

Multi-fidelity constraints in blackbox optimization

S. Alarie, C. Audet, M. Diago, S. Le Digabel, X. Lebeuf

G-2026-01

January 2026

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 : S. Alarie, C. Audet, M. Diago, S. Le Digabel, X. Lebeuf (Janvier 2026). Multi-fidelity constraints in blackbox optimization, Rapport technique, Les Cahiers du GERAD G-2026-01, GERAD, HEC Montréal, Canada.

Avant de citer ce rapport technique, veuillez visiter notre site Web (<https://www.gerad.ca/fr/papers/G-2026-01>) 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, 2026
– Bibliothèque et Archives Canada, 2026

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: S. Alarie, C. Audet, M. Diago, S. Le Digabel, X. Lebeuf (January 2026). Multi-fidelity constraints in blackbox optimization, Technical report, Les Cahiers du GERAD G-2026-01, GERAD, HEC Montréal, Canada.

Before citing this technical report, please visit our website (<https://www.gerad.ca/en/papers/G-2026-01>) 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, 2026
– Library and Archives Canada, 2026

Multi-fidelity constraints in blackbox optimization

Stéphane Alarie ^{a, c}

Charles Audet ^{b, c}

Miguel Diago ^a

Sébastien Le Digabel ^{b, c}

Xavier Lebeuf ^{b, c}

^a Centre de recherche d'Hydro-Québec, Varennes (Qc), Canada, J3X 1S1

^b Polytechnique Montréal, Montréal (Qc), Canada, H3T 0A3

^c GERAD, Montréal (Qc), Canada, H3T 1J4

alarie.stephane@hydroquebec.com

diagomartinez.miguel@hydroquebec.com

charles.audet@gerad.ca

sebastien.le.digabel@gerad.ca

xavier.lebeuf@etud.polymtl.ca

January 2026

Les Cahiers du GERAD

G-2026-01

Copyright © 2026 Alarie, Audet, Diago, Le Digabel, Lebeuf

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:

- May download and print one copy of any publication from the public portal for the purpose of private study or research;
- May not further distribute the material or use it for any profit-making activity or commercial gain;
- May 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 : This work studies constrained blackbox optimization problems that cannot be solved in reasonable time due to prohibitive computational costs. This challenge is especially prevalent in industrial applications, where blackbox evaluations are costly. However, constraints can be evaluated at various fidelities at a lower computational cost. More specifically, this work targets situations in which the infeasibility of each individual constraint can be detected at lower fidelities, and where a large discrete number of fidelities are available. Moreover, highly discontinuous problems which may fail to evaluate are considered, such that direct search methods are preferred to model-based ones. To this effect, the Interruptible Direct Search (IDS) and the Dynamic Interruptible Direct Search (DIDS) algorithms are proposed to leverage feasibility assessments from various fidelity levels to avoid high cost evaluations. The results show highly increased performances from NOMAD when it is paired with IDS or DIDS.

Keywords: Blackbox optimization; derivative-free optimization; multi-fidelity; constrained optimization; direct search methods; static surrogates

Résumé : Ce travail étudie les problèmes d'optimisation de boîte noire contraints qui ne peuvent pas être résolus dans un temps raisonnable en raison de coûts de calcul prohibitifs. Ce défi est particulièrement présent dans les applications industrielles, où les évaluations de boîte noire sont coûteuses. Cependant, les contraintes peuvent être évaluées à différents niveaux de fidélité, pour un coût de calcul inférieur. Plus précisément, ce travail cible des situations dans lesquelles l'irréalisabilité de chaque contrainte peut individuellement être détectée à des niveaux de fidélité plus faibles, et où un grand nombre discret de niveaux de fidélité est disponible. De plus, des problèmes fortement discontinus pouvant échouer lors de l'évaluation sont considérés, de sorte que les méthodes de recherche directe sont préférables aux méthodes basées sur des modèles. À cette fin, les algorithmes Interruptible Direct Search (IDS) et Dynamic Interruptible Direct Search (DIDS) sont proposés afin d'exploiter les estimations de réalisabilité issues de différents niveaux de fidélité pour éviter des évaluations à coût élevé. Les résultats montrent une haute augmentation des performances de NOMAD lorsqu'il est couplé avec IDS ou DIDS.

Mots clés : Optimisation de boîte noire; optimisation sans dérivées; multifoldérité; optimization sous contraintes; méthodes par recherche directe; substituts statiques

Acknowledgements: The solar problems are available at <https://github.com/bbopt/solar>. The NOMAD software package is available at <https://github.com/bbopt/nomad>. This work is partly supported by the NSERC Alliance-Mitacs Accelerate grant ALLRP 571311-21 (“Optimization of future energy systems”) in collaboration with Hydro-Québec. It is also funded by X. Lebeuf’s doctoral Hydro-Québec excellence grant.

1 Introduction

This work studies constrained multi-fidelity optimization problems of the form

$$\mathcal{P}(\phi) \quad \min_{x \in \Omega} \quad f(x, \phi), \quad \text{where} \quad \Omega = \{x \in X : c_j(x, \phi) \leq 0, j \in J\},$$

in which $X \subseteq \mathbb{R}^n$ is the bound constrained domain of the objective function $f : X \times [0, 1] \rightarrow \overline{\mathbb{R}} = \mathbb{R} \cup \{\infty\}$ and of the m relaxable quantifiable [25] constraint functions $c_j : X \times [0, 1] \rightarrow \overline{\mathbb{R}}$, for $j \in J = \{1, 2, \dots, m\}$. The set of feasible points is denoted by Ω . The parameter $\phi \in [0, 1]$ is the *fidelity level* at which these functions are evaluated. Hence, by convention, the problem that one wishes to solve is denoted by $\mathcal{P}(1)$ and is called the *truth*. The objective and constraint functions are provided by a blackbox process. They have no accessible analytical formulation, they might be highly discontinuous, and their derivatives are unavailable or may be non-existent [10]. This is why a direct search approach is adopted [17]. Moreover, the blackbox is expensive to evaluate, and may fail to execute. The set $\overline{\mathbb{R}}$ allows the use of common blackbox optimization tools. Notably, f is assigned an infinite value by an extreme barrier method at infeasible points, and the vector $c(x, \phi) = (c_1(x, \phi), c_2(x, \phi), \dots, c_m(x, \phi))$ may possess one or more infinite values when an evaluation fails.

The specificity of this work is that Problem $\mathcal{P}(\phi)$ is *multi-fidelity*, meaning that evaluating the blackbox requires not only specifying a trial point $x \in X$, but also selecting a fidelity value $\phi \in [0, 1]$ that controls the accuracy and computational cost of the evaluation. Lower fidelities correspond to lower precision, and generally, lower evaluation cost, and vice-versa. An evaluation using fidelity $\phi < 1$ can be interpreted as calling a static surrogate model [10, Ch. 13] that provides point-wise approximations of f and c . The computational cost required to evaluate a trial point $x \in X$ using fidelity level $\phi \in [0, 1]$ is denoted by $\lambda(x, \phi)$, where $\lambda : X \times [0, 1] \rightarrow \mathbb{R}_+$. This cost function is assumed to be unknown, but is usually increasing with respect to ϕ .

This work considers a finite and discrete subset of $L \in \mathbb{N}$ fidelities. It is described by the set of fidelity indices $I = \{1, 2, \dots, L\}$, where the sequence $\{\phi_i\}_{i \in I}$ is strictly increasing and $\phi_L = 1$. Two novel optimization algorithms are introduced: Interruptible Direct Search (IDS) and Dynamic Interruptible Direct Search (DIDS). They propose a strategy to reduce the computational cost of solving the optimization problem by identifying and exploiting the minimal fidelities required to determine that a trial point is deemed infeasible. This approach avoids costly high-fidelity evaluations in a context where solving the true Problem $\mathcal{P}(1)$ directly is impossible, as the computational effort is prohibitive. While IDS is applicable to any multi-fidelity problem, DIDS targets problems with *intermediary outputs*.

Definition 1. A multi-fidelity blackbox with domain X is said to return *intermediary outputs* for I if during an evaluation of trial point $x \in X$ at fidelity 1, all $f(x, \phi_i)$ and $c(x, \phi_i)$ values for $i \in I$ become sequentially available during an evaluation. The evaluation can be interrupted after reaching ϕ_1, ϕ_2 up until the maximal fidelity $\phi_L = 1$.

Such a blackbox can easily be built for a stochastic problem where the fidelity controls the number of Monte-Carlo (MC) draws using Sample Average Approximation (SAA) [15]. When one queries a stochastic blackbox function $f(x)$ at a point $x \in X$ for a MC draw, a noisy value $f_\xi(x)$ is obtained, where ξ is a random variable. A set of L increasing amounts of MC draws $\{\eta_i\}_{i \in I}$ that correspond to L fidelities is determined, where η_L is considered as the truth and $\phi_i = \frac{\eta_i}{\eta_L} \in [0, 1]$ for each $i \in I$. When using SAA, a predefined set of noise observations $\{\xi_\omega\}_{\omega \in \{1, 2, \dots, \eta_L\}}$ randomly sampled from the distribution of ξ is considered. This allows the definition of a deterministic multi-fidelity value for each $x \in X$ and each $i \in I$.

$$f(x, \phi_i) = \begin{cases} \frac{1}{\eta_1} \sum_{\omega=1}^{\eta_1} f_{\xi_\omega}(x) & \text{if } i = 1, \\ \frac{1}{\eta_i} \left(\eta_{i-1} f(x, \phi_{i-1}) + \sum_{\omega=\eta_{i-1}+1}^{\eta_i} f_{\xi_\omega}(x) \right) & \text{otherwise.} \end{cases} \quad (1)$$

The same process applies to $c_j(x, \phi)$ for each $j \in J$. A non-deterministic framework is presented in [2], where the observations ω are not predefined. When this new multi-fidelity blackbox is called at a trial point $x \in X$, the first η_1 draws are performed and $f(x, \phi_1)$ and $c(x, \phi_1)$ are made available. If the evaluation is not interrupted, only the $\eta_2 - \eta_1$ next draws are performed, and the outputs for ϕ_2 are computed using $f(x, \phi_1)$, $c(x, \phi_1)$ and these new draws. Then, to continue the evaluation, the next $\eta_3 - \eta_2$ draws are performed, and so on until ϕ_L is reached. This behaviour corresponds to [Definition 1](#).

