of two points yields the smallest function value [7]. LOWESS models build a linear regression of kernel functions around the point to estimate. They have been shown to be suitable for surrogate-based optimization [31]. Their parameters are chosen using an error metric called “aggregate order error with cross-validation” (AOECV), which favors equivalence between the original problem ( $P$ ) and the surrogate problem ( $\hat{P}$ ) [7]. We use the SGTELIB implementation of the surrogate models, which is now integrated as a surrogate library in NOMAD version 3.8 [30]. Specifically, the considered LOWESS model is defined in SGTELIB as follows.

```
TYPE Lowess DEGREE 1 RIDGE 0 SHAPE_COEF OPTIM KERNEL_TYPE OPTIM,
```

which means that the local regression is linear, the ridge (regularization) coefficient is 0, and the kernel shape and kernel type are optimized to minimize the aggregate order error (see [31]).

The LOWESS model is built as described in Appendix A. Only the Gaussian kernel was considered in [31]. Six additional kernel functions have meanwhile been implemented in SGTELIB. Accordingly, not only  $\lambda$  (kernel shape) is chosen to minimize AOECV, but also  $\phi$  (kernel type).

## 4.2 Surrogate problem solution

The surrogate problem ( $\hat{P}$ ) is solved by means of an inner instance of MADS; it is initialized by a Latin hypercube search (LHS) [25, 29] and uses variable neighborhood search (VNS) [19, 26] as the SEARCH step and a large budget of function evaluations (10,000). The POLL step is performed using the ORTHO 2N directions option. This inner MADS is implemented in the SEARCH step of the outer MADS.

The LHS guarantees that there are surrogate cache points widely spread over the design space. To ensure this, 30% of all function evaluations (i.e., 3,000) are devoted to the LHS. Four additional points are considered: The current best feasible point of the original problem, the current best infeasible point of the original problem, the best feasible point obtained by solving the most recent surrogate problem instantiation, and the best infeasible point by solving the most recent surrogate problem instantiation. These points are used as initial guesses of the inner MADS problem, which will be run until the remaining evaluation budget is exhausted. This budget will be shared between the POLL step and the VNS in a default proportion where 75% is devoted to VNS, which favors the exploration of multiple local attraction basins. A large number of evaluations that build the surrogate cache  $\hat{\mathbf{X}}$  favors an accurate solution of the surrogate problem. Using LHS, VNS, and a large number of function evaluations ensures that  $\hat{\mathbf{X}}$  contains highly promising candidates for the solution of the original problem ( $P$ ).

## 4.3 The modified MADS algorithm

Recall that each iteration of MADS includes a SEARCH step (performed first) and a POLL step. Let  $t$  denote the current iteration. Then,  $\mathbf{X}_t$ ,  $\hat{\mathbf{X}}_t$ , and  $\mathbf{S}_t$  denote the sets  $\mathbf{X}$ ,  $\hat{\mathbf{X}}$ , and  $\mathbf{S}$  considered by MADS at iteration  $t$ . Let  $q$  be the number of available CPUs, i.e., the number of blackbox evaluations that can be performed in parallel. The proposed MADS for exploiting  $q$  CPUs is as follows. First, the SEARCH step proceeds by solving the surrogate problem to populate the set  $\hat{\mathbf{X}}_t$ . From that set,  $q$

candidates are selected and returned to be evaluated by the blackbox(es) in parallel. The selection is made by cycling through a user-defined subset of the six proposed selection methods (Section 3.2) until a total of  $q$  candidates are selected, or until all selection methods consecutively failed to add a candidate to  $\mathbf{S}_t$ . If  $q$  is smaller than the number of selection methods retained by the user, we do not necessarily go through all the methods, but stop as soon as we get  $q$  candidates.

If/when the SEARCH step fails to return a better objective function value, the MADS algorithm proceeds to the POLL step. Let  $\mathbf{P}_t$  be the set of candidates produced by the polling directions at the iteration  $t$ . The cardinality of  $\mathbf{P}_t$  is denoted by  $|\mathbf{P}_t|$ . To be consistent with our need to fulfill continuously all the available CPUs with evaluations, additional candidates are added so that  $|\mathbf{P}_t|$  is at least  $q$  or a multiple of  $q$ . This is accomplished by means of NOMAD's intensification mechanism ORTHO 1. If  $|\mathbf{P}_t| = q$ , all poll candidates of  $\mathbf{P}_t$  are evaluated concurrently, eliminating the need to order them. If  $|\mathbf{P}_t| > q$ , then the points in  $\mathbf{P}_t$  are regrouped in several blocks of  $q$  candidates. The blocks are then evaluated sequentially and opportunistically, which means that if a block evaluation leads to a success, the remaining blocks are not evaluated. To increase the probability of success from the first block, and hence avoiding to proceed with the remaining ones, the candidates of  $\mathbf{P}_t$  are sorted using the surrogate models and distributed in the blocks so that the more promising ones are in the first block.

Recall that  $\mathcal{M}_t$  is the current mesh,  $\Delta_t^{\mathcal{M}}$  is the associated mesh size parameter,  $\Delta_t^{\mathbf{P}}$  is the corresponding mesh poll parameter, and  $\mathbf{x}_t^*$  is the best solution found at iteration  $t$ . Finally, the set  $\mathbf{X}_t$  is updated with all the points  $\mathbf{x}$  in  $\mathbf{S}_{t-1}$  and  $\mathbf{P}_{t-1}$  that have been evaluated during the previous iteration. The process is summarized in Algorithm 7.

---

**Algorithm 7** The proposed MADS optimization algorithm
 

---

- [1] **Initialization**  
 $t \leftarrow 0$   
 Set initial poll and mesh sizes  $\Delta_0^{\mathbf{P}} \geq \Delta_0^{\mathcal{M}} > 0$   
 Initialize  $\mathbf{X}_0$  with starting points  
 Evaluate  $\{f(\mathbf{x}), c_1(\mathbf{x}), c_2(\mathbf{x}), \dots, c_m(\mathbf{x})\} \forall \mathbf{x} \in \mathbf{X}_0$
  - [2] **Model search**  
 Use  $\mathbf{X}_t$  to build  $\hat{f}$  and  $\{\hat{c}_j\}_{j \in J}$   
 Solve surrogate problem ( $\hat{P}$ ) using the inner MADS instance  
 $\hat{\mathbf{X}}_t \leftarrow$  Set of points evaluated with surrogate model while solving ( $\hat{P}$ )  
 $\mathbf{S}_t \leftarrow$  Cycle through selection steps to select  $q$  points of  $\hat{\mathbf{X}}_t$   
 $\mathbf{S}_t \leftarrow$  Projection of the points of  $\mathbf{S}_t$  onto mesh  $\mathcal{M}_t$   
 Parallel evaluation of  $\{f(\mathbf{x}), c_1(\mathbf{x}), c_2(\mathbf{x}), \dots, c_m(\mathbf{x})\} \forall \mathbf{x} \in \mathbf{S}_t$   
 If success, goto [4]
  - [3] **Poll**  
 Build poll set  $\mathbf{P}_t$   
 Sort  $\mathbf{P}_t$  according to  $\hat{f}$  and  $\{\hat{c}_j\}_{j \in J}$   
 Parallel evaluation of  $\{f(\mathbf{x}), c_1(\mathbf{x}), c_2(\mathbf{x}), \dots, c_m(\mathbf{x})\} \forall \mathbf{x} \in \mathbf{P}_t$
  - [4] **Updates**  
 $t \leftarrow t + 1$   
 Update  $\Delta_t^{\mathcal{M}}$ ,  $\Delta_t^{\mathbf{P}}$ ,  $\mathbf{x}_t^*$  and  $\mathbf{X}_t$   
 If no stopping condition is met, goto [2]
- 

## 5 Numerical investigation

The proposed SEARCH step technique is tested using five optimization problems. We first describe the algorithms considered for benchmarking. Next, numerical results are presented and discussed for the five engineering design problems.

## 5.1 Compared algorithms

Five solvers are compared in our numerical experiments, all based on the MADS algorithm and implemented using NOMAD 3.8 [23]. This ensures avoiding coding biases since features are identical among solvers.