[Figure 1](#) illustrates a classification of some known problems with respect to [Definition 1](#) and the applicability of SAA, as well as the classes for which the IDS and DIDS algorithms are applicable. PRIAD is an electrical equipment maintenance optimization problem, and solar is a concentrated solar power plant design optimization problem.

- **Multi-fidelity problems.** In finite element analysis, fidelity usually controls the coarseness of the mesh [20]. The `solar2` problem [6] is deterministic when its seed and number of MC draws are predetermined. Then, the fidelity controls the convergence criteria of numerical methods.
- **With intermediary outputs.** A neural network hyper-parameter optimization problem can be formulated such that the accuracy is returned after different number of epochs (different fidelities) during training [24]. Sequential blackboxes are problems where each output is given by a distinct blackbox which corresponds to a fidelity level, and they are each called sequentially [5].
- **Stochastic problems.** PRIAD's blackbox problem [16], `solar2` [6] and the cookie recipe optimization problem [31] are stochastic multi-fidelity problems where SAA may be applied. Additionally, this last problem is a laboratory experiment where intermediary outputs are easily accessible during a blackbox evaluation.




Figure 1: Classification of a few multi-fidelity problems with domains of application of IDS and DIDS.

1.1 Motivation

The primary motivation for this work is an asset management blackbox optimization problem encountered at Hydro-Québec as part of the PRIAD project [14, 16, 18, 23]. Its objective is to develop a periodic maintenance strategy optimization framework that relies on blackbox optimization methods. As the project is under development, the blackbox simulator involved is not yet available. Preliminary tests suggest that a single high fidelity evaluation of this simulator could take up to 145 days [23]. For an optimization of 2,000 evaluations, parallel computing can reduce the optimization time to one week if 45,000 CPUs are used [23]. To this effect, this work presents a new cost reduction method

based on the idea that the feasibility of some constraints may be estimated with low fidelities. When it is estimated that a point is infeasible from low fidelity information, high fidelity information is not computed to avoid high evaluation costs.

The second motivation of this work is to improve on the `Inter-DS` algorithm presented in [3]. `Inter-DS` served as a first step towards an algorithmic approach that exploits information from a broad range of available fidelities in the context of direct search methods for constrained blackbox optimization. The `IDS` algorithm of the present work is a strict upgrade from `Inter-DS` in a theoretical analysis sense. Many assumptions necessary to [3] are lifted for `IDS`, and in addition, some pathological cases are avoided without negatively impacting the algorithmic performance. In Figure 1, `Inter-DS` would lie in the same class as `IDS`. An important observation from [3] is that the sample point selection, required to solve an assignment subproblem, is crucial to the method. A second algorithm, `DIDS`, is proposed to periodically solve the subproblem as the optimization process proceeds. This new method is expected to perform better, but it requires the blackbox problem to have intermediary outputs, as described in Definition 1.

1.2 Contribution

Scientific literature on the subject of costly multi-fidelity blackbox problems predominantly studies the unconstrained case, or considers constraints along with a penalty in the objective function. Moreover, multi-fidelity frameworks sometimes allow for more than two fidelity levels, but a single high fidelity and a single low fidelity are almost always considered in benchmarks. Lastly, to the knowledge of the authors, all current multi-fidelity approaches use a low fidelity source to approximate gradients. This is usually achieved by fitting a model on the low fidelity data [26]. Conversely, this research approaches multi-fidelity problems under radically different lens:

- A finite and discrete set of fidelities is considered. The benchmarks in Section 5 use up to 11 fidelity levels.
- Direct-search methods are preferred to model-based approaches due to noisy, highly discontinuous and unpredictable blackbox problems which may fail to evaluate.
- Constraints are handled directly rather than being penalized in the objective function. As a matter of fact, the handling of constraints is the main focus of this work, and multi-fidelity information of the objective function is not used to improve the optimization process.

Future work will integrate multi-fidelity information of the objective function into the presented methods.

1.3 Organization

The document is structured as follows. Section 2 contains a short literature review on interruptible methods in blackbox optimization and on the use of multi-fidelity in the constrained case. Section 3 presents each of the new algorithmic components. They are assembled in Section 4 to create the two distinct optimization algorithms, `IDS` and `DIDS`. Finally, Section 5 compares `IDS` and `DIDS` to other methods on problems from the solar benchmarking collection [6]. On average, `DIDS` outperforms `IDS` which in turn outperforms `Inter-DS`. Concluding remarks follow in the final section.

2 Literature review

This work uses the KARQ (Known/hidden, A priori/simulated, Relaxable/unrelaxable, Quantifiable/nonquantifiable) taxonomy of constraints [25]. Notably, a constraint is said to be a priori if it has an explicit analytical formulation. When evaluating a point, a priori constraints are checked first at a negligible cost, and the blackbox process is only launched if they are satisfied. When counting

evaluations during an optimization, points where the blackbox process is not launched are not considered. This work also considers direct search methods for blackbox optimization, which are categorized into directional, mesh-based and line-search algorithms [17]. For these methods, the quantitative constraints $c_j(x) \leq 0, j \in J$ are often handled through the constraint violation function $h : \mathbb{R}^m \rightarrow \bar{\mathbb{R}}$ introduced in [8]:

$$h(x) := \begin{cases} \sum_{j=1}^m (\max\{c_j(x), 0\})^2 & \text{if } x \in X \\ \infty & \text{otherwise.} \end{cases}$$

The extreme barrier (EB) [10] algorithm is a two-phase method that first minimizes $h(x)$ in X for as long as the incumbent solution is infeasible. When it reaches feasibility, the second phase minimizes the extended-value function

$$f_\Omega(x) := \begin{cases} f(x) & \text{if } x \in \Omega \\ \infty & \text{otherwise} \end{cases}$$

on X from the feasible starting point. This approach is employed to reject infeasible points. The progressive barrier [8] is a more advanced constraint handling method. It consists of maintaining two incumbent solutions, a feasible and an infeasible one. The infeasible solution is progressively pushed towards the feasible domain by decreasing a threshold on $h(x)$, above which trial point $x \in X$ is rejected. The infeasible incumbent often has a better objective function value, which leads to the discovery of new feasible incumbent solutions.

The two phase interruptible EB [5] is a barrier method that applies to sequential blackboxes problems described above Figure 1. During an evaluation of trial point $x \in X$, the constraint violation function $h(x)$ is updated after each constraint value is computed. When this value is greater than $h(x^k)$ where x^k is the incumbent solution, the evaluation is interrupted. In the case where the ordering of the blackboxes can be chosen, the hierarchical satisfiability EB [5] suggests a sequence of optimization problems to solve with a similar interruption mechanism. These methods are built on the principle that costly evaluations should be interrupted when it is recognized through intermediate processes that a solution will not benefit the search strategy [30]. This principle is relevant to multi-fidelity blackbox problems, since fidelity levels correspond to such intermediate processes. The Inter-DS algorithm [3] is designed to leverage multi-fidelity information to quickly identify infeasible points and interrupt their evaluation. An important contribution introduced a trust region algorithm that incorporates multi-fidelity constraints and objective function into a sufficient decrease framework [27]. An in-depth review of the relevant blackbox optimization and multi-fidelity literature is presented in [3].

In the recent literature, a trust region method for stochastic problems where each level of fidelity corresponds to a different stochastic simulation subject to different numbers of MC draws is presented in [21]. Constrained stochastic problems are also tackled under a multi-fidelity framework by the Scout-Nd algorithm [1]. In the field of Bayesian optimization, the expected constrained improvement [19] was introduced as an extension to the expected improvement for constrained problems. It consists of multiplying the expected improvement with a probability of feasibility, and it was recently used in the context of multi-fidelity constrained Bayesian optimization [22, 29, 32]. Building on this work, new acquisition functions which are easier to implement by not solely relying on the expected constraint improvement were proposed [33].

3 Algorithmic components

The methods presented herein exploit multi-fidelity information from a blackbox problem with the primary goal of reducing evaluation costs during the optimization process. They are designed to work in conjunction with an existing direct search solver that does not natively support multi-fidelity evaluations. In order to efficiently leverage information across fidelity levels, it is essential that data from previous evaluations at various fidelities is accessible. Section 4 proposes the IDS and DIDS

algorithms, each applicable under different assumptions regarding the nature of this accessibility, and to different problems as classified in [Figure 1](#). These algorithms share several components, which are described individually in the present section, before introducing the complete frameworks.

[Section 3.1](#) introduces the *fidelity controller algorithm*, which performs the feasibility estimations. These estimations are guided by a key parameter of the fidelity controller: the *assignment vector*. Then, [Section 3.2](#) describes the assignment vector computation process by solving an optimization subproblem. [Section 3.3](#) presents theoretical guarantees of the subproblem.

3.1 The fidelity controller algorithm

The fidelity controller algorithm is a wrapper around the blackbox. Its purpose is to reduce evaluations costs by estimating the feasibility of a candidate using different fidelities with indices stored in I , for the case when the solver does not natively handle fidelity. Hence, rather than solving Problem $\mathcal{P}(1)$ directly, the wrapper problem $\bar{\mathcal{P}}$ is provided to the solver. Problem $\bar{\mathcal{P}}$ is an accelerated version of Problem $\mathcal{P}(1)$, with the intention of conserving feasibility. Through this document, for any optimization problem \mathcal{P} , *evaluating* \mathcal{P} is a shorthand used for *evaluating the blackbox from which Problem \mathcal{P} is derived*.

$$\bar{\mathcal{P}} \quad \min_{x \in X} \quad \bar{f}(x) \quad \text{s.t.} \quad \bar{c}_j(x) \leq 0, j \in J.$$

The functions $\bar{f}(x)$ and $\bar{c}(x)$ correspond to values of $f(x, \phi)$ and $c(x, \phi)$, where $\phi \in [0, 1]$ is determined by the fidelity controller.

Consider an assignment of the constraint $c_j \leq 0, j \in J$, to the fidelity level $\phi_i, i \in I$, with vector $a \in I^m$, where $a_j = i$ if constraint $c_j \leq 0$ is assigned to ϕ_i . An assignment of a constraint to a fidelity indicates the lowest fidelity for which it is estimated that the constraint's violation is accurately asserted. [Section 3.2](#) shows how this assignment vector a is computed. As I can be large, a subset of fidelities are selected, which depends on this vector a . This subset is denoted by $\Phi(a) \subseteq \{\phi_i\}_{i \in I}$, the set of fidelities of interest for the fidelity controller. Two distinct descriptions of this subset and of I are given in [Section 4.1](#) for IDS and [Section 4.2](#) for DIDS.

This assignment is at the core of the fidelity controller algorithm. Whenever $\bar{\mathcal{P}}$ is evaluated at a point $x \in X$ during the optimization, $\mathcal{P}(\phi)$ is sequentially evaluated at increasing fidelity levels $\phi \in \Phi(a)$ by the fidelity controller algorithm. After an evaluation at fidelity ϕ , only the constraints assigned to ϕ or a lower fidelity are checked. If any of those constraints is violated, the sequence of evaluations is interrupted, and the evaluated point is deemed infeasible. The goal is to identify infeasible points as cheaply as possible, and to stop investing computational costs into their evaluation. When all evaluations of $\mathcal{P}(\phi)$ are completed without interruption, the evaluated point is deemed feasible.

When an interruption occurs, the most recent outputs are returned to the solver. If the subset of fidelities $\Phi(a)$ does not contain 1, then it is possible that a point deemed feasible becomes the new incumbent solution, while being infeasible in reality. The assignment computation method proposed in [Section 3.2](#) is such that this event is rare, but it may occur. Since finishing the optimization with an infeasible point as the incumbent is highly undesirable, an additional evaluation at $\phi_L = 1$ is performed whenever this situation could arise. The complete process is described in [Algorithm 3.1](#) and schematized in [Figure 2](#).

This interruption mechanism is such that the solver may unknowingly receive blackbox output values from various fidelities. Nevertheless, outputs for points deemed feasible are only returned to the solver at the highest fidelity level in $\Phi(a)$. It may help solver performance to use the extreme barrier function f_Ω , to reject points deemed infeasible, which are more susceptible to have outputs returned at low fidelity. Note that the use of the extreme barrier with the fidelity controller algorithm allows for the use of some direct search unconstrained optimization solvers.

Algorithm 3.1: Fidelity controller algorithm (Blackbox wrapper problem $\bar{\mathcal{P}}$).

Inputs: trial point $x \in X$; assignment vector $a \in I^m$; incumbent value $f^* \in \bar{\mathbb{R}}$
Outputs: $\bar{f}(x) \in \bar{\mathbb{R}}$; $\bar{c}(x) \in \bar{\mathbb{R}}^m$

For each fidelity level $\phi \in \Phi(a)$ in increasing order

```

  | Evaluate  $f(x, \phi)$  and  $c(x, \phi)$ 
  | If there exists a  $j \in J$  such that  $\phi \geq \phi_{a_j}$  and  $c_j(x, \phi) > 0$ 
  |   | Return  $\bar{f}(x) = f(x, \phi)$ ,  $\bar{c}(x) = c(x, \phi)$ 
  | If  $\phi < 1$  and  $f(x, \phi) < f^*$ 
  |   | Evaluate  $f(x, 1)$  and  $c(x, 1)$ 
  | Return  $\bar{f}(x) = f(x, 1)$ ,  $\bar{c}(x) = c(x, 1)$ 
```

The fidelity controller algorithm is presented with the intention of reducing computational costs by causing interruptions on infeasible points while avoiding costly evaluations at $\phi_L = 1$. Constraints should be assigned to minimal fidelities, while being assigned to fidelities sufficiently high to perform feasibility estimations. Moreover, it may occur that estimations are incorrect, which can have an impact on the solver's performance. For these reasons, choosing a good assignment is not a trivial task.




Figure 2: Flow chart diagram of the fidelity controller algorithm. The for loop is completed when ϕ is greater or equal to the greatest element of $\Phi(a)$, where $a \in I^m$ is the assignment vector.

3.2 The assignment computation

Before using the fidelity controller algorithm from Section 3.1, the assignment vector $a \in I^m$ must be computed. This section proposes a method to compute this assignment that minimizes the computa-

tional cost of an evaluation of $\bar{\mathcal{P}}$, subject to constraints on the probability that a feasibility estimation is erroneous. Consider a finite set of sample points H . The indicator function

$$\mathbb{1}(c_j(x, \phi) > 0) := \begin{cases} 1 & \text{if } c_j(x, \phi) > 0 \\ 0 & \text{otherwise} \end{cases} \quad \forall j \in J$$

is used to define *representativity*, a useful concept to determine the lowest fidelity at which the feasibility of a constraint can be correctly identified.

Definition 2. Fidelity $\phi \in [0, 1]$ is said to be representative for a constraint $c_j(x, \phi) \leq 0$, $j \in J$ at point $x \in X$ if

$$\mathbb{1}(c_j(x, \phi_i) > 0) = \mathbb{1}(c_j(x, 1) > 0) \quad \forall i \in I \text{ where } \phi_i \geq \phi.$$

For example, if the constraint values of c_1 for $L = 4$ fidelities are $(c_1(x, \phi_i))_{i \in I} = (-0.2, 10, 0, -3)$, then ϕ_3 and ϕ_4 are the only two representative fidelities for constraint $c_1 \leq 0$ at point $x \in X$.

To compute the assignment vector a , statistical estimations are performed using the sample set $H \subset X$. Define X^{ap} as the subset of points from X where all a priori constraints are satisfied, and $H^{\text{ap}} = H \cap X^{\text{ap}}$. When $H \cap \Omega = \emptyset$, all constraints are assigned to $\phi_L = 1$, i.e., $a_j = L$ for each $j \in J$. In that case, $\bar{\mathcal{P}}$ is identical to $\mathcal{P}(1)$. If no information about feasible points is available, then multi-fidelity information as a whole should not be exploited. This allows DIDS and IDS to avoid cases that are pathological for Inter-DS. Otherwise, given Algorithm 3.1's interruption mechanism, the vector $a \in I^m$ that minimizes the expected computational cost of an evaluation of $\bar{\mathcal{P}}$ is found by solving the assignment Subproblem \mathcal{Q} .

$$\mathcal{Q} \quad \min_{a \in I^m} \quad f_{\mathcal{Q}}(a) = \sum_{i: \phi_i \in \Phi(a)} \left(\lambda_i \prod_{j: a_j < i} p_{a_j j} \right) \quad (2a)$$

$$\text{s.t.} \quad a_j \geq i(j) \quad \forall j \in J \quad (2b)$$

where, each $p_{a_j j}$, $j \in J$ is assumed independent, and for each $i \in I$ and $j \in J$,

$$\lambda_i := \frac{1}{|H^{\text{ap}}|} \sum_{x \in H^{\text{ap}}} \lambda(x, \phi_i) \approx \mathbb{E}_{x \in X^{\text{ap}}} [\lambda(x, \phi_i)], \quad (3)$$

$$p_{ij} := \frac{1}{|H^{\text{ap}}|} |\{x \in H^{\text{ap}} : c_j(x, \phi_i) \leq 0\}| \approx \mathbb{P}[c_j(x, \phi_i) \leq 0, x \in X^{\text{ap}}], \quad (4)$$

$$r_{ij} := \frac{1}{|H \cap \Omega|} |\{x \in H \cap \Omega : \phi_i \text{ is representative for } c_j \leq 0\}| \quad (5)$$

$$\approx \mathbb{P}[\text{fidelity } \phi_i \text{ is representative for constraint } c_j \leq 0, x \in \Omega],$$

$$i(j) := \min\{i \in I : r_{ij} = 1\}. \quad (6)$$

In \mathcal{Q} , λ_i is an estimation of the expected computational cost of evaluating $\mathcal{P}(\phi_i)$, p_{ij} is an estimation of the probability that $c_j \leq 0$ is satisfied at ϕ_i , r_{ij} is an estimation of the probability that ϕ_i is representative for $c_j \leq 0$, and $i : J \rightarrow I$ is a function that returns the lowest fidelity index $i \in I$ for which ϕ_i is representative for $c_j \leq 0$ for all feasible points in H . In Equation (2a), $f_{\mathcal{Q}}(a)$ expresses an expected evaluation cost of $\bar{\mathcal{P}}$, which depends on the evaluated fidelities. As an illustrative example, suppose $L = m = 4$, $a = [1, 2, 2, 4]$ and $\Phi(a) = [\phi_1, \phi_2, \phi_4]$. Then, $f_{\mathcal{Q}}([1, 2, 2, 4]) = \lambda_1 + \lambda_2 p_{11} + \lambda_4 p_{11} p_{22} p_{23}$. To evaluate $x \in H$, the evaluation at ϕ_1 of cost λ_1 always occurs, then the evaluation at ϕ_2 of cost λ_2 occurs if there was no interruption when evaluation at ϕ_1 , i.e., if $c_1(x, \phi_1) \leq 0$, which happens with probability p_{11} , and so on for ϕ_4 .

Minimal expected evaluation costs are obtained by assigning constraints to low fidelities. This is balanced with Equation (2b) to ensure constraints are assigned to fidelities with maximal estimated

representativity, therefore avoiding erroneous feasibility estimations. Let us then study the effects of such errors on the optimization process.

- **An infeasible point $x \in X \setminus \Omega$ is deemed feasible.** This only occurs if 1 is not an element of $\Phi(a)$, all constraints are satisfied at all fidelities in $\Phi(a)$, and $\bar{f}(x) \geq f^*$. Hence, an infeasible point can not become an incumbent solution. The only possible harm an infeasible point deemed feasible can cause occurs when a solver mistakenly uses it to compute search directions.
- **A feasible point $x \in \Omega$ is deemed infeasible.** This can occur after any evaluation of $\mathcal{P}(\phi)$ where $\phi < 1$. The main issue this can cause is when x would, if correctly identified as feasible, become the new incumbent solution. Generally, misidentifying feasible points can cause the solver to omit them to compute search directions. This is particularly harmful if an omitted point has a good objective function value, missing an opportunity to point the solver towards a minimum.

The drawbacks are more important and more frequent in the second case compared to the first. To mitigate this, [Equation \(2b\)](#) imposes that a blackbox constraint can only be assigned to a fidelity that is representative for all feasible points in H .