- **MADS.** Refers to the POLL step of MADS, without any SEARCH step, where  $2n$  directions are generated and evaluated in parallel. If needed,  $k$  additional directions are generated such that  $2n + k$  is a multiple of  $q$ .
- **Multi-Start.** Consists of  $q$  parallel runs of MADS. They are totally independent and each instance runs on its own CPU. Each instance proceeds to its evaluations sequentially, one after the other. Only the POLL step is executed and no cache is shared between running instances.
- **LH Search.** The MADS solver mentioned above using a Latin hypercube search (LHS) at the SEARCH step, where  $q$  candidates are generated and evaluated in parallel.
- **Lowess-A.** The MADS solver mentioned above with the described surrogate optimization conducted at the SEARCH step. The  $q$  candidates are selected by cycling through Methods 1 and 2, and then evaluated in parallel.
- **Lowess-B.** The MADS solver mentioned above with the proposed surrogate optimization conducted at the SEARCH step. The  $q$  search candidates are selected by cycling through Methods 3, 4, 5, and 6, and then evaluated in parallel.

Both LOWESS solvers are exactly like Algorithm 7, excepted for the used selection methods. The only difference between them and the LHS solver is that the surrogate optimization approach is replaced by a LHS at the SEARCH step. This should allow us to determine whether surrogate optimization has any advantage over a random search. The MADS solver is used as the baseline. Finally, the Multi-Start solver is considered to ensure that one should not proceed with  $q$  independent narrow trajectories instead of one single trajectory having  $q$  wide evaluations.

## 5.2 Engineering design optimization problems

The above solvers are compared on five engineering design application problems. A short description follows below for each problem. More details are provided in Appendix B.

- **TCSD.** The Tension/Compression Spring Design problem consists of minimizing the weight of a spring under mechanical constraints [3, 9, 16]. This problem has three variables and four constraints. The design variables define the geometry of the spring. The constraints concern shear stress, surge frequency, and minimum deflection.
- **Vessel.** This problem considers the design of a compressed air storage tank and has four design variables and four constraints [16, 21]. The variables define the geometry of the tank and the constraints are related to the volume, pressure, and solidity of the tank. The objective is to minimize the total cost of the tank, including material and labour.
- **Welded.** The welded beam design problem (Version I) has four variables and six constraints [16, 28]. It aims at minimizing the construction cost of a beam, under shear stress, bending stress, and deflection constraints. The design variables define the geometry and the characteristics of the welded joint.
- **Solar 1.** This optimization problem aims at maximizing the energy received over a period of 24 hours under five constraints related to budget and heliostat field area [17]. It has nine variables, including an integer one without an upper bound.
- **Solar 7.** This problem aims at maximizing the efficiency of the receiver over a period of 24 hours for a given heliostats field under six binary constraints [17]. It has seven design variables, including an integer one without an upper bound.

A progressive barrier is used to deal with the aggregated constraints [6]. The three first problems are easier relative to the last two ones. However, it is difficult to find a feasible solution for the TCSD problem. Among all the considered problems, Solar 1 is certainly the most difficult one.

### 5.3 Numerical experiments

We compare the efficiency of each solver for different values of block size  $q \in \{1, 2, 4, 8, 16, 32, 64\}$ . As an example, we will use “Lowess-A 16” to refer to the solver that relies on LOWESS models cycling over Methods 1 and 2 considering a block size  $q = 16$ . For each problem, we generated 50 sets of 64 starting points with Latin hypercube sampling [25]. For all solvers other than “Multi-Start”, only the first point of each set is used to perform optimizations. Doing so, we get 50 runs from the same starting points for each solver, each problem, and each value of  $q$ . For “Multi-Start”, since  $q$  independent and parallel sequential runs of MADS must be performed, we use the  $q$  first points of each set. Doing so, we still get 50 runs for each problem and each  $q$ , while ensuring that all starting points are the same for all solvers.

To avoid that all “LH Search” runs end up with nearly identical solutions for a given  $q$ , we use a random seed for initializing each LHS.

For the relatively three simpler problems (TCSO, Vessel, and Welded), a budget of 100 block evaluations is allocated. For the two relatively difficult problems (Solar 1 and Solar 7), the budget is increased to 200 block evaluations. This means that, for a given problem, all solvers will have the same “wall-clock” time, but not necessarily the same resources (number of CPUs available for block evaluations) nor the same total number of blackbox evaluations.

#### Solution quality

Figures 1 and 2 represent the distribution of the final objective function over the 50 runs for each problem, each solver, and each block size  $q$ . The minimum and maximum objective values that were obtained from the runs are indicated by circles in the figures. Lower and upper quantiles are delimited by boxes. Median values are represented by a bar into the boxes. The more a distribution is on the left side, the better the combination of solver and  $q$  is. Since we are mostly interested in the best combinations, the figures only focus on the smallest values. Otherwise, it would be difficult to discern the difference among the best combinations. All combinations for which the distribution is cut on the right side are performing poorly.

For the three simpler problems (TCSO, Vessel, and Welded, Figure 1), the LOWESS solvers (and in particular “Lowess-B”) are by far superior to the solvers that do not rely on surrogate optimization.

For the TCSO problem, the “MADS” solver often failed to find a feasible design (thus leading to infinite objective values), even with a large number of evaluations per block. The four other solvers always managed to find a feasible point for at least 75% of the runs. The “Lowess-B” solver performs better than any of the other ones. We see that “Lowess-A 64” is outperformed by “Lowess-B 8”. As the TCSO problem is very constrained, the final objective function value depends on the initial guess. This is why the “Multi-Start” solver performs quite well on this problem.

The same trend is observed for the Vessel, and Welded problems. “Lowess-B” performs better than “Lowess-A”, which outperforms “LHS” or “MADS”. In particular, “Lowess-B 8” outperforms the solvers “Lowess-A 8/16/32”. As expected, increasing the block size improves performance. However, for the “Lowess-B” solver these three problems are easy to solve, so it is difficult to see an advantage of using parallel computing because the global optimum is found most of the time within 100 block evaluations for a block size of 16 or more. The “Multi-Start” solver performs rather poorly on these two problems.

The numerical results generally follow the same trend for the two Solar problems (Figure 2). For a block of equal size, the LOWESS solvers outperform the other solvers while “Lowess-B” outperforms “Lowess-A”.

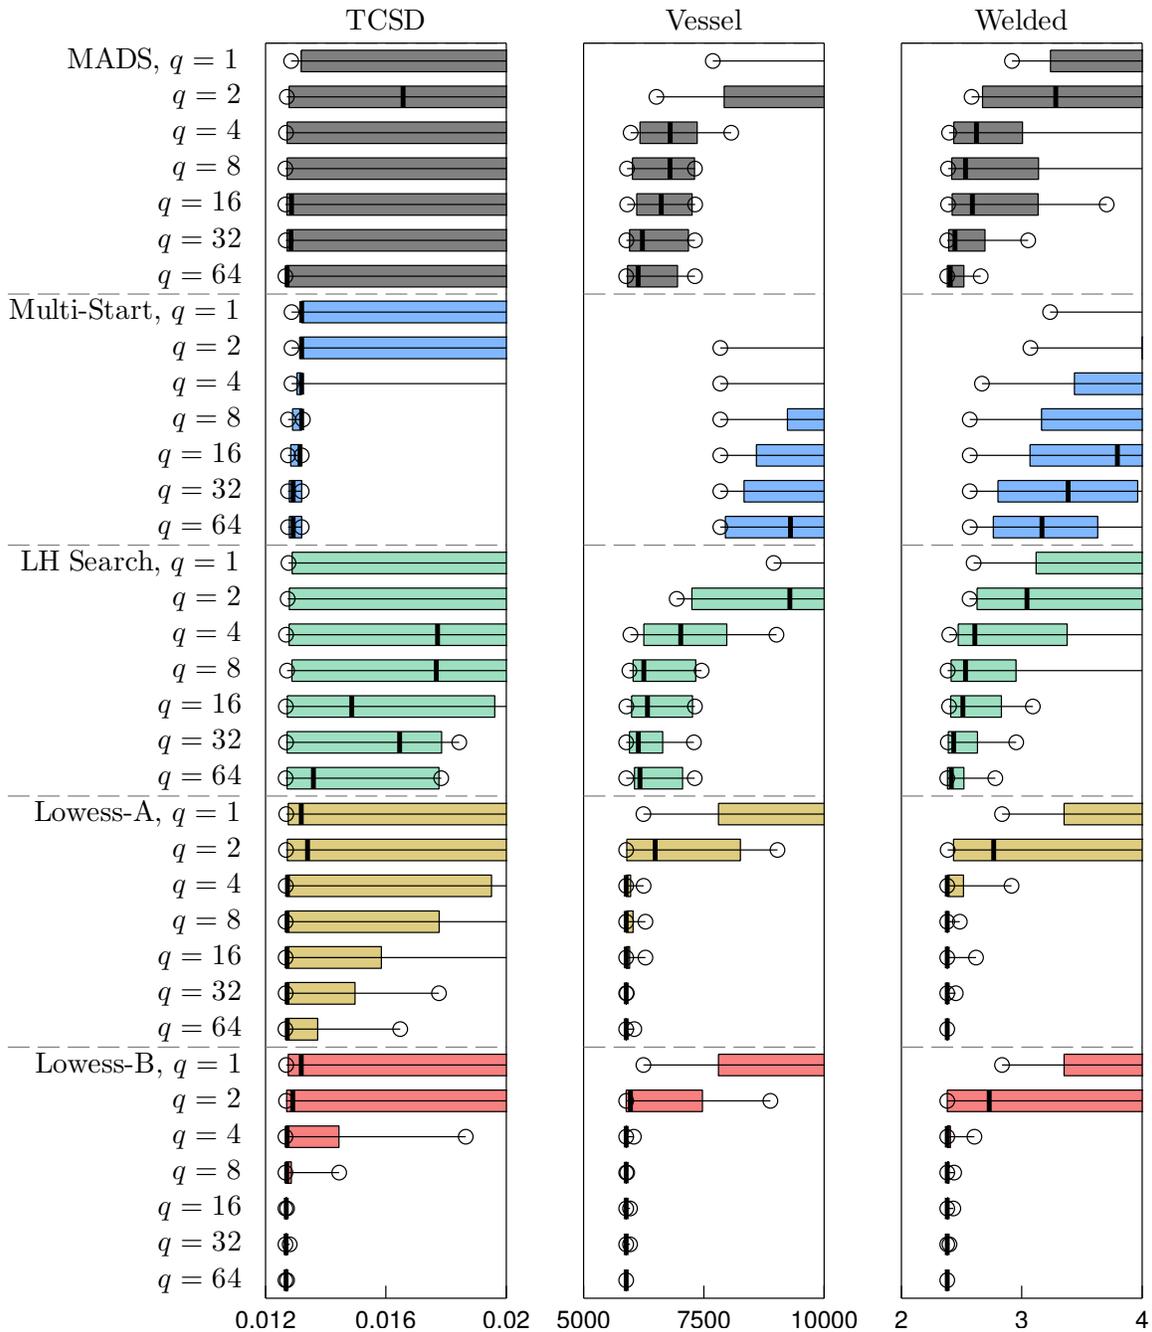


Figure 1: Performance summary for the TCSD, Vessel, and Welded problems over 50 runs

### Convergence rate

We now examine the convergence rate of the solvers for the case where  $q = 64$ . Figures 3 and 4 depict the evolution of the median objective function value of 50 runs as a function of the number of block evaluations. For each problem, the plots on the left compare the convergence of the five solvers with blocks of size  $q = 64$  while the plots on the right compare the convergence of the best-performing solver i.e., “Lowess-B”, for block sizes ranging from  $q = 1$  to 64.

We can conclude that “Lowess-B” yields the best solutions faster than any other solver (for  $q = 64$ ). The worst-performing solvers are “Multi-Start” for problems TCSD and Solar 1 and “LH Search” for

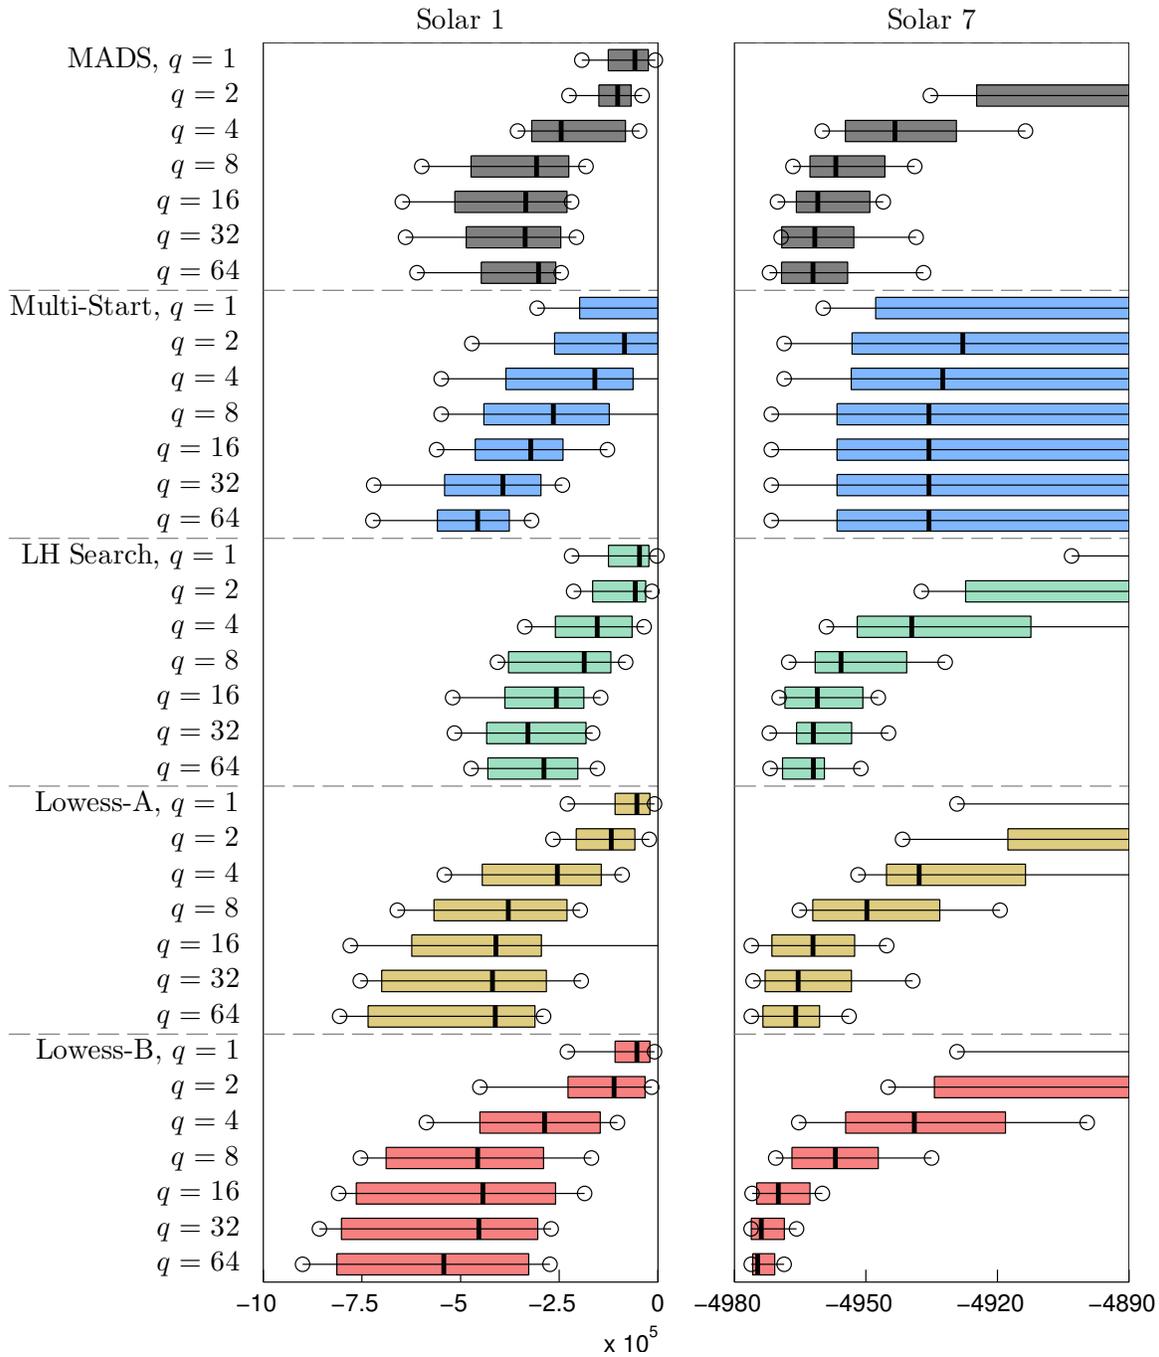


Figure 2: Performance summary for the Solar 1 and Solar 7 problems over 50 runs

problems Vessel, Welded, and Solar 7. It is also notable that although “MADS” does not use the SEARCH step, it performs generally well, except for Solar 1. “Lowess-A” performed well but does not clearly outperform other solvers.

Considering the performance of “Lowess-B” as a function of  $q$ , we observe that, as expected, convergence improves for larger values of  $q$ . Depending on the problem, there may be a saturation point beyond which an increase of  $q$  does not effect an improvement. E.g., a saturation point arises around  $q = 8$  or  $16$  for TCSD, Vessel, and Welded. On the contrary,  $q$  could be even larger than 64 for Solar 1 as more CPUs can be utilized.