3.3 Theoretical analysis of the assignment problem

Subproblem \mathcal{Q} is constructed to compute an assignment vector from a sample set $H \subset X$ that ensures some results when $\bar{\mathcal{P}}$ evaluates $x \in H$. These results are shown here.

Lemma 1. When $H \cap \Omega \neq \emptyset$, for each $j \in J$, $\{r_{ij}\}_{i \in I}$ is monotone increasing with respect to i .

Proof. For each $i^1, i^2 \in I$ where $i^1 < i^2$, and for each $j \in J$,

$$R_{ij} := \{x \in H \cap \Omega : \mathbb{1}(c_j(x, \phi_\ell) > 0) = \mathbb{1}(c_j(x, 1) > 0) \forall \ell \in I, \phi_\ell \geq \phi_i\} \quad \forall i \in I$$

$$\Rightarrow r_{i^1 j} = \frac{|R_{i^1 j}|}{|H \cap \Omega|} \quad \text{and} \quad r_{i^2 j} = \frac{|R_{i^2 j}|}{|H \cap \Omega|}.$$

Because $\{\phi_i\}_{i \in I}$ is strictly increasing with respect to i , $\{\phi_\ell : \ell \in I, \phi_\ell \geq \phi_{i^1}\} \subset \{\phi_\ell : \ell \in I, \phi_\ell \geq \phi_{i^2}\}$, and therefore $R_{i^1 j} \subseteq R_{i^2 j}$. As a result, $|R_{i^1 j}| \leq |R_{i^2 j}|$ and $r_{i^1 j} \leq r_{i^2 j}$. \square

Theorem 2. For a given $H \subset X$, when the fidelity controller performs an evaluation of $x \in H \cap \Omega$ with an assignment vector given by Subproblem \mathcal{Q} , feasibility is conserved between $\bar{\mathcal{P}}$ and $\mathcal{P}(1)$.

Proof. If $H \cap \Omega = \emptyset$, then $\bar{\mathcal{P}}$ is identical to $\mathcal{P}(1)$. Otherwise, [Equation \(2b\)](#) ensures that each constraint $c_j \leq 0$ is assigned to a fidelity ϕ_i where $r_{ij} = 1$, and [Lemma 1](#) indicates that $r_{\ell j} = 1$ for each $\ell \geq i, \ell \in I$. A representativity of 1 means that the violation of a constraint is correctly identified at points in $H \cap \Omega$. This implies that if a constraint is satisfied at the fidelity it is assigned to for a point $x \in H^{\text{ap}}$, then it is also satisfied at any higher fidelity, including the truth. \square

Moreover, if $x \in H \cap \Omega$ and $1 \in \Phi(a)$, $\bar{c}(x) = c(x, 1)$, and $\bar{f}(x) = f(x, 1)$.

Assumption 1. In [Algorithm 3.1](#), the evaluation of $\mathcal{P}(1)$ that is conditional to $1 \notin \Phi(a)$ and $\bar{f} < f^*$ has no impact on the expected cost of evaluating $\bar{\mathcal{P}}$.

This evaluation ensures that only feasible points can become incumbent solutions in the solver. The assumption is false, but since the moments where a solution could become a new incumbent are unpredictable, it is necessary for the following theorem.

Theorem 3. Under [Assumption 1](#), for a given $H \subset X$ where $H \cap \Omega \neq \emptyset$ and a given assignment vector $a \in I^m$, the objective function $f_Q(a)$ of Subproblem \mathcal{Q} expresses the expected cost of evaluating $\bar{\mathcal{P}}$ at a point $x \in H^{\text{ap}}$.

Proof. For any $H \subset X$ where $H \cap \Omega \neq \emptyset$ and any $a \in I^m$,

$$\begin{aligned} \mathbb{E}_{x \in H^{\text{ap}}} [\text{cost of evaluation of } \bar{\mathcal{P}} \text{ at } x] &= \sum_{i=1}^L \lambda_i \mathbb{P}[\text{evaluation of } \mathcal{P}(\phi_i) \text{ at } x \text{ occurs for an } x \in H^{\text{ap}}] \\ &= \sum_{i: \phi_i \in \Phi(a)} \lambda_i \prod_{\ell: \phi_\ell \in \Phi(a), \ell < i} \mathbb{P}[\text{no interruption occurs at } (x, \phi_\ell) \text{ for an } x \in H^{\text{ap}}]. \end{aligned}$$

This last result holds because the probability that an evaluation of $\mathcal{P}(\phi)$ occurs is 0 if $\phi \notin \Phi(a)$, and otherwise, it is the probability that no interruption happens beforehand.

According to Theorem 2, considering the fidelity at which a constraint is assigned is sufficient to verify if this constraint would cause an interruption at higher fidelities. Therefore, assuming that all $p_{a_j j}$ for each $j \in J$ are independent and that the product of the elements of an empty set has value 1,

$$\begin{aligned} \mathbb{P}[\text{no interruption occurs at } (x, \phi_\ell) \text{ for an } x \in H^{\text{ap}}] &= \prod_{j: a_j = \ell} p_{\ell j} \quad \forall \ell \in I \\ \implies \mathbb{E}_{x \in H^{\text{ap}}} [\text{cost of evaluation of } \bar{\mathcal{P}} \text{ at } x] &= \sum_{i: \phi_i \in \Phi(a)} \lambda_i \prod_{\ell: \phi_\ell \in \Phi(a), \ell < i} \left(\prod_{j: a_j = \ell} p_{\ell j} \right) \\ &= \sum_{i: \phi_i \in \Phi(a)} \lambda_i \prod_{j: a_j < i} p_{a_j j} = f_Q(a). \quad \square \end{aligned}$$

A priori constraints are ignored in Theorem 3 because their violation causes the cost of an evaluation to be virtually null. For a new candidate $x \in X \setminus H$ that shares a similar behaviour with the points in H , $f_Q(a)$ is an approximation of the expected cost of $\bar{\mathcal{P}}$, and this problem approximates the feasibility of x . As noted in [3], richer sample sets H yield significantly better results.

4 Two interruptible direct search algorithms

For the approximations based on H to be effective, information about many evaluated points at many fidelity levels must be available. In this section, two optimization algorithms are proposed, each applicable under different assumptions regarding this availability. The first is the Interruptible Direct Search (IDS) algorithm, which is presented in Section 4.1. The second is its dynamic counter-part, the Dynamic Interruptible Direct Search (DIDS) algorithm, which is presented in Section 4.2. DIDS relies on a stronger assumption, in order to exploit more fidelities and reach greater computational cost reductions than IDS.

4.1 Interruptible Direct Search (IDS)

The IDS algorithm is applicable under Assumption 2a:

Assumption 2a. A set of sample points $H \subset X$, each evaluated at multiple fidelities with indices forming the set I , is provided with Problem $\mathcal{P}(\phi)$. An evaluation of $\mathcal{P}(\phi)$ only returns information for this fidelity level ϕ .

A set H may be the result of Latin Hypercube Sampling (LHS) or previous experiments. Using LHS ensures an unbiased sample set, but it may represent a high computational cost. IDS is applicable to any problem illustrated in Figure 1 where Assumption 2a is verified. The best solution from H is chosen as the initial optimization point. Whenever the fidelity controller calls $\mathcal{P}(\phi_i)$ at trial point $x \in X$ for some $i \in I$ during the optimization, it only receives the values $f(x, \phi_i)$ and $c(x, \phi_i)$ at cost λ_i . Fidelities

with no assigned constraints are not used by IDS to avoid their cost. Thus, for a given assignment a , the set $\Phi^{\text{IDS}}(a)$ is given by

$$\Phi^{\text{IDS}}(a) = \{\phi_i : i \in a\}. \quad (7)$$

The definition of $\Phi^{\text{IDS}}(a)$ implies that for a single evaluation of $\bar{\mathcal{P}}$, the fidelity controller could perform calls to $\mathcal{P}(\phi)$ at several fidelities, with total cost exceeding the cost of evaluating $\mathcal{P}(1)$. Conversely, using various fidelities may be relevant to interrupt evaluations on infeasible points. The optimal solution of Subproblem \mathcal{Q} is the assignment vector a that finds the best compromise.

Although a differentiable formulation of Subproblem \mathcal{Q} exists [3], it remains a mixed-integer problem with a polynomial objective function. To help solving, Algorithm 4.1 is proposed to add constraints to \mathcal{Q} that cut down the size of the set of feasible solutions, without removing an optimal solution. Then, an exhaustive search on this feasible set is sufficient in practice to solve the optimal assignment problem. This cutting algorithm introduces the set $J_\nu \subseteq J$, given by (8). It is the set of constraint indices that are not significantly affected by multi-fidelity. These constraints may have values that change with fidelity levels, but never such that their feasibility is affected.

$$J_\nu = \{j \in J : \forall i \in I, p_{ij} = p_{Lj} \text{ and } r_{ij} = 1\}. \quad (8)$$

Algorithm 4.1: Cutting algorithm.

Inputs: Subproblem \mathcal{Q} ; subset of constraint indices J_ν

Outputs: Updated Subproblem \mathcal{Q}

1. *Blackbox constraints related cuts*

If $J = J_\nu$, **Add** constraints “ $a_j = 1, \forall j \in J$ ” to Subproblem \mathcal{Q}

Else if $p_{Lj} < 1$ for some $j \in J_\nu$

Add constraints “ $a_j \leq a_j, \forall (j, j) \in J_\nu \times J$ ” to Subproblem \mathcal{Q}

Else if J_ν is not empty

Remove some decision variables of Subproblem \mathcal{Q} with $a \leftarrow \{a_j \in a : j \notin J_\nu\}$

2. *Fidelity related cuts*

For each fidelity index $i \in I$

If there exists a fidelity index $\ell \in I, \ell > i$, such that $\lambda_\ell \leq \lambda_i$

Add constraints “ $a_j \neq i, \forall j \in J$ ” to Subproblem \mathcal{Q}

If $i \notin \bigcup_{j \in J} i(j)$

Add constraints “ $a_j \neq i, \forall j \in J$ ” to Subproblem \mathcal{Q}

The first category of cuts from Algorithm 4.1 follows from the observation that the feasibility of constraints of indices in J_ν can be asserted at any fidelity $\phi_i, i \in I$. If all constraints are of this nature, they are simply assigned to the lowest fidelity, ϕ_1 . Else, if $p_{Lj} < 1$ for some $j \in J_\nu$, constraints of indices in J_ν are certainly assigned to the lowest fidelity in $\Phi^{\text{IDS}}(a)$. Finally, if all constraints $c_j \leq 0, j \in J_\nu$ are such that $p_{Lj} = 1$, they will never cause an interruption according to the parameters of Subproblem \mathcal{Q} . As a result, in an optimal solution, these constraints can be assigned to any fidelity where a constraint $c_j \leq 0, j \notin J_\nu$ is also assigned, therefore not causing a new call to $\mathcal{P}(\phi)$. Hence, Subproblem \mathcal{Q} is solved without assigning them, and they are assigned to the lowest fidelity where another constraint is assigned afterwards, as shown later in Algorithm 4.2.