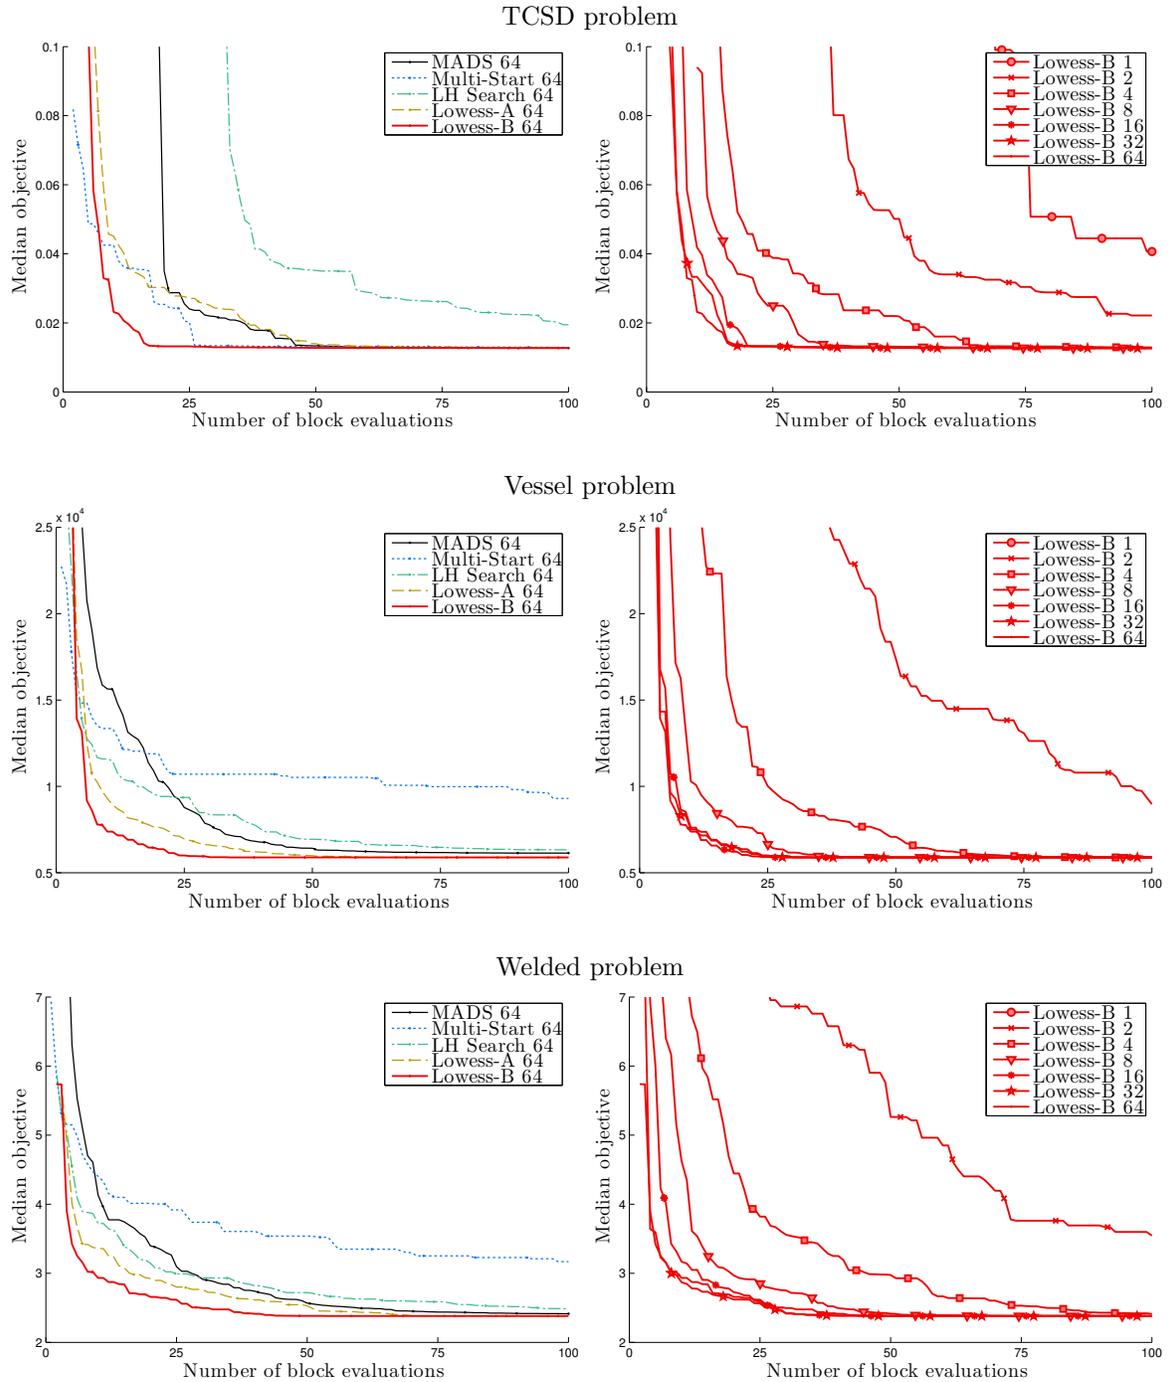


Figure 3: Results for the TCSD, Vessel, and Welded problems; median objective value of 50 runs

### Performance profiles

We now consider performance profiles, which indicate the percentage of runs where the problem is solved within a deviation from the best known solution  $\tau$  under a budget of function evaluations [14]. Specifically, for each solver  $s$ , each instance  $r$  and problem  $p$ , we compute the number of block evaluations  $b_{s,p,r}(\tau)$  such that

$$\frac{|f_{s,b,p,r} - f_p^*|}{|f_p^*|} \leq \tau, \tag{10}$$

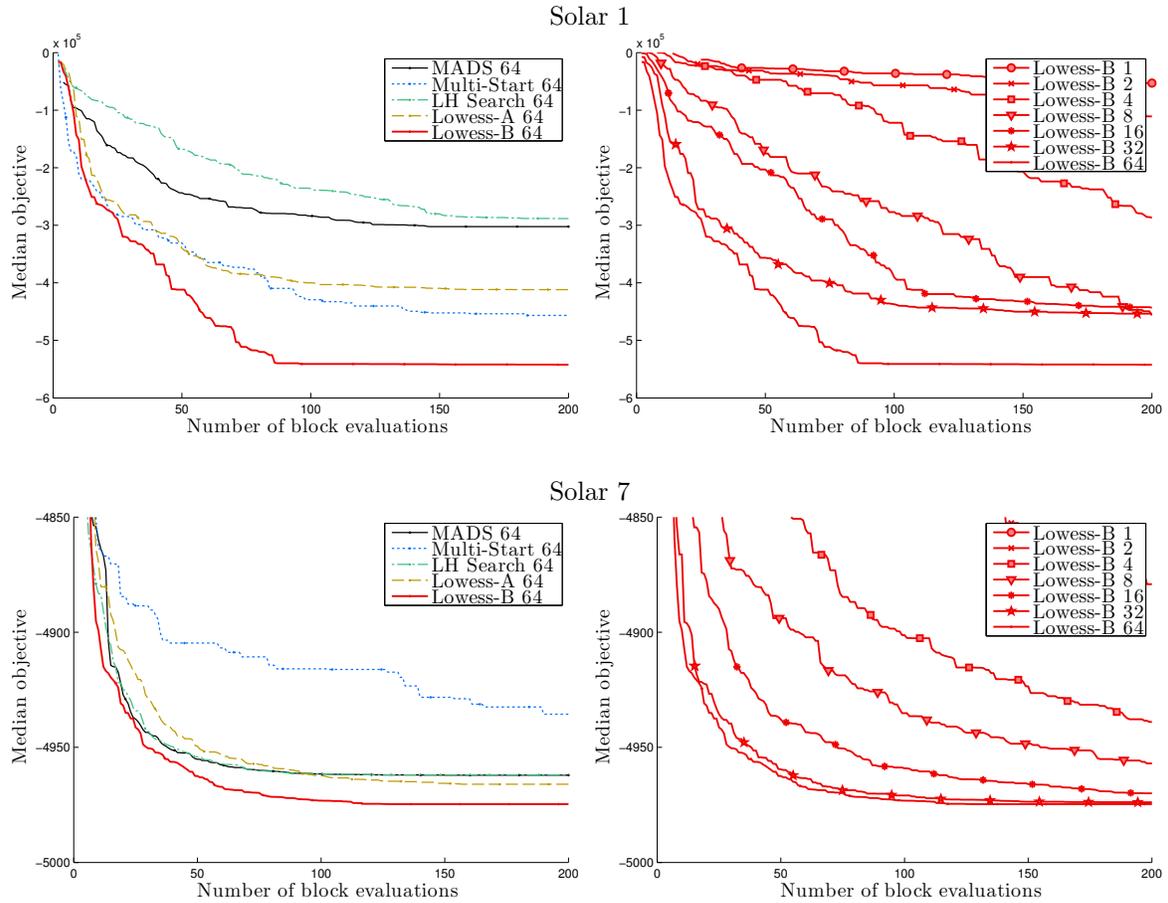


Figure 4: Results for the Solar 1 and Solar 7 problems; median objective value of 50 runs

where  $f_p^*$  is the best known objective value for problem  $p$  and  $f_{s,b,p,r}$  is the value obtained with the solver  $s$  after  $b$  block evaluations. Let  $b_{p,r}^{\min}(\tau)$  be the smallest budget for solving the instance  $r$  of problem  $p$  with deviation  $\tau$ , i.e.,

$$b_{p,r}^{\min}(\tau) = \min_s b_{s,p,r}(\tau).$$

Then, we can plot the proportion of runs of solver  $s$  that satisfy Equation (10) at a multiple  $\alpha$  of the smallest budget, i.e.,  $\alpha b_{p,r}^{\min}(\tau)$  block evaluations. Figure 5 depicts the performance profiles of the five considered problems for  $q = 64$  over the 50 runs (instances) for  $\tau$  values that range from  $10^{-1}$  to  $10^{-4}$ .

Note that higher curves imply better solver performance. Moreover, a solver that performs well for small values of  $\alpha$  is a solver that can solve, for the considered  $\tau$ , a large number of problems with a small evaluation budget. Figure 5 confirms our previous observations: “Lowess-B” outperforms all solvers, followed by “Lowess-A” and “LH Search” in second and third position. “MADS” and “Multi-Start” are in the last position. For large tolerances, “MADS” does better than “Multi-Start”, and better than “LH Search” for small values of  $\alpha$  ( $\leq 4$ ). For small precisions, “MADS” is outperformed by the other solvers.

The most interesting observation is the significant gap between the performance curves of LOWESS solvers and the ones from the three other solvers. The gap increases as  $\tau$  decreases. From moderate to lower  $\tau$  values ( $\leq 10^{-2}$ ), “Lowess-B” systematically solves at least twice more problems than “LH Search”, “Multi-Start” and “MADS”. It is unusual to observe such clear differences on performance profiles. “Lowess-A” performs almost as well as “Lowess-B”, particularly when  $\tau$  is small (around  $10^{-4}$ ), but needs at least four times more block evaluations to achieve this.

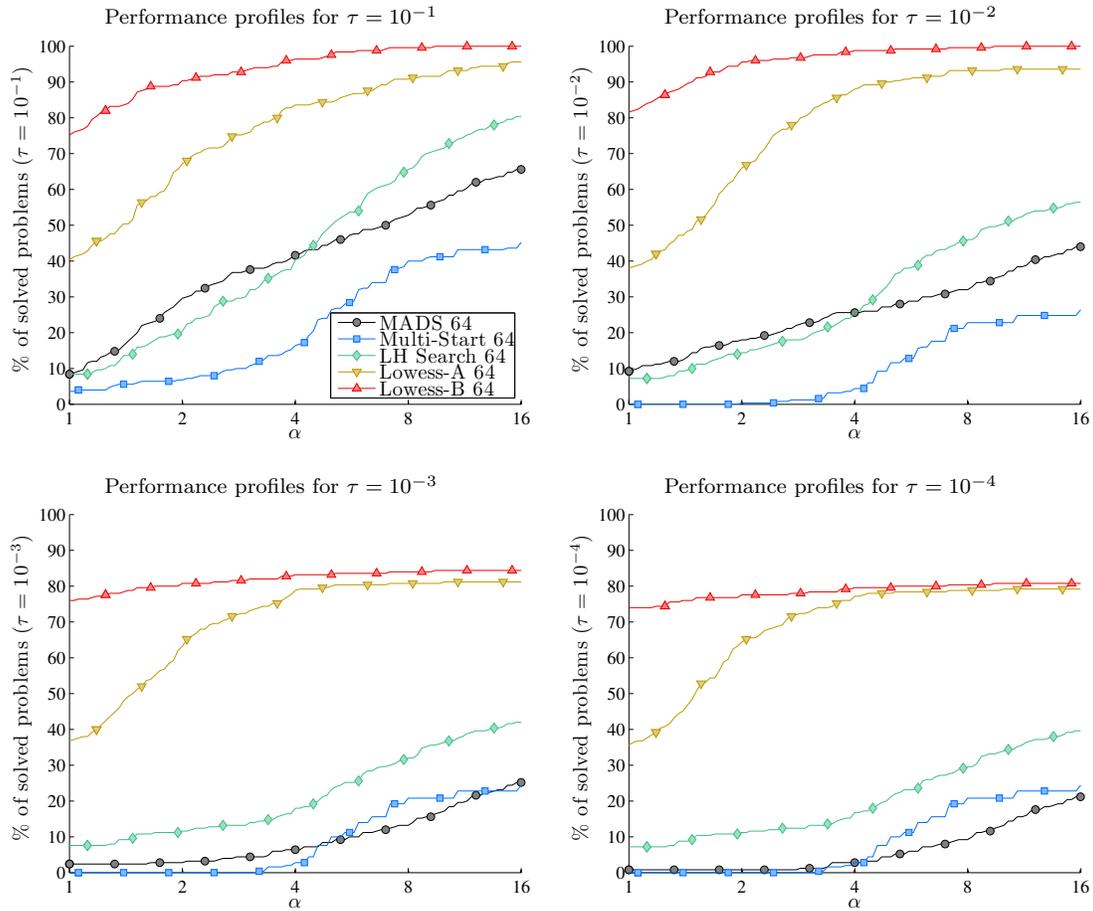


Figure 5: Performance profiles for  $q = 64$  over the 50 runs of all five problems

### Scalability analysis

We wish to establish the reduction of wall-clock time when using additional resources for each solver. To that end, we follow the methodology proposed in [20]. We define  $f_{s,b,p,r,q}$  as the value of the objective function obtained by solver  $s$  after  $b$  block evaluations on instance  $r$  of problem  $p$  when using  $q$  CPUs. We also define the *reference* objective value as the best value achieved with only one CPU ( $q = 1$ ), i.e.,

$$f_{s,p,r}^{\text{ref}} = \min_b f_{s,b,p,r,1}, \tag{11}$$

and  $b_{s,p,r,q}^{\text{ref}}$  the number of block evaluations necessary to reach  $f_{s,p,r}^{\text{ref}}$  when  $q$  CPUs are used, i.e.,

$$b_{s,p,r,q}^{\text{ref}} = \min\{b : f_{s,p,r}^{\text{ref}} \leq f_{s,b,p,r,q}\}. \tag{12}$$

The *speed-up* of solver  $s$  when solving with  $q$  CPUs is defined as

$$\text{speed-up}(s, q) = \text{geomean}_{p,r} \left( \frac{b_{s,p,r,q}^{\text{ref}}}{b_{s,p,r,1}^{\text{ref}}} \right) \tag{13}$$

and its *efficiency* as

$$\text{efficiency}(s, q) = \frac{\text{speed-up}(s, q)}{q}. \tag{14}$$

Figure 6 depicts the speed-up and efficiency values obtained by our numerical experiments.

Perfect scalability is obtained when the speed-up is equal to  $q$  and the efficiency is equal to 1. The speed-up curves show that the power introduced by new CPUs decreases as their number increases.

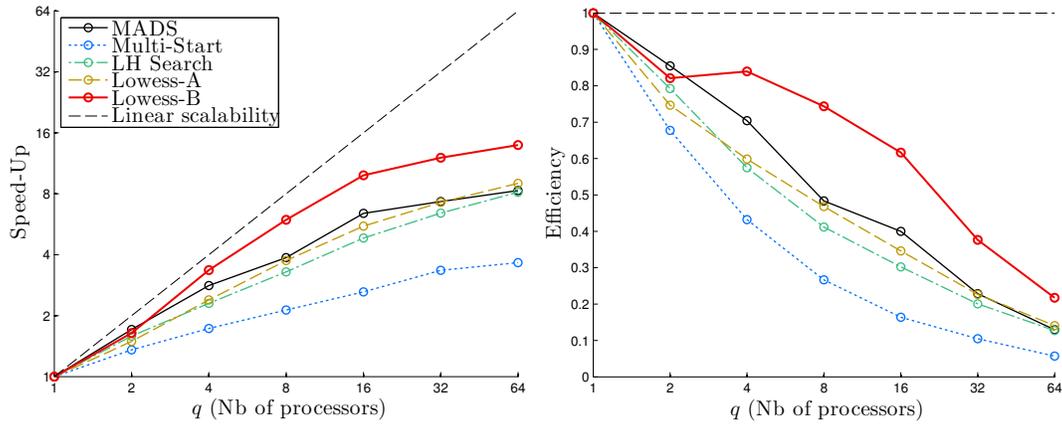


Figure 6: Speed-up and efficiency

This was observed in Figures 3 and 4 where problems exhibited saturation around  $q = 8$  and 16. “Lowess-B” achieves the best speed-up, followed by “MADS”. “Lowess-A” and “LH Search” come next, followed by the worst-performing solver, namely “Multi-Start”. We conclude that it is better and more productive to proceed with one search performing  $q$  parallel evaluations instead of conducting  $q$  independent searches consisting of a single evaluation.