The second category of cuts are related to fidelity levels. The first cuts have the effect of discarding fidelities that are not cheaper than higher fidelity levels. This causes the remaining λ_i values to be strictly increasing with respect to $i \in I$. Then, the last cuts follow from Theorem 5.

In the optimizations results presented in Section 5, with the cutting algorithm, Subproblem \mathcal{Q} is solved in at most one second.

Lemma 4. For any sample $H \subset X$ and for each constraint index $j \in J$, $p_{ij} = p_{Lj}$ for each fidelity index $i \in I$ where $i \geq i(j)$.

Proof. For any sample $H \subset X$ and for each constraint index $j \in J$, Equation (6) and Lemma 1 imply $r_{ij} = 1$ for each fidelity index $i \in I$ where $i \geq i(j)$. For those indices, for each $x \in H$, $\mathbb{1}(c_j(x, \phi_i) > 0) = \mathbb{1}(c_j(x, 1) > 0)$, meaning that the proportion of points where constraint $c_j(x, \phi_i) \leq 0$ is violated is the same across all fidelity indices $i \in I$ where $i \geq i(j)$. \square

Theorem 5. In the particular case of Subproblem \mathcal{Q} where $\Phi^{\text{IDS}}(a)$ is given by (7), there exists an optimal solution such that all blackbox constraints are assigned to fidelity of indices in $\bigcup_{j \in J} i(j)$.

Proof. Two cases are considered. First, if there exists a constraint index $j \in J$ such that $p_{i(j)j} = 0$, then constraint $c_j \leq 0$ is always violated for a feasible solution $a \in I^m$. Consequently, a is optimal if this constraint is assigned to $\phi_{i(j)}$, the lowest fidelity satisfying Equation (2b).

In the second case, let a be a feasible solution where there exists a fidelity index $\hat{i} \in I \setminus \bigcup_{j \in J} i(j)$ such that at least one constraint is assigned to this fidelity. If a does not exist, then each feasible solution is such that all constraints are assigned to fidelity indices in $\bigcup_{j \in J} i(j)$. Otherwise, let \hat{a} be a solution identical to a , except that all constraints assigned to ϕ_i are assigned to ϕ_{i-1} instead, as defined by

$$\hat{a}_j = \begin{cases} \hat{i} - 1 & \text{if } a_j = \hat{i} \\ a_j & \text{otherwise} \end{cases} \quad \forall j \in J. \quad (9)$$

Solution \hat{a} is feasible because for each $j \in J$, if $a_j = \hat{i}$ and a is feasible, then $\hat{i} - 1 \geq i(j)$. Define $P_i(a) = \prod_{j: a_j < i} p_{a_j j}$ for each $i \in I$. Lemma 4 implies that $P_i(\hat{a}) = P_i(a)$ for each $i \in I \setminus \{\hat{i}\}$. Additionally, definition (9) implies that $\Phi^{\text{IDS}}(\hat{a}) \setminus \{\hat{i} - 1, \hat{i}\} = \Phi^{\text{IDS}}(a) \setminus \{\hat{i} - 1, \hat{i}\}$, meaning that except for indices $\hat{i} - 1$ and \hat{i} , the terms of the sum in the objective function (2a) are identical for a and \hat{a} . Moreover, if no constraint is assigned to $\hat{i} - 1$ in a , $P_{\hat{i}-1}(\hat{a}) = P_{\hat{i}}(a)$. Otherwise, $P_{\hat{i}-1}(\hat{a}) = P_{\hat{i}-1}(a)$ and the summation terms for $i = \hat{i} - 1$ are also identical. As a result, with the indicator function $\mathbb{1}(\hat{i} - 1 \notin a)$ returning 1 if $\hat{i} - 1 \notin a$ and 0 otherwise,

$$\begin{aligned} f_{\mathcal{Q}}(a) - f_{\mathcal{Q}}(\hat{a}) &= \lambda_{\hat{i}} P_{\hat{i}}(a) - \lambda_{\hat{i}-1} P_{\hat{i}-1}(\hat{a}) \mathbb{1}(\hat{i} - 1 \notin a) = (\lambda_{\hat{i}} - \lambda_{\hat{i}-1} \mathbb{1}(\hat{i} - 1 \notin a)) P_{\hat{i}}(a) \\ &> 0 \implies a \text{ is not optimal.} \end{aligned}$$

This inequality holds because $P_{\hat{i}}(a) = 0$ only happens in the first case, and λ_i is strictly increasing with $i \in I$ thanks to Algorithm 4.1. The contrapositive of this result is that a solution a is optimal if there exists no fidelity index $\hat{i} \in I \setminus \bigcup_{j \in J} i(j)$ such that at least one constraint is assigned to this fidelity. \square

If there is no constraint $c_j \leq 0, j \in J$, such that $i(j) = L$, then the highest fidelity is not an element of $\Phi^{\text{IDS}}(a)$, meaning that evaluations of the truth by the fidelity controller only happen for potential future incumbent solutions. This allows IDS to save computational costs. However, it may be preferable to ensure that $1 \in \Phi^{\text{IDS}}(a)$, such that points deemed feasible are always truly feasible. Indeed, if $1 \in \Phi^{\text{IDS}}(a)$, the true values $f(x, 1)$ and $c(x, 1)$ are necessarily returned by $\bar{\mathcal{P}}$ when there is no interruption for point $x \in X$. To this effect, IDS posses the `include_truth` boolean parameter. When it is true, the objective function (2a) of the assignment Subproblem \mathcal{Q} is replaced with

$$\lambda_L + \sum_{i: \phi_i \in \Phi^{\text{IDS}}(a), i \neq L} \lambda_i \prod_{j: a_j < i} p_{a_j j} \quad (10)$$

to account for the mandatory evaluation using $\phi_L = 1$. The IDS algorithm is shown in Algorithm 4.2.

Algorithm 4.2: Interruptible Direct Search (IDS).**Inputs:**

$\mathcal{P}(\phi)$: optimization problem defined by f, c and $X \subseteq \mathbb{R}^n$;
 $H \subset X$: finite sample set;
solver : direct search blackbox optimization solver;
include_truth : boolean parameter that imposes $1 \in \Phi^{\text{IDS}}(a)$ when true. False by default.

Output: best solution found in X **1. Optimal assignment vector computation**

Initialize $f^* \leftarrow \infty$, the incumbent value
If $H \cap \Omega$ is empty, $a_j = L, \forall j \in J$ (equivalent to solving $\mathcal{P}(1)$ without interruptions)
Else
From H , (5), (4) and (3), evaluate the parameters of Subproblem \mathcal{Q}
If **include_truth** is true, use (10) as objective function for Subproblem \mathcal{Q}
Compute J_ν using (8), define Subproblem \mathcal{Q} and apply Algorithm 4.1
Find a by performing an exhaustive search on the feasible set of Subproblem \mathcal{Q}
If constraints of index in J_ν are unassigned, $a_j \leftarrow \min_{j \in J \setminus J_\nu} \{a_j\}$ for each $j \in J_\nu$

2. Direct search optimization

Launch solver on \mathcal{P} with the blackbox provided by Algorithm 3.1 (parametrized by a and f^*)
- Use the best point of H as the initial point
- Update f^* after each new incumbent
- **If** **include_truth** is true, use $\Phi^{\text{IDS}}(a) \cup \{1\}$ instead of $\Phi^{\text{IDS}}(a)$

Return the solver output

As the incumbent value f^* varies during the optimization, Problem $\bar{\mathcal{P}}$ evaluated at the same point at different moments may return different values. As a result, $\bar{f}(x)$ and $\bar{c}(x)$ require f^* as a third parameter to be functions, but they are considered as such to simplify the expressions.

4.2 Dynamic Interruptible Direct Search (DIDS)

The DIDS algorithm is applicable under an assumption regarding the accessibility of multi-fidelity information that differs from IDS. The initial sample H is not necessary. Assumption 2a is replaced with Assumption 2b. It is also later shown that Assumption 1 is no longer necessary.

Assumption 2b. Problem $\mathcal{P}(\phi)$ has intermediary outputs for a set of multiple fidelity indices I , as given by Definition 1.

The fidelity controller algorithm is given a more specific interpretation for an evaluation using $\phi_i, i \in I \setminus \{1\}$. In the case of available intermediary outputs, an evaluation of $\mathcal{P}(\phi_i)$ at trial point $x \in X$ means that the blackbox continues its evaluation of x from ϕ_{i-1} until fidelity ϕ_i is reached. Figure 3 illustrates this case where Definition 1 is met with SAA. The cost λ_i represents the additional cost to reach ϕ_i from ϕ_{i-1} .

Contrary to IDS, evaluating many fidelities cannot be more costly than evaluating $\mathcal{P}(1)$. When evaluating $\bar{\mathcal{P}}$ at point $x \in X$, fidelities are reached sequentially in increasing order, until $\phi_L = 1$ is reached. Then, $c(x, \phi_i)$ and $f(x, \phi_i)$ are known for every fidelity $\phi_i, i \in I$. Consequently,

$$\Phi^{\text{DIDS}}(a) = \{\phi_i\}_{i \in I} \quad \implies \quad f_{\mathcal{Q}}(a) = \sum_{i=1}^L \left(\lambda_i \prod_{j: a_j < i} p_{a_j j} \right). \quad (11)$$

This implies that a point is only deemed feasible when it is truly feasible. As a result, Algorithm 3.1's safeguard against an infeasible becoming an incumbent is never run, and Assumption 1 is unnecessary.

Without Assumption 2a, no initial sample $H \subset X$ is required. Rather, evaluated points deemed feasible during the optimization are stored in a cache \mathcal{V} , along with a vector of length m containing




Figure 3: Flow chart diagram showing a particular case of the fidelity controller with the DIDS algorithm and a stochastic blackbox with SAA.

the lowest representative fidelity for each constraint for $i(j)$ computations. Then, periodically during the optimization process, points from \mathcal{V} are selected to form a sample set H from which a is updated. Before iteration k of solver, the smallest radius Δ such that $B_\Delta(x^k) \cap \mathcal{V} \cap \Omega$ contains $n + 1$ feasible points is computed. This ensures H is formed with points in close proximity to x^k , while its size is at least the size of an n -dimensional simplex. When this radius does not exist, every constraint is assigned to $\phi_L = 1$. The new optimal assignment a follows from the following theorem.

Theorem 6. The particular case of Subproblem \mathcal{Q} where the objective function is given by (11) has optimal solution a^* defined by

$$a_j^* = i(j). \quad (12)$$

Proof. Two cases are considered. First, if $p_{i(j)j} = 1$ for each $j \in J$, then no blackbox constraint is ever violated for any feasible solution $a \in I^m$. In that case, every feasible solution, notably a^* , is optimal. If there exists a constraint index $j \in J$ such that $p_{i(j)j} = 0$, then constraint $c_j \leq 0$ is always violated. Consequently, a solution a is optimal if this constraint is assigned to the lowest possible fidelity for a feasible solution, which is the case for a^* .

In the second case, let a be a feasible solution where there exists a constraint index $\hat{j} \in J$ such that $p_{i(\hat{j}),\hat{j}} < 1$ and $a_{\hat{j}} = \hat{i} > i(\hat{j})$. If such an a does not exist, then a^* is the only feasible solution, and it is therefore optimal. Let \hat{a} be a solution identical to a , except that constraint $c_{\hat{j}} \leq 0$ is assigned to $\phi_{i(\hat{j})}$ instead of $\phi_{\hat{j}}$, as defined by

$$\hat{a}_j = \begin{cases} i(\hat{j}) & \text{if } j = \hat{j} \\ a_j & \text{otherwise} \end{cases} \quad \forall j \in J.$$

Solution \hat{a} is feasible because a satisfies [Equation \(2b\)](#). Define $P_i(a) = \prod_{j:a_j < i} p_{a_j j}$ for each $i \in I$. [Definition \(12\)](#) and [Lemma 4](#) imply that

$$P_i(\hat{a}) = \begin{cases} P_i(a)p_{i(j),\hat{j}} & \text{if } i \in \hat{I} := \{i(j) + 1, i(j) + 2, \dots, \hat{i}\} \\ P_i(a) & \text{otherwise} \end{cases} \quad \forall i \in I.$$

As a result, except for indices in \hat{I} , the terms of the sum in the objective function [\(11\)](#) are identical for a and \hat{a} .

$$\begin{aligned} f_Q(a) - f_Q(\hat{a}) &= \sum_{i \in \hat{I}} \lambda_i P_i(a) - p_{i(j),\hat{j}} \sum_{i \in \hat{I}} \lambda_i P_i(a) = (1 - p_{i(j),\hat{j}}) \sum_{i \in \hat{I}} \lambda_i P_i(a) > 0 \\ \implies a &\text{ is not optimal.} \end{aligned}$$

The inequality holds because $p_{i(j),j} > 0$ for each $j \in J$, implying that the summation is non zero, and $p_{i(j),\hat{j}} < 1$. The contrapositive of this result is that a solution a is optimal if $a_j = i(j)$ for each $j \in J$ where $p_{i(j),\hat{j}} < 1$, which describes a^* . \square

As the assignment vector a varies during the optimization, Problem $\bar{\mathcal{P}}$ evaluated at the same point at different moments may return different values. As a result, $\bar{f}(x)$ and $\bar{c}(x)$ require a as a third parameter to be functions, but they are considered as such to simplify the expressions. The DIDS algorithm is shown in [Algorithm 4.3](#).

Algorithm 4.3: Dynamic Interruptible Direct Search (DIDS).

Inputs:

$\mathcal{P}(\phi)$:	problem containing $X \subseteq \mathbb{R}^n$, f and c ;
x^0 :	optimization starting point (if <code>solver</code> requires it);
<code>solver</code> :	direct search blackbox optimization solver

Output: best solution found in X

1. Initialization

Initialize $\mathcal{V} \leftarrow \emptyset$, the cache of evaluated points
Initialize $k \leftarrow 1$, the <code>solver</code> iteration counter
Initialize $f^* \leftarrow \infty$, the incumbent value
Initialize $a \leftarrow a^0$, the initial assignment vector in which all constraints are assigned to ϕ_L

2. Direct search optimization

Launch <code>solver</code> on $\bar{\mathcal{P}}$ with the blackbox provided by Algorithm 3.1 (parametrized by a and f^*)
- Use x^0 as the initial point
- After each evaluation of $\mathcal{P}(\phi)$, update \mathcal{V}

After iteration k of `solver` with incumbent solution x^k

$f^* \leftarrow f(x^k)$
Find the minimal radius Δ such that the ball $B_\Delta(x^k) \cap \mathcal{V} \cap \Omega$ has at least $n + 1$ points.
If Δ does not exist, $a \leftarrow a^0$

Else

$H \leftarrow B_\Delta(x^k) \cap \mathcal{V}$
From H , compute the representativity probability r_{ij} for each $(i, j) \in I \times J$ with (5)
$a \leftarrow [i(j)]_{j \in J}$, with (6) , as indicated by (12)
$k \leftarrow k + 1$

Return the `solver` output

The set H is updated dynamically such that it contains points in close proximity to x^k , while $H \cap \Omega$ is at least the size of a simplex.

5 Computational results

In this section, the proposed algorithms are compared on blackbox optimization benchmark problems. Since the blackbox simulator from the Hydro-Québec project motivating this study is still under development, benchmarks are conducted on the `solar` collection of test problems [6]. To the authors' knowledge, the `solar` family of benchmarks is the only one containing realistic industrial problems in which a fidelity parameter affects the constraint values. Moreover, selecting benchmark problems that reflect real engineering optimization challenges is crucial. This suite provides ten blackbox simulators of a concentrated solar power plant (CSP). Table 1 shows relevant characteristics for the three problems selected for this study: `solar2`, `solar3` and `solar4`. Problem `solar7` is the only other problem with multi-fidelity constraints. However, preliminary tests show that the optimal assignment vector computed by IDS and DIDS is almost always $a_j = L$ for each $j \in J$. This means that fidelities that allow relevant interruptions are rarely found, and both algorithms almost always simply solve $\mathcal{P}(1)$ directly. Problem `solar7` is not studied further.

Table 1: Number of variables and constraints for the three studied solar problems. In parentheses, the number of categorical variables and multi-fidelity constraints are indicated.

Problem	Number of variables			Number of constraints			m
	continuous	integer (categorical)	n	a priori	simulated (multi-fidelity)		
<code>solar2</code>	12	2 (0)	14	5	7 (4)		12
<code>solar3</code>	17	3 (1)	20	5	8 (5)		13
<code>solar4</code>	22	7 (1)	29	7	9 (6)		16

To conduct extensive benchmarks, the computational cost of using SAA with the selected `solar` problems showed to be prohibitive. Instead, the `solar` problems are used deterministically by setting the seed value to 0. For DIDS, intermediary outputs are made available by adding code to `solar`. To this effect, the fidelity parameter, described in [6], controls the convergence criteria of numerical methods. In PRIAD, evaluation times may be measured in days [23]. To emulate this context where solver computing times are negligible, only blackbox evaluation times are considered. Time dependant data profiles are used [11, 13, 28] for comparisons on numerous problem instances. For an algorithm, a data profile plots the portion of τ -solved problem instances with respect to time. A problem instance is said to be τ -solved at time T by algorithm A with initial point x^0 if

$$\frac{f(x^0) - f^T}{f(x^0) - f^*} \geq 1 - \tau,$$

where f^T is the best feasible objective function value found by algorithm A as of time T , f^* is the best feasible solution found among all algorithms on the problem instance, and $\tau \in (0, 1)$ is a tolerance.

The IDS and DIDS algorithms are paired with the NOMAD blackbox optimization solver, version 4.4.0 [12]. NOMAD is a freely available implementation of the MADS algorithm [7, 8] that has shown to be successful on real engineering and industrial problems [4]. It is among the best suited solvers for constrained, noisy and highly discontinuous problems where direct-search methods are preferred to model-based approaches [9]. The NOMAD parameters are left to their default values for all tested algorithms, except that the extreme barrier is used for all constraints to ensure the rejection of points for which the evaluation is interrupted by the fidelity controller. The fidelity indices used in this section are $I = \{1, 2, \dots, 11\}$ with corresponding fidelity parameters $\{\phi_i\}_{i \in I} = \{10^{-10}, 0.1, 0.2, 0.3, \dots, 1\}$. A unique feasible initial solution is provided for each `solar` problem.

IDS and DIDS are compared to two methods: a base case where NOMAD solves $\mathcal{P}(1)$ directly without the fidelity controller, and the Inter-DS algorithm from [3], also paired with NOMAD. For IDS

and **Inter-DS**, the set H is obtained via Latin Hypercube Sampling (LHS) with 1000 points for each problem, and the LHS bounds are selected close to the starting point, as recommended in [3]. For **DIDS**, the LHS set is not considered. Rather, intermediary outputs are available, as per Definition 1. As the LHS is included in the **Inter-DS** algorithm and **IDS** assumes the LHS results already exist, no LHS computation time is considered for fair comparison.

First, Section 5.1 shows benchmarks on **solar3**. Second, Section 5.2 presents results for **solar4**, a problem where constraint behaviour with respect to fidelity is highly correlated with the blackbox input. Last, Section 5.3 studies benchmarks on **solar2**, a problem where the impact of erroneous feasibility estimations by the fidelity controller defies expectations. All optimizations are performed using multiple Intel Xeon Gold 6150 CPUs operating at 2.70 GHz.