The efficiency curves demonstrate rapid decrease except for “Lowess-B”; its rate exhibits a bump on its efficiency curve at  $q = 4$ . For  $q = 2$ , only methods 3 and 4 are used to generate candidates. For  $q \geq 4$ , methods 5 and 6 are also used. The aforementioned bump highlights the important contributions of these methods to the efficiency of “Lowess-B”.

## 6 Conclusion

Linear LOWESS models with optimized kernel shapes and coefficients seem to provide high-performing surrogates of the blackboxes. The use of diverse selection methods (3 to 6) enables an efficient exploration of the design space, accelerates local convergence, and makes optimal use of additional CPU resources. Methods 5 and 6 are particularly efficient, outperforming the other selection methods. This means that the way surrogates are used by method 1 is not effective. Similarly, the diversification strategy of method 2 is not adequate to select points that lie far enough from the ones already evaluated.

We cannot draw a definite conclusion about which of the methods 5 or 6 is better than the other. We believe that the good performance of “Lowess-B” is due to using method 5.

The proposed selection methods are not specific to the LOWESS model considered here; they are applicable to any surrogates. We believe that they will work well with reduced-fidelity (or variable-fidelity) physical-based models since high- and low-fidelity models typically have similarly structured solution domains. The selection methods are also applicable to other algorithms using surrogates to identify promising points to evaluate.

# Appendix

## A LOWESS predictions

As a convention, we denote with  $\boldsymbol{\xi} \in \mathcal{X} \subseteq \mathbb{R}^n$  the point of the design space where we want to predict the value of the blackbox output. Locally weighted scatterplot smoothing (LOWESS) models build a local linear regression at the point  $\boldsymbol{\xi}$  where the blackbox output  $[f \ c_1 \dots c_m]$  are to be estimated [4, 10, 11, 12, 13, 31]. This local regression emphasizes data points that are close to  $\boldsymbol{\xi}$ . The interested reader can refer to [31] for details about the method described below. We consider here only local linear regressions; local quadratic regressions and Tikhonov regularization are considered in [31]. On the contrary, while only a Gaussian kernel was considered in [31], six others are added here as kernel functions.

We define the output matrix  $\mathbf{Y} \in \mathbb{R}^{p \times (m+1)}$ , the design matrix  $\mathbf{Z}_\boldsymbol{\xi} \in \mathbb{R}^{p \times (n+1)}$ , and the weight matrix  $\mathbf{W}_\boldsymbol{\xi} \in \mathbb{R}^{p \times p}$ :

$$\mathbf{Y} = \begin{bmatrix} f(\mathbf{x}_1) & c_1(\mathbf{x}_1) & \dots & c_m(\mathbf{x}_1) \\ \vdots & \vdots & & \vdots \\ f(\mathbf{x}_p) & c_1(\mathbf{x}_p) & \dots & c_m(\mathbf{x}_p) \end{bmatrix}, \quad \mathbf{Z}_\boldsymbol{\xi} = \begin{bmatrix} 1 & (\mathbf{x}_1 - \boldsymbol{\xi})^\top \\ \vdots & \vdots \\ 1 & (\mathbf{x}_p - \boldsymbol{\xi})^\top \end{bmatrix}, \quad \mathbf{W}_\boldsymbol{\xi} = \begin{bmatrix} w_1(\boldsymbol{\xi}) & & & \\ & \ddots & & \\ & & \ddots & \\ & & & w_p(\boldsymbol{\xi})^\top \end{bmatrix}. \quad (15)$$

The details of the computation of  $w_i(\boldsymbol{\xi})$  are described in Section A.1. Then, we define  $\mathbf{u}_\boldsymbol{\xi} \in \mathbb{R}^{n+1}$  as the first column of  $(\mathbf{Z}_\boldsymbol{\xi}^\top \mathbf{W}_\boldsymbol{\xi} \mathbf{Z}_\boldsymbol{\xi})^{-1}$ , which means that  $\mathbf{u}_\boldsymbol{\xi}$  is the solution of the linear system  $\mathbf{Z}_\boldsymbol{\xi}^\top \mathbf{W}_\boldsymbol{\xi} \mathbf{Z}_\boldsymbol{\xi} \mathbf{u}_\boldsymbol{\xi} = \mathbf{e}_1$ . The prediction of the blackbox outputs at  $\boldsymbol{\xi}$  is then

$$\hat{\mathbf{y}}(\boldsymbol{\xi}) = [ \hat{f}(\boldsymbol{\xi}) \quad \hat{c}_1(\boldsymbol{\xi}) \quad \dots \quad \hat{c}_m(\boldsymbol{\xi}) ] = \mathbf{u}_\boldsymbol{\xi}^\top \mathbf{Z}_\boldsymbol{\xi}^\top \mathbf{W}_\boldsymbol{\xi} \mathbf{Y}. \quad (16)$$

The cross-validation value  $\hat{y}(\mathbf{x}_i)$  (i.e., the value of the LOWESS model at  $\mathbf{x}_i$  when the data point  $\mathbf{x}_i$  is not used to build the model) are computed by setting  $w_i$  to 0. Unfortunately, unlike for radial basis function (RBF) models or polynomial response surfaces (PRSs), we do not know any computational shortcut allowing a more efficient computation of the values of  $\hat{y}$ . However, each value  $\hat{y}(\mathbf{x}_i)$  is computed at the same computational cost as a prediction  $\hat{y}(\mathbf{x}_i)$ .

### A.1 Weights computation in LOWESS models

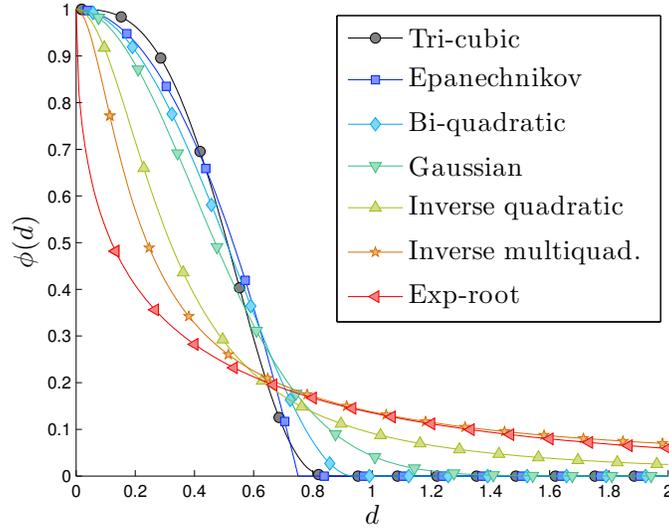
The weight  $w_i(\boldsymbol{\xi})$  quantifies the relative importance of the data point  $\mathbf{x}_i$  in the construction of the local regression at  $\boldsymbol{\xi}$ . Like for kernel smoothing, it relies on a kernel function  $\phi$  and depends on the distance between  $\boldsymbol{\xi}$  and  $\mathbf{x}_i$ . In our method, we use

$$w_i(\boldsymbol{\xi}) = \phi \left( \lambda \frac{\|\boldsymbol{\xi} - \mathbf{x}_i\|_2}{d_{n+1}(\boldsymbol{\xi})} \right), \quad (17)$$

where  $\phi(d)$  is one of the kernel functions described in Table 1 and Figure 7. All kernel functions are normalized so that  $\phi(0) = 1$  and, if applicable,  $\int_{\mathbb{R}} \phi = 1$ . As the integral of the inverse multi-quadratic kernel does not converge, the normalization constant 52.015 is introduced to minimize the  $\mathcal{L}^2$  distance between the inverse multi-quadratic and inverse quadratic kernel. The parameter  $\lambda > 0$  controls the general shape of the model, and  $d_{n+1}(\boldsymbol{\xi})$  is a local scaling coefficient that estimates the distance of the  $n + 1^{\text{th}}$  closest data point to  $\boldsymbol{\xi}$ . The kernel function  $\phi$  and the shape parameter  $\lambda$  are chosen to minimize the aggregate order error with cross-validation (AOECV) described in Section A.2. The fact that some of the available kernel function have a compact domain gives to LOWESS models the ability to ignore outliers or aberrant data points. As an example, if the blackbox fails to compute correctly the objective function for a given data point, the value returned by the blackbox might be an arbitrarily high value (e.g.,  $1.8 \cdot 10^{308}$  for a C++ code returning the standard max double). With non-compact kernel function, this would perturb the LOWESS model on the entire design space. However, if there is no such aberrant data points, non-compact kernel functions tend to yield better results.

**Table 1: Possible values for the kernel function  $\phi$** 

#	Kernel name	$\phi : \mathbb{R} \rightarrow \mathbb{R}^+$	Compact domain
1	Tri-cubic	$\phi(d) = (1 -  \frac{162}{140}d ^3)\mathbb{1}_{ d  \leq \frac{140}{162}}$	Yes
2	Epanechnikov	$\phi(d) = (1 - \frac{16}{9}d^2)\mathbb{1}_{ d  \leq \frac{3}{4}}$	Yes
3	Bi-quadratic	$\phi(d) = (1 -  \frac{16}{15}d ^2)^2\mathbb{1}_{ d  \leq \frac{15}{16}}$	Yes
4	Gaussian	$\phi(d) = \exp(-\pi d^2)$	No
5	Inverse quadratic	$\phi(d) = \frac{1}{1 + \pi^2 d^2}$	No
6	Inverse multi-quadratic	$\phi(d) = \frac{1}{\sqrt{1 + 52.015d^2}}$	No
7	Exp-root	$\phi(d) = \exp(-2\sqrt{ d })$	No

**Figure 7: Representation of the 7 kernels listed in Table 1**

To obtain a model  $\hat{y}$  that is differentiable everywhere, [31] defines  $d_{n+1}(\boldsymbol{\xi})$  such that the expected number of training points in a ball of center  $\boldsymbol{\xi}$  and radius  $d_{n+1}(\boldsymbol{\xi})$  is  $n + 1$ :

$$\mathbb{E} \left[ \text{card} \left\{ \mathbf{x}_i : \mathbf{x}_i \in \mathbf{X}, \|\boldsymbol{\xi} - \mathbf{x}_i\|_2 \leq d_{n+1}(\boldsymbol{\xi}) \right\} \right] = n + 1. \quad (18)$$

Moreover, [31] observes that the values  $\{\|\boldsymbol{\xi} - \mathbf{x}_i\|_2^2\}_{i=1, \dots, p}$  can be fitted well by a Gamma distribution and therefore defines the local scaling parameter as

$$d_{n+1}(\boldsymbol{\xi}) = \sqrt{g^{(-1)} \left( \frac{\mu_{\boldsymbol{\xi}}^2}{\sigma_{\boldsymbol{\xi}}^2}, \frac{\sigma_{\boldsymbol{\xi}}^2}{\mu_{\boldsymbol{\xi}}}; \frac{n+1}{p} \right)}, \quad (19)$$

where  $\mu_{\boldsymbol{\xi}}$  (resp.  $\sigma_{\boldsymbol{\xi}}^2$ ) denotes the mean (resp. variance) of  $\|\boldsymbol{\xi} - \mathbf{x}_i\|_2^2$  over  $\mathbf{X}$  and  $g^{(-1)}(k, \theta; \cdot)$  is the inverse function of the cumulative density function of a Gamma distribution with shape parameter  $k$  and scale parameter  $\theta$ .

## A.2 Aggregate Order Error with Cross-Validation

The AOECV is an error metric that aims at quantifying the quality of a *multi-output* surrogate model. Specifically, it aims at quantifying the discrepancy between problems  $(P)$  and  $(\hat{P})$  for a given surrogate model. We first define the aggregate constraint violation function [15]  $h(\mathbf{x}) = \sum_{j=1}^m \max\{0, c_j(\mathbf{x})\}^2$ . Note that other definitions of  $h$  are possible (notably: number of violated constraints, most violated constraint, etc.) but as the previous definition of  $h$  is used in the main MADS instance to solve  $(P)$ , we need to use the same aggregate constraint in our definition of the AOECV.

We then define the order operators

$$\mathbf{x} \prec \mathbf{x}' \Leftrightarrow \begin{cases} h(\mathbf{x}) < h(\mathbf{x}') \\ \text{or} \\ h(\mathbf{x}) = h(\mathbf{x}') \text{ and } f(\mathbf{x}) < f(\mathbf{x}'), \end{cases} \quad (20)$$

$$\mathbf{x} \preceq \mathbf{x}' \Leftrightarrow \text{not}(\mathbf{x}' \prec \mathbf{x}) \quad (21)$$

which are transitive. In particular, the incumbent solution  $\mathbf{x}^t$  of the original problem ( $P$ ) is such that  $\mathbf{x}^t \preceq \mathbf{x}$ ,  $\forall \mathbf{x} \in \mathbf{X}$ . Similarly, a global minimizer  $\mathbf{x}^*$  is such that  $\mathbf{x}^* \preceq \mathbf{x}$ ,  $\forall \mathbf{x} \in \mathcal{X}$ . By the same principle, we define the operator  $\hat{\preceq}$  by using the cross-validation values  $\hat{f}$  and  $\hat{h} = \sum_{j=1}^m \max\{0, \hat{c}_j(\mathbf{x})\}^2$  instead of  $f$  and  $h$ . We then define the aggregated order error with cross-validation (AOECV) metric:

$$\mathcal{E}_{AOECV} = \frac{1}{p^2} \sum_{i=1}^p \sum_{j=1}^p \text{xor} \left( \mathbf{x}_i \prec \mathbf{x}_j, \mathbf{x}_i \hat{\preceq} \mathbf{x}_j \right). \quad (22)$$

where  $\text{xor}$  is the exclusive or operator (i.e.,  $\text{xor}(A, B) = 1$  if the booleans  $A$  and  $B$  differ and 0 otherwise). The metric allows to quantify how often the model is able to correctly decide which of two points is better.

The shape parameter  $\lambda$  and the kernel function  $\phi$  are then chosen to minimize  $\mathcal{E}_{AOECV}(\lambda, \phi)$ . If two couples  $(\lambda, \phi)$  lead to the same metric value (because of the piecewise-constant nature of the metric), the couple with the smallest value of  $\lambda$  (i.e., the smoother model) is preferred.

## B Detailed description of the test problems

The five engineering design application problems considered are listed in Table 2. Problem size is reflected by  $n$  and  $m$ , where  $n$  denotes the number of design variables and  $m$  the number of general nonlinear inequality constraints. Table 2 also indicates whether any variables are integer or unbounded and reports the best known value of the objective function.

**Table 2: Summary of the five engineering design optimization problems**

Problem name	$n$	$m$	Integer variables	Infinite bounds	Best objective function value
TCSD	3	4	No	No	0.0126652
Vessel	4	4	No	No	5,885.332
Welded	4	6	No	No	2.38096
Solar 1	9	5	Yes	Yes	-900,417
Solar 7	7	6	Yes	Yes	-4,976.17

The Tension/Compression Spring Design (TCSD) problem consists of minimizing the weight of a spring under mechanical constraints [3, 9, 16]. The design variables define the geometry of the spring. The constraints concern shear stress, surge frequency and minimum deflection. The best known solution, denoted  $\mathbf{x}^*$ , and the bounds on the variables, denoted by  $\underline{\mathbf{x}}$  and  $\bar{\mathbf{x}}$ , are given in Table 3.

**Table 3: Variables of the TCSD problem**

Variable description	$\underline{\mathbf{x}}$	$\bar{\mathbf{x}}$	$\mathbf{x}^*$
Mean coil diameter	0.05	2	0.051686696913218
Wire diameter	0.25	1.3	0.356660815351066
Number of active coil	2	15	11.292312882259289

The Vessel problem considers the optimal design of a compressed air storage tank [16, 21]. The design variables define the geometry of the tank. The constraints are related to the volume, pressure, and solidity of the tank. The objective is to minimize the total cost of the tank, including material and labour. Table 4 lists the variable bounds and the best known solution.