5.1 solar3 benchmarks

The optimization results of 20 **solar3** problem instances, obtained with different NOMAD seeds and a 16.3 hours budget, are shown in Figure 4 with two different values of τ . A first observation is that pairing NOMAD with **DIDS** or **IDS** with `include_truth` set to true necessarily improves its performance. At any time, with $\tau = 0.01$, **DIDS** τ -solves the most problem instances and **IDS-truth** τ -solves the second most. With tolerance $\tau = 0.2$, **IDS** performs the best with 100% of instances τ -solved. With $\tau = 0.01$, this percentage drops to 30%, which is significantly lower than the 70% of both **DIDS** and **IDS-truth**. **Inter-DS** presents the same data profile as the base case with $\tau = 0.2$, and τ -solves 5% less problem instances than the base case with the lower tolerance.




Figure 4: Data profiles on Problem **solar3** with 20 different problem instances.

To delve further into these results, Figure 5 illustrates the frequency at which evaluations throughout all 20 optimizations ended. In **solar3**, the constraint that has lead to the most interruptions for all algorithms throughout all optimizations is $c_2 \leq 0$, which verifies compliance to an energy demand. For this constraint, higher fidelities best indicate if the energy produced is sufficient. As a result, low fidelities are rarely used. **DIDS** leverages six different fidelity levels, resulting in the best performance with $\tau = 0.01$. **IDS** only uses $\phi_L = 1$ in 28.86% of its evaluations, while the other algorithms use it for at least 85.05% of their evaluations.

As discussed at the end of Section 3.2, **IDS** and **DIDS** are designed to ensure low probabilities that feasibility errors occur. Ideally, an interruption method performs more evaluations than the base case in the same time budget, and interruptions always occur on infeasible points. However, this can not be guaranteed without additional hypotheses, and despite low probabilities of errors, they are expected to occur during the optimization process (16.3 hours). Such errors sometimes cause the evaluation sequence to differ from the base case. To study this phenomenon, Table 2 displays, for each interruption algorithm, the average factor of number of base case evaluations, the percentage of problem instances

where the sequence of points up to the base case's last evaluation are different, and, for these instances, the average percentage of base case evaluations before the first different point in the sequence.




Figure 5: Occurrence (%) of last used fidelities in solar3 evaluations.

Table 2: Evaluation sequence comparison with the base case for solar3.

Value compared with the base case	Inter-DS	IDS	IDS-truth	DIDS
Average number of evaluations factor	1.001	1.006	1.003	1.013
% of problem instances where an x differs	0	100	0	0
Average % of evaluations before an x differs	-	2.26	-	-

Where all other algorithms followed the same evaluation sequence as the base case for all optimizations, IDS has differences for all 20 optimizations. These differences started after only 2.26% of the base case evaluations on average. This is the result of IDS's prevalent usage of fidelities inferior to $\phi_L = 1$. This way, IDS performs the most average number of evaluations at the cost of causing numerous feasibility errors, and these errors lead to worse search directions, resulting in the worst performance with $\tau = 0.01$. Additionally, IDS's average number evaluations is only 1.006 times that of the base case. Activating the `include_truth` parameter for IDS successfully prevents these feasibility errors. With $\tau = 0.01$, Inter-DS's profile is below the base case's at any time. This is because 98.97% of its evaluation cost $\lambda_1 + \lambda_5 + \lambda_{11}$. The use of fidelity $\phi_5 = 0.4$ is not sufficiently effective at causing interruptions to overcome the λ_5 cost, causing a constant delay in the optimization.

5.2 solar4 benchmarks

The optimization results of 20 solar4 problem instances, obtained with different NOMAD seeds and a 16.4 hours budget, are shown in Figure 6 with two different values of τ . For solar4, the assignment vector found by IDS contains $\phi_L = 1$, meaning that the `include_truth` parameter is redundant. At any time, the data profiles of Inter-DS, IDS and the base case are almost identical, even with low tolerance values such as 10^{-3} and 10^{-5} (which usually help differentiating the profiles). The only exceptions to this similitude are the last few minutes with tolerance $\tau = 10^{-5}$, where IDS suddenly reached 45% of τ -solved problem instances, compared to 25% for Inter-DS and the base case. Conversely, DIDS dominates both plots.

Further insights are found in Figure 7. Similarly to solar3, the constraint that has lead to the most interruptions for all algorithms throughout all optimizations is $c_2 \leq 0$, which verifies compliance to an energy demand and requires high fidelities to estimate. Inter-DS and IDS use $\phi_L = 1$ for 99.51% and 99.56% of their evaluations respectively. In that matter, they behave almost identically to the base case. The leveraging of intermediary outputs by DIDS allowed it to interrupt evaluations at various fidelities.

Table 3 describes some evaluation sequence statistics for these optimizations. Notice that Inter-DS achieves less evaluations than unpaired NOMAD on average. Overall, the ineffectiveness of the non-




Figure 6: Data profiles on Problem solar4 with 20 different problem instances.




Figure 7: Occurrence (%) of last used fidelities in solar4 evaluations.

dynamic interruption algorithms at leveraging multi-fidelity information is attributed to the fact that the behaviour of the constraints relative to fidelity is heavily dependant on the solution space for solar4 [3]. As such, the assignment vectors of Inter-DS and IDS can become ineffective, whereas the dynamic adaptations of DIDS ensure a high quality assignment vector during the whole optimization. This explains why DIDS is the only algorithm that leverages multiple fidelity levels other than 10^{-10} and 1 in Figure 7. This use of multi-fidelity for interruptions causes DIDS to not follow the base case's evaluation sequence for 90% of the problem instances, and the first difference occurs after only 6.61% of base case evaluations on average. Nonetheless, deeper analysis shows the different search directions result in similar solutions to those of the base case. Considering all factors, the reason for DIDS's highest portion of τ -solved instances is the 1.211 average factor in number of evaluations.

Table 3: Evaluation sequence comparison with the base case for solar4.

Value compared with the base case	Inter-DS	IDS	DIDS
Average number of evaluations factor	0.991	1.039	1.211
% of problem instances where an x differs	70	70	90
Average % of evaluations before an x differs	10.91	10.37	6.61

5.3 solar2 benchmarks

The optimization results of 20 solar2 problem instances, obtained with different NOMAD seeds and a 16.1 hours budget, are shown in Figure 8 with two different values of τ . All interruption methods except Inter-DS allow NOMAD to reach better or equal solutions than by directly solving $\mathcal{P}(1)$ at any time. The IDS algorithm dominates both plots by τ -solving 85% of instances with a tolerance of 0.1 and 60% of instances with a tolerance of 0.01. By decreasing the tolerance, DIDS and the base case

are the algorithms with the largest decrease of τ -solved instances, going from 70% to 10% for DIDS and from 50% to 0% for the base case.




Figure 8: Data profiles on Problem solar2 with 20 different problem instances.

To explain these results, Figure 9 illustrates how each algorithm used the available fidelities. It indicates that IDS interrupted evaluations at the lowest fidelity the most: 75.93% of evaluations interrupted at $\phi_1 = 10^{-10}$. Throughout all problem instances and all algorithms, every interruption at $\phi_1 = 10^{-10}$ was caused by $c_6 \leq 0$, and by it only. It ensures that the number of heliostats to place, x_6 , can fit in the field, for which the dimensions are computed from x_1 , x_2 and x_3 . This is not an a priori constraint, but its value is computed before any ray-tracing MC draw is realized. This explains why $\phi_1 = 10^{-10}$ is so heavily used by all interruption algorithms. The usage of all fidelity levels by DIDS is not sufficient to overcome the other interruption methods. This is not surprising, as the behaviour of the constraints relative to fidelity is fairly constant throughout the solution space of solar2 [3], meaning that dynamic adjustments to the assignment vector are almost irrelevant.




Figure 9: Occurrence (%) of last used fidelities in solar2 evaluations.

Table 4 displays how the different multi-fidelity approaches affected the optimizations. A direct correlation between the frequency of evaluations interrupted at $\phi_1 = 10^{-10}$ in Figure 9 and the number of evaluations in this table is observed. This is the first element explaining IDS's success: it performed 6.86 times more evaluations than the base case. The second is that IDS's evaluation sequence differed from the base case for all problem instances, and it does so the fastest: only the first 3.84% of base case evaluations are identical before the first difference on average. A deeper analysis revealed that the feasibility mistakes resulted in NOMAD exploring better solutions. This challenges the idea that interruption methods should be designed to minimize feasibility estimation errors. Most errors occurred on infeasible points $x \in X$ where $f(x, 1)$ is close to f^* and the true $h(x)$ value is close to 0.

The fact that this lead to better solutions is most likely the result of chance. Data profiles with a greater number of problem instances are required to confirm this observation.

Table 4: Evaluation sequence comparison with the base case for solar2.

Value compared with the base case	Inter-DS	IDS	IDS-truth	DIDS
Average number of evaluations factor	4.05	6.86	2.94	4.65
% of problem instances where an x differs	100	100	0	5
Average % of evaluations before an x differs	5	3.84	-	29

In contrast, DIDS differed from the base case in only 5% of instances, and for this 5%, the first 29% of base case evaluations are identical before the first difference on average. It also performed 4.65 times more evaluations than the base case on average. Yet, DIDS τ -solves only 10% of the problem instances with tolerance $\tau = 0.01$. Again, more tests would be needed to confirm these results. Activating the `include_truth` parameter in IDS decreases the number of optimizations with a different sequence from 100% to 0%, all while performing 2.94 times more evaluations than the base case on average. Similarly to solar3, this demonstrates that the parameter is efficient at reducing feasibility errors, but contrary to solar3, IDS performs better here without the parameter. This is because it reduces the number of evaluations factor from 6.86 to 2.94, but NOMAD's performance is still improved when using IDS with the parameter activated. In comparison, Inter-DS performed more evaluations, with a factor of 4.05, but also had numerous feasibility errors (all problem instances differed from the base case), resulting in a similar data profile with tolerance $\tau = 0.01$.