**Table 4: Variables of the Vessel problem**

Variable description	$\underline{x}$	$\bar{x}$	$\mathbf{x}^*$
Thickness of the vessel	0.0625	6.1875	0.778168641330718
Thickness of the head	0.0625	6.1875	0.384649162605973
Inner radius	10	200	40.319618721803231
Length of the vessel without heads	10	200	199.99999998822659

The Welded (or welded beam design) problem (Version I) consists of minimizing the construction cost of a beam under shear stress, bending stress, load and deflection constraints [16, 28]. The design variables define the geometry of the beam and the characteristics of the welded joint. Table 5 lists the variable bounds and the best known solution.

**Table 5: Variables of the Welded problem**

Variable description	$\underline{x}$	$\bar{x}$	$\mathbf{x}^*$
Thickness of the weld	0.1	2	0.244368407428265
Length of the welded joint	0.1	10	6.217496713101864
Width of the beam	0.1	10	8.291517255567012
Thickness of the beam	0.1	2	0.244368666449562

The Solar1 and Solar7 problems consider the optimization of a solar farm, including the heliostat field and/or the receiver [17]. The Solar1 optimization problem aims at maximizing the energy received over a period of 24 hours under several constraints of budget and heliostat field area. This problem has one integer variable that has no upper bound. Table 6 lists the variable bounds and the best known solution.

**Table 6: Variables of the Solar1 problem**

Variable description	$\underline{x}$	$\bar{x}$	$\mathbf{x}^*$
Heliostat height	1	40	6.165258994385601
Heliostat width	1	40	10.571794049143792
Tower height	20	250	91.948461670428486
Receiver aperture height	1	30	6.056202026704944
Receiver aperture width	1	30	11.674984434929991
Max number of heliostats (Integer)	1	$+\infty$	1507
Field maximum angular span	1	89	51.762281627953051
Minimum distance to tower	0.5	20	1.347318830713629
Maximum distance to tower	1	20	14.876940809562798

The Solar7 problem aims at maximizing the efficiency of the receiver over a period of 24 hours, for a given heliostats field, under 6 binary constraints [17]. This problem has one integer variable that has no upper bound. The objective function is the energy transferred to the molten salt. Table 7 lists the variable bounds and the best known solution.

**Table 7: Variables of the Solar7 problem**

Variable description	$\underline{x}$	$\bar{x}$	$\mathbf{x}^*$
Aperture height	1	30	11.543687848308958
Aperture width	1	30	15.244236061098078
Outlet temperature	793	995	803.000346734710888
Number of tubes (Integer)	1	$+\infty$	1292
Insulation thickness	0.01	5	3.399190219909724
Tubes inside diameter	0.005	0.1	0.010657067457678
Tubes outside diameter	0.0055	0.1	0.011167646941518

## References

- [1] E. Alba, G. Luque, and S. Nasmachnow. Parallel metaheuristics: recent advances and new trends. *International Transactions in Operational Research*, 20(1):1–48, 2013.
- [2] R. Alizadeh, J.K. Allen, and F. Mistree. Managing computational complexity using surrogate models: a critical review. *Research in Engineering Design*, 2020.
- [3] J. Arora. *Introduction to Optimum Design*. Elsevier Science, 2004.
- [4] C.G. Atkeson, A.W. Moore, and S. Schaal. Locally weighted learning. *Artificial Intelligence Review*, pages 11–73, 1997.
- [5] C. Audet and J.E. Dennis, Jr. Mesh adaptive direct search algorithms for constrained optimization. *SIAM Journal on Optimization*, 17(1):188–217, 2006.
- [6] C. Audet and J.E. Dennis, Jr. A progressive barrier for derivative-free nonlinear programming. *SIAM Journal on Optimization*, 20(1):445–472, 2009.
- [7] C. Audet, M. Kokkolaras, S. Le Digabel, and B. Talgorn. Order-based error for managing ensembles of surrogates in derivative-free optimization. *Journal of Global Optimization*, 70(3):645–675, 2018.
- [8] C. Audet, S. Le Digabel, C. Tribes, and V. Rochon Montplaisir. The NOMAD project. Software available at <https://www.gerad.ca/nomad>.
- [9] A.D. Belegundu. *A Study of Mathematical Programming Methods for Structural Optimization*. University of Iowa, 1982.
- [10] W.S. Cleveland. Robust locally weighted regression and smoothing scatterplots. *Journal of the American Statistical Association*, 74:829–836, 1979.
- [11] W.S. Cleveland. LOWESS: A Program for Smoothing Scatterplots by Robust Locally Weighted Regression. *The American Statistician*, 35(1), 1981.
- [12] W.S. Cleveland and S.J. Devlin. Locally weighted regression: An approach to regression analysis by local fitting. *Journal of the American Statistical Association*, 83:596–610, 1988.
- [13] W.S. Cleveland, S.J. Devlin, and E. Grosse. Regression by local fitting: methods, properties, and computational algorithms. *Journal of Econometrics*, 37(1):87–114, 1988.
- [14] E.D. Dolan and J.J. Moré. Benchmarking optimization software with performance profiles. *Mathematical Programming*, 91(2):201–213, 2002.
- [15] R. Fletcher and S. Leyffer. Nonlinear programming without a penalty function. *Mathematical Programming, Series A*, 91:239–269, 2002.
- [16] H. Garg. Solving structural engineering design optimization problems using an artificial bee colony algorithm. *Journal of Industrial and Management Optimization*, 10(3):777–794, 2014.
- [17] Mathieu Lemyre Garneau. Modelling of a solar thermal power plant for benchmarking blackbox optimization solvers. Master’s thesis, École Polytechnique de Montréal, 2015.
- [18] Raphael T. Haftka, Diane Villanueva, and Anirban Chaudhuri. Parallel surrogate-assisted global optimization with expensive functions – a survey. *Structural and Multidisciplinary Optimization*, 54(1):3–13, Jul 2016.
- [19] P. Hansen and N. Mladenović. Variable neighborhood search: principles and applications. *European Journal of Operational Research*, 130(3):449–467, 2001.
- [20] Prasad Jogalekar and Murray Woodside. Evaluating the scalability of distributed systems. *IEEE Trans. Parallel Distrib. Syst.*, 11(6):589–603, June 2000.
- [21] B. K. Kannan and S. N. Kramer. Augmented Lagrange multiplier based method for mixed integer discrete continuous optimization and its applications to mechanical design. *Journal of Mechanical Design*, 65:103–112+, 1993.

- [22] O. Kramer, D.E. Ciaurri, and S. Koziel. Derivative-free optimization. In S. Koziel and X.S. Yang, editors, *Computational Optimization, Methods and Algorithms*, volume 356 of *Studies in Computational Intelligence*, pages 61–83. Springer, 2011.
- [23] S. Le Digabel. Algorithm 909: NOMAD: Nonlinear optimization with the MADS algorithm. *ACM Transactions on Mathematical Software*, 37(4):44:1–44:15, 2011.
- [24] S. Le Digabel, M.A. Abramson, C. Audet, and J.E. Dennis, Jr. Parallel versions of the MADS algorithm for black-box optimization. In *Optimization days, Montreal, May 2010*. GERAD. Slides available at [http://www.gerad.ca/Sebastien.Le.Digabel/talks/2010\\_JOPT\\_25mins.pdf](http://www.gerad.ca/Sebastien.Le.Digabel/talks/2010_JOPT_25mins.pdf).
- [25] M.D. McKay, R.J. Beckman, and W.J. Conover. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 21(2):239–245, 1979.
- [26] N. Mladenović and P. Hansen. Variable neighborhood search. *Computers and Operations Research*, 24(11):1097–1100, 1997.
- [27] Mahdi Pourbagian, Bastien Talgorn, WagdiG. Habashi, Michael Kokkolaras, and Sébastien Le Digabel. Constrained problem formulations for power optimization of aircraft electro-thermal anti-icing systems. *Optimization and Engineering*, pages 1–31, 2015.
- [28] Singiresu S. Rao. *Engineering Optimization: Theory and Practice*, 3rd Edition. Wiley-Interscience, 1996.
- [29] T.J. Santner, B.J. Williams, and W.I. Notz. The Design and Analysis of Computer Experiments, chapter 5.2.2, Designs Generated by Latin Hypercube Sampling, pages 127–132. Springer, New York, NY, 2003.
- [30] B. Talgorn. SGTELIB: Surrogate model library for derivative-free optimization. <https://github.com/bbopt/sgtelib>, 2019.
- [31] B. Talgorn, C. Audet, M. Kokkolaras, and S. Le Digabel. Locally weighted regression models for surrogate-assisted design optimization. *Optimization and Engineering*, 19(1):213–238, 2018.
- [32] B. Talgorn, S. Le Digabel, and M. Kokkolaras. Statistical Surrogate Formulations for Simulation-Based Design Optimization. *Journal of Mechanical Design*, 137(2):021405–1–021405–18, 2015.