5.4 Remarks on the computational effort

Recall that solver computing times are omitted in the data profiles to create conditions similar to PRIAD's. This explains the irregular optimization times in Figures 4, 6 and 8, as they are computed a posteriori. Note that a total of 6333 hours (8.65 months, divided between multiple machines) of optimization time was required to generate these figures. To this effect, only pairings with NOMAD are tested. Numerical tests in [3] show that solvers implementing the Particle Swarm Optimization and the Nelder-Mead algorithms are successfully used with fidelity controller-based algorithms.

When Inter-DS was introduced in [3], it was noted that when H contains no feasible points, Inter-DS is extremely harmful to the optimization. These pathological cases are solved in IDS and DIDS by assigning all constraints to $\phi_L = 1$ when $H \cap \Omega = \emptyset$. Because the solution is trivial, no additional test is conducted.

6 Discussion

This work introduces two new algorithms for constrained multi-fidelity blackbox optimization problems based on interruptible evaluations. They both perform feasibility estimations at various fidelity levels, and interrupt the process when a point is deemed infeasible. The first, IDS, exploits a set of points evaluated prior to the optimization to construct its interruption mechanism. The second, DIDS, relies on intermediary information from previously evaluated points during the optimization. IDS is presented as a theoretical improvement of Inter-DS, which is also based on pre-optimization evaluations. The results show that for most problems, IDS also successfully improves on Inter-DS practically.

For the problems tested, IDS with the `include_truth` parameter activated and DIDS both strictly allow NOMAD to reach better solutions at the end of the same time budget. Without the `include_truth` parameter, IDS's performance is highly variable because it performs interruptions more aggressively. The solar3 Problem punishes this greediness, solar4 is indifferent to it, and solar2 rewards it. A shortcoming of this study is that it is still unclear whether feasibility estimation mistakes are desirable for solar2, or if more optimizations should compose the data profiles to expose different global trends.

Future work will address phenomenon, as well as study if a bias is introduced by successful interruptions at low fidelities, leading to a lack of high fidelity information for infeasible points, which is used for assignment vector updates.

An important takeaway is that the behaviour of the constraints relative to fidelity must be fairly constant throughout the solution space for **IDS** to perform well. Otherwise, only **DIDS** can leverage multi-fidelity information to an advantage.

The current literature in the field of multi-fidelity blackbox optimization mainly studies the unconstrained case or model-based optimization approaches. This research proposes methods to handle multi-fidelity constraints in the context of direct search methods. Later work will incorporate the aforementioned objective function focused literature with the presented methods to consider the generalized case where all available multi-fidelity information is leveraged.

References

- [1] A. Agrawal, K. Ravi, P.S. Koutsourelakis, and H.J. Bungartz. Stochastic black-box optimization using multi-fidelity score function estimator. *Machine Learning: Science and Technology*, 6(1):015024, 2025.
- [2] S. Alarie, C. Audet, P.-Y. Bouchet, and S. Le Digabel. Optimisation of stochastic blackboxes with adaptive precision. *SIAM Journal on Optimization*, 31(4):3127–3156, 2021.
- [3] S. Alarie, C. Audet, M. Diago, S. Le Digabel, and X. Lebeuf. Inter-DS: A cost saving algorithm for expensive constrained multi-fidelity blackbox optimization. *Computational Optimization and Applications*, 90(3):607–629, 2025.
- [4] S. Alarie, C. Audet, A.E. Gheribi, M. Kokkolaras, and S. Le Digabel. Two decades of blackbox optimization applications. *EURO Journal on Computational Optimization*, 9:100011, 2021.
- [5] S. Alarie, C. Audet, P. Jacquot, and S. Le Digabel. Hierarchically constrained blackbox optimization. *Operations Research Letters*, 50(5):446–451, 2022.
- [6] N. Andrés-Thiό, C. Audet, M. Diago, A.E. Gheribi, S. Le Digabel, X. Lebeuf, M. Lemyre Garneau, and C. Tribes. Solar: a solar thermal power plant simulator for blackbox optimization benchmarking. *Optimization and Engineering*, 26(3):1815–1861, 2025.
- [7] C. Audet and J.E. Dennis, Jr. Mesh Adaptive Direct Search Algorithms for Constrained Optimization. *SIAM Journal on Optimization*, 17(1):188–217, 2006.
- [8] C. Audet and J.E. Dennis, Jr. A Progressive Barrier for Derivative-Free Nonlinear Programming. *SIAM Journal on Optimization*, 20(1):445–472, 2009.
- [9] C. Audet, M. Diago, S. Le Digabel, X. Lebeuf, P. Brillion, and C. Tribes. Benchmarking derivative-free optimization solvers for CSP plant design using the SOLAR simulator. Technical Report G-2025-70, Les cahiers du GERAD, 2025.
- [10] C. Audet and W. Hare. Derivative-Free and Blackbox Optimization. Springer Series in Operations Research and Financial Engineering. Springer, Cham, Switzerland, 2017.
- [11] C. Audet, W. Hare, and C. Tribes. Benchmarking constrained, multi-objective and surrogate-assisted derivative-free optimization methods. Technical Report G-2025-36, Les cahiers du GERAD, 2025.
- [12] C. Audet, S. Le Digabel, V. Rochon Montplaisir, and C. Tribes. Algorithm 1027: NOMAD version 4: Non-linear optimization with the MADS algorithm. *ACM Transactions on Mathematical Software*, 48(3):35:1–35:22, 2022.
- [13] V. Beiranvand, W. Hare, and Y. Lucet. Best practices for comparing optimization algorithms. *Optimization and Engineering*, 18(4):815–848, 2017.
- [14] A. Côté, O. Blancke, S. Alarie, A. Dems, D. Komljenovic, and D. Messaoudi. Combining Historical Data and Domain Expert Knowledge Using Optimization to Model Electrical Equipment Reliability. In 2020 International Conference on Probabilistic Methods Applied to Power Systems (PMAPS), pages 1–6, Liege, Belgium, 2020.
- [15] T. Homem de Mello and G. Bayraksan. Monte Carlo sampling-based methods for stochastic optimization. *Surveys in Operations Research and Management Science*, 19(1):56–85, 2014.
- [16] M. Diago, X. Lebeuf, T.J.M. Meango, A. Côté, and C. Rajotte. Risk-Based Optimization of Periodic Maintenance for Power Grid Equipment. In 2025 Annual Reliability and Maintainability Symposium (RAMS), pages 1–6. IEEE, 2025.

- [17] K.J. Dzahini, F. Rinaldi, C.W. Royer, and D. Zeffiro. Direct-search methods in the year 2025: Theoretical guarantees and algorithmic paradigms. *EURO Journal on Computational Optimization*, 13:100110, 2025.
- [18] M. Gaha, B. Chabane, D. Komljenovic, A. Côté, C. Hébert, O. Blancke, A. Deleva, and G. Abdul-Nour. Global Methodology for Electrical Utilities Maintenance Assessment Based on Risk-Informed Decision Making. *Sustainability*, 13(16):9091, 2021.
- [19] J.R. Gardner, M.J. Kusner, Z. Xu, K.Q. Weinberger, and J.P. Cunningham. Bayesian optimization with inequality constraints. Technical Report 3044805.3044997, *Journal of Machine Learning Research*, 2014.
- [20] Q. Guo, J. Hang, S. Wang, W. Hui, and Z. Xie. Design optimization of variable stiffness composites by using multi-fidelity surrogate models. *Structural and Multidisciplinary Optimization*, 63(1):439–461, 2021.
- [21] Y. Ha and J. Mueller. Multi-Fidelity Stochastic Trust Region Method with Adaptive Sampling. Technical Report 2508.03901, arXiv, 2025.
- [22] H. Khatouri, T. Benamara, P. Breitkopf, J. Demange, and P. Feliot. Constrained multi-fidelity surrogate framework using Bayesian optimization with non-intrusive reduced-order basis. *Advanced Modeling and Simulation in Engineering Sciences*, 7(43), 2020.
- [23] D. Komljenovic, D. Messaoudi, A. Côté, M. Gaha, L. Vouliagny, S. Alarie, A. Dems, and O. Blancke. Asset Management in Electrical Utilities in the Context of Business and Operational Complexity. In 14th WCEAM Proceedings, pages 34–45, 2021.
- [24] D. Lakhmiri and S. Le Digabel. Use of static surrogates in hyperparameter optimization. *Operations Research Forum*, 3(11), 2022.
- [25] S. Le Digabel and S.M. Wild. A taxonomy of constraints in black-box simulation-based optimization. *Optimization and Engineering*, 25(2):1125–1143, 2024.
- [26] K. Li and F. Li. Multi-Fidelity Methods for Optimization: A Survey. Technical Report 2402.09638, arXiv, 2024.
- [27] A. March and K. Willcox. Constrained multifidelity optimization using model calibration. *Structural and Multidisciplinary Optimization*, 46(1):93–109, 2012.
- [28] J.J. Moré and S.M. Wild. Benchmarking Derivative-Free Optimization Algorithms. *SIAM Journal on Optimization*, 20(1):172–191, 2009.
- [29] L. Quan, H. Jiexiang, Z. Qi, S. Leshi, and Z. Anfu. A Multi-Fidelity Bayesian Optimization Approach for Constrained Multi-Objective Optimization Problems. *Journal of Mechanical Design*, 146(7):071702, 2024.
- [30] S. Razavi, B. Tolson, L. Matott, N. Thomson, A. MacLean, and F. Seglenieks. Reducing the computational cost of automatic calibration through model preemption. *Water Resources Research*, 46(11), 2010.
- [31] B. Solnik, D. Golovin, G. Kochanski, J.E. Karro, S. Moitra, and D. Sculley. Bayesian Optimization for a Better Dessert. In Proceedings of the 2017 NIPS Workshop on Bayesian Optimization, Long Beach, USA, 2017.
- [32] A. Tran, T. Wildey, and S. McCann. sMF-BO-2CoGP: A Sequential Multi-Fidelity Constrained Bayesian Optimization Framework for Design Applications. *Journal of Computing and Information Science in Engineering*, 20(3):031007–1–031007–11, 2020.
- [33] J. Wang, N.Y. Chiang, T. Hartland, J.L. Peterson, J. Solberg, and C.G. Petra. A Constrained Multi-Fidelity Bayesian Optimization Method. Technical Report 2508.03901, arXiv, 2025